

Quadratic Vote Buying*

E. Glen Weyl[†]

April 2013

Abstract

A group of individuals with access to transfers seeks to make a binary collective decision. All known mechanisms they might use are either are often inefficient (e.g. voting), subject to severe collusion problems (e.g. the Vickrey-Clarke-Groves mechanism) or require the planner being informed about the distribution of valuations (e.g. the Expected Externality mechanism). I propose a simple, budget-balanced mechanism inspired by the work of Hylland and Zeckhauser (1979). Individuals purchase votes with the cost of a *marginal* vote being linear in the number of votes purchased; thus the total cost of votes is quadratic in the number purchased. The revenues earned from that individual are then refunded to other individuals.

When there are a large number of individuals, this *Quadratic Vote Buying* mechanism is efficient in any Bayesian equilibrium under symmetric independent private values and is usually nearly efficient even with aggregate uncertainty. Collusion by a small group or individuals' taking on (a small number of) multiple identities does not significantly reduce efficiency.

Nota bene: Many of the results in this paper are conjectural, and full formal proofs of many results are missing. They will be added, hopefully, by mid-summer.

*I am grateful to David Ahn, Eduardo Azevedo, Eric Budish, Drew Fudenberg, Jacob Goeree, Scott Duke Kominers, Steven Levitt, Paul Milgrom, Stephen Morris, Michael Ostrovsky, Azeem Shaikh and especially Roger Myerson, as well as seminar participants at the University of Chicago Booth School of Business, for helpful comments and for the financial support of the Institut D'Économie Industrielle and the Social Sciences Division at the University of Chicago. Tim Rudnicki, Daichi Ueda and especially Joshua Bosshardt supplied excellent research assistance. All errors, and I expect there are many in this draft, are my own.

[†]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

As Churchill's quote emphasizes, existing mechanisms for collective decision making are unsatisfying. Voting, while robust in many dimensions (Bierbrauer and Hellwig, 2011), does not incorporate intensity of preferences and thus is often Pareto-inefficient. Vickrey (1961), Clarke (1971) and Groves (1973) (VCG) proposed a mechanism that has a unique equilibrium in weakly dominant strategies that implements the efficient choice, but is highly sensitive to even two individuals colluding (even in equilibrium) and often requires destroying resources (Ausubel and Milgrom, 2005). All other mechanisms I am aware of either require the administrator to have detailed information on the distribution of valuations or have no reasonably general benefits over voting or the VCG scheme. This paper proposes a novel mechanism for binary collective decision making with transfers that is approximately efficient in any Bayesian equilibrium with a large number of individuals that have symmetric independent private values, is budget-balanced, puts strong limits on the potential benefits from collusion by a small group and has no information requirements for the planner.

This *Quadratic Vote Buying* (QVB) mechanism, partly inspired by the under-appreciated and unpublished work of Hylland and Zeckhauser (1979), is simple both to describe and implement. Each individual may purchase any continuous number of votes for her favored outcome at a cost of the square of that number of votes. Whichever alternative has more votes in its favor is selected, and the expenditures that individual is returned to all individuals other than her according to some pre-specified rule, such as even division, ensuring budget balance.¹

Why is QVB (limit-)efficient and why the quadratic rule exactly? The derivative of a quadratic function is linear, so under QVB the the marginal cost of an additional vote is proportional to the number of votes already purchased. The marginal benefit an individual derives from buying an additional vote is proportional to her cardinal value from changing the outcome as well as the chance (density) that this marginal vote is pivotal in changing the outcome. If this density of pivotality is common across individuals and insensitive to the number of votes purchased, as I show it must be in equilibrium in a large population, this leads individuals to buy a number of votes proportional to their value, linearly aggregating preferences and leading to efficiency. As I discuss more extensively in Subsection 6.1 and 6.2, any such (detail-free, scale-invariant, symmetric and separable) rule that is not quadratic (at least near the origin) would not aggregate value linearly and thus would not achieve efficiency robustly in the limit as does QVB.

Why is QVB resilient against collusion? Collusion cannot achieve much: the best any group of individuals can do is pool all their values and act as if they were a single individual, behaving

¹In fact under equal division the individual may receive back her own revenue as this only changes the arbitrary multiplicative constant on the cost of votes.

optimally for that individual and economizing the convexity of the quadratic vote cost function. This leads (approximately in large markets) to an aggregate outcome differing from that absent collusion only by multiplying the influence of the group members by the group's size. So long as the group is small compared to the total population, this effect is relatively small. Furthermore, because of the quadratic nature of costs, individual incentives to deviate from a collusive arrangement are proportional to the deviation collusion demands of individuals as compared to unilateral behavior, and the most effective forms of collusion, if anticipated, also create incentives for other individuals to increase their vote purchases. Thus, similar to collusion in oligopoly competition, effective (and thus harmful) collusion among a large group of individuals creates strong unilateral incentives for deviation. This contrasts sharply with VCG where any two individuals have an *equilibrium* in which they achieve their desired outcome and pay nothing. For similar reasons, the amount of harm an individual can cause by “de-merging” and pretending to be a small number of individuals is small under QVB but unlimited under VCG.

While I show that QVB is in some sense uniquely limit-efficient, my goal is *not*, as in many mechanism design papers, to argue that QVB is exactly optimal or perfectly robust under a well-defined but relatively narrow set of assumptions or criteria. Instead it is to persuade you that, unlike other proposals economists have made (Subsection 6.6), QVB has the robustness and simplicity necessary to be a practical alternative to voting that is superior in many settings. As with all standard analysis, tractability and brevity require me to focus on stylized cases that miss many realistic features. However, by considering many dimensions (de-mergers, collusion, budget balance, simplicity, etc.) and different environments (i.i.d. values and a setting with aggregate uncertainty) I aim to “stress test” QVB sufficiently to demonstrate its robustness.

By its nature, any mechanism with such robustness is unlikely to be perfectly optimized to any finite set of criteria and settings, just as for any finite data set a flexible class of curves can be over-fitted better than can a linear regression even if the true model is linear. For example, under aggregate uncertainty, QVB is not typically perfectly limit-efficient and sometimes even underperforms voting, but appears in a range of special cases I have been able to analyze to always be nearly limit-efficient. By contrast, voting is often highly inefficient even in the limit. Furthermore, I cannot guarantee QVB's performance is robust to a broader range of realistic features (in terms of e.g. interdependent values, higher-order beliefs or other potential manipulations) that I cannot hope to analyze in a single paper. I must hope that Occam's Razor, combined with the small range of relatively standard analyses that are within my capacity to convey here, under which QVB performs both reasonably well and is more robust than all other alternatives thus far proposed, will be sufficient to persuade you that QVB is a promising mechanism worth of further investigation.

Closely-related work

QVB is partly inspired by the work of Hylland and Zeckhauser (1979) who propose allowing individuals to allocate a given set of points across a variety of public goods, exerting influence on each public good equal to the square root of their expenditure on that good. Hylland and Zeckhauser consider an iterative “price-taking” process analogous to Tâtonement and argue, in a way that inspired my heuristic analysis in Section 2, that any equilibrium in this process must be efficient. While I consider a very different setting (an explicit game theoretic model and analysis with incomplete information, a binary rather than continuous decision, etc.) and thus obtain different results, QVB is in an important sense an application of their insight.²

Additionally, I have become aware of a related paper by Goeree and Zhang (2012), which was posted independently several months after my first draft. Goeree and Zhang show that when the mean of the distribution of valuations for changing the outcome, μ , is 0 the limit of the Expected Externality mechanism of Arrow (1979) and d’Aspremont and Gérard-Varet (1979) calls for payments quadratic in revealed type and inversely proportional to the number of individuals in the population and the standard deviation of their valuations. This insight is related to that underlying QVB in ways I detail in Subsection 6.5. However, because their mechanism is detail-dependent and applies only when $\mu = 0$ is commonly known, unlike QVB, it is subject to the same critiques for mechanisms economists have proposed for collective decision-making that I discuss above. Partly as a result, and also because their paper is primarily experimental, there is little overlap between the analysis in the two papers.³

Outline of paper

Following this introduction, the paper is divided into nine sections. Section 1 describes the model environment and formalizes the QVB mechanism. Section 2 assumes there is no aggregate uncertainty and separately analyzes three cases in which the limiting behavior of the equilibria of the mechanism is qualitatively different. When valuations are independently and identically distributed and $\mu = 0$, inefficiency dies off at a rate $\frac{1}{n}$ in any equilibrium. When μ is a constant different from 0, the rate at which inefficiency dies off depends on the tail index of the value distribution, but if

²The Hylland and Zeckhauser mechanism was recently revived by Benjamin et al. (2013) in the context of aggregating subjective well-being surveys, though the authors do not appear to be aware of Hylland and Zeckhauser’s work (nor of mine).

³First, Goeree and Zhang take a revelation principle-based approach to analyzing their mechanism, focusing on showing that as the population grows large there exists an approximate equilibrium in which individuals reveal their cardinal value. By contrast I study the (potentially large) *full set of* Bayesian equilibria of a mechanism in which individuals typically do not have an incentive to report their type. Second, Goeree and Zhang study only the limiting behavior of the mechanism (and a single example with a population size of 2) and not the behavior for large, but finite, populations, which is the focus of my formal results. I analyze not only the (zero-measure) case of $\mu = 0$ but also cases where μ takes on arbitrary values. While both my mechanism and theirs were inspired by the limit of the Expected Externality mechanism in the $\mu = 0$ case, in the $\mu \neq 0$ case neither mechanism is the large-population limit of the Expected Externality mechanism.

the distribution is thin-tailed it dies off at a rate $\frac{1}{n}$. When μ itself shrinks with $\frac{1}{\sqrt{n}}$, inefficiency does not vanish but instead asymptotes to a level that never exceeds 5% of the efficiency of the first-best.

Section 3 merges these results into the headline, take-away results of the paper; those not interested in the details of equilibrium may wish to skip directly to this section. There I consider the case in which μ is drawn randomly from a well-behaved distribution representing an “average case” analysis of welfare. Inefficiency in that case dies off at the rate $\frac{1}{n^{\frac{\bar{\alpha}-1}{\bar{\alpha}+1}}}$, where $\bar{\alpha}$ is the largest Pareto tail index of the distribution.

Section 4 departs from my emphasis on non-cooperative equilibrium and studies the incentives for and effects of collusion under QVB. I consider the effects of collusion under a range of settings where the collusive group is drawn either at random or systematically from the most extreme individuals and when $\mu = 0$ or $\mu \neq 0$. Collusion by extremists poses the greatest challenges to efficiency. In the “average case” collusion is both ineffective and very hard to enforce given unilateral deviation incentives in moderately large populations, while even in the “worst case” the incentives created by anticipated collusion for other individuals to increase their offsetting vote purchases are important to keep collusion from damaging efficiency. I also show that de-mergers, individuals pretending to have multiple identities, do not pose as large of a problem as does collusion and is unlikely to be a major concern for QVB in practice.

Section 5 discusses the case when the distribution of valuations is not commonly known. In this case, analysis and examples led me to conjecture that under reasonable conditions, QVB is efficient in the limit if and only if the ex-ante expectation of value is 0. On the other hand, if there is an “expected favorite”, QVB is biased against it and in favor of the “underdog” so that the underdog wins whenever it is efficient for her to do so but also wins in some cases where the favorite should. Voting may thus be more efficient than QVB in some cases. However, some examples I consider suggest that the potential and typical inefficiency of voting is much larger than that of QVB, which appears to be tightly bounded.

Section 6 discusses the relationship between QVB and a range of other mechanisms. A broader class of “convex power vote buying” mechanisms where the cost of votes is $|v|^x$ for $x > 2$ is shown to nest democracy for $x \rightarrow \infty$ and dictatorship for $x \rightarrow 1$. I show that among rules that are symmetric, detail-free and have transfers that are scale-invariant and separable between an individual’s own action and those of other individuals, all in senses formalized in Subsection 6.2, only rules that are approximately quadratic near 0 votes (of which QVB is arguably the simplest rule) are robustly limit-efficient. I argue in Subsection 6.3 that while QVB does not dominate voting along all dimensions, it is often a more attractive and equally practical alternative. In Subsections 6.4 and 6.5 I argue that while QVB is not as “clean” as VCG, the Expected Externality mechanism and the Goeree and Zhang (2012) mechanism along many dimensions, it avoids the crippling defects that limit the practical applicability of these mechanisms. I then argue in Subsection 6.6 that

several other theoretical alternatives to QVB are unlikely to be practically useful. Finally, I discuss in Subsection 6.7 analogs to QVB in practice and why (despite the fascinating analysis of Krishna and Morgan (2012) that helped inspire my analysis of aggregate uncertainty) costly and voluntary voting is unlikely to solve many of the inefficiencies of voting in practice.⁴

Section 7 discusses practical applications of QVB to committees, holdout problems, corporate governance and public choices. Section 8 discusses results I hope to derive in a future draft of this paper and in future, follow-on papers. Section 9 concludes with a discussion of potential potential avenue for future research by others. Extended calculations are collected into appendices following the main text. In the current draft, most of these arguments are sketches rather than full formal proofs. I hope by the end of the summer to have a version where all proofs are fully formal.

1 Model

n individuals must collectively decide whether to take an action A that impacts all of their utilities. Individual i receives a utility $u_i 1_A - t_i$, where 1_A is indicator for whether the action is undertaken, u_i is her *value* and t_i is the (net) *transfer* she makes. Individuals evaluate uncertain prospects as risk-neutral expected utility maximizers. Individuals know their own values, but not the values of other individuals. Individuals values are drawn independently and identically contingent on some aggregate statistic γ according to a continuously differentiable probability distribution functions $f(u|\gamma)$ with support on an open interval in \mathbb{R} . In what follows, with the exception of Section 5, γ is commonly known.

I assume that, given any γ , the first two moments of f exist and are $\mu(\gamma)$ and $\sigma^2(\gamma)$ respectively. When they exist, I denote the n th raw moment of f given γ as $\mu_n(\gamma)$. I drop all dependence on γ where it is not needed for clarity.

I define the *welfare* of a rule for choosing A as $(2 \cdot 1_A - 1) \sum_i u_i \equiv (2 \cdot 1_A - 1) U$ in a particular instance and its *expected welfare* as the expectation of this expression. The *inefficiency* of a rule is the amount by which the expected welfare falls short of that of the first-best rule $1_{U>0}$ and the *relative efficiency* of a rule is the ratio of its expected welfare to that of the first best.

Definition 1. *The Quadratic Vote Buying (QVB) mechanism with shares $\{s_i\}_{i=1}^n$ has $\sum_i s_i = 1$ and allows each individual to choose how many (positive or negative) votes v_i to purchase among any real number. It sets $1_A = 1_{\sum_i v_i > 0}$ and $t_i = v_i^2 - s_i \sum_{j \neq i} \frac{v_j^2}{1-s_j}$.*

Note that QVB is trivially budget balanced as

⁴While the Krishna and Morgan (2012) analysis directly influenced my analysis under aggregate uncertainty, I obtain results in that case that are inconsistent with their interpretation, though consistent with the letter, of their results. Additionally my mechanism and my analysis under aggregate certainty is independent of theirs and the first draft of my paper was posted several months before the first draft of their paper was.

$$\sum_i t_i = \sum_i v_i^2 - s_i \sum_{j \neq i} \frac{v_j^2}{1 - s_j} = \sum_i v_i^2 - \sum_{j \neq i} \frac{s_j}{1 - s_i} v_i^2 = 0.$$

2 Analysis under Aggregate Certainty

A simple way to analyze the equilibrium properties of QVB, which parallels the first-order analysis of Hylland and Zeckhauser (1979) for continuous public goods, is to consider it from the perspective of any individual i 's optimization. Let G_{-i} be the cumulative distribution function of equilibrium values of $V_{-i} \equiv \sum_{j \neq i} v_j$, with corresponding probability density function g_{-i} . Because values are drawn conditionally i.i.d., V_{-i} is independent of u_i . If individual i purchases votes v_i the probability of the action occurring is $1 - G_{-i}(-v_i)$. Thus individual i 's ex-ante expected payoff is, up to the receipts from others vote purchases which she cannot influence,

$$u_i [1 - G_{-i}(-v_i)] - v_i^2.$$

The first-order condition for maximization is then

$$u_i g_{-i}(-v_i) - 2v_i = 0 \iff v_i = \frac{u_i g_{-i}(-v_i)}{2}. \quad (1)$$

Suppose that g_{-i} is (approximately) independent of i and v_i , at least over the “relevant” range of v_i .⁵ Then $v_i \approx k u_i$ for some $k > 0$, and any equilibrium would be efficient as the outcome is determined by the sign of $\sum_i v_i$, which is the same as the sign of $\sum_i u_i$. Intuitively this would seem to be the case when the number of individuals is large, as each individual is a small part of a population, and the event of a tie is an aggregate event. The analysis in this section formalizes this strategy in three distinct contexts where γ is commonly known, in all of which I assume that the only moment of the distribution that depends on γ or n is μ : the case in which $\mu = 0, \frac{\hat{\mu}}{\sqrt{n}}$, where $\hat{\mu}$ is a constant, and the case when μ is a non-zero constant.

2.1 $\mu = 0$

When $\mu = 0$, as also noted by Goeree and Zhang (2012), the limit of the Expected Externality mechanism is a payment scheme that is quadratic in the reported value. The logic follows that applied by Morris and Shin (2003) closely: the normal distribution (to which the sum of the values of all but one individual converges by the central limit theorem) is well-approximated about its peak by a uniform distribution when the number of individuals, and thus the variance of the distribution, grows large. If such a linear-in-value equilibrium exists for one such quadratic payment

⁵The usefulness of the commonality of probability assessments was first observed by Thompson (1966). See Subsection 6.6 for more on his mechanism.

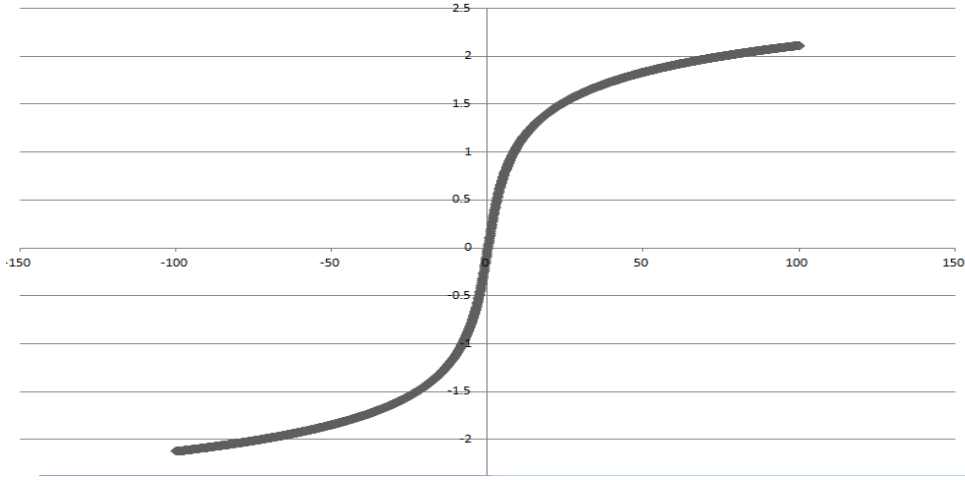


Figure 1: Exact optimal vote purchases when g_{-i} is distributed $N(0,1)$.

scheme, it should exist for any, as the multiplier applied to value to generate votes may be “un-done” by individuals rescaling the number of votes they purchase linearly. Thus one should expect an asymptotically linear and thus efficient equilibrium to exist. The following proposition shows something stronger: that *any* equilibrium of the mechanism must be of this form and that under such an equilibrium inefficiency dies off inversely with the number of individuals.

Lemma 1. *In any Bayesian equilibrium of QVB, if $\mu = 0$ the relative efficiency is $\geq 1 - \frac{\zeta}{n}$ where ζ is a constant in n that will be calculated in a future draft and will likely depend on smoothness properties of the distribution of valuations.*

Proof. Will appear in next draft. □

Here I provide a brief sketch of the proof, which is not yet available. Suppose that every individual uses a strategy that has a linear and a cubic term in their value $v = a(n)u + b(n)u^2 + c(n)u^3$; clearly no constant term appears, as it is a dominant strategy for an individual with value 0 to chose $v = 0$. To motivate this approximation, note that as long as all individuals use the same mapping between values and votes, g_{-i} is approximately normal in a large population given that values are drawn independently and identically. It furthermore seems reasonable to conjecture that when $\mu = 0$ the mean of the distribution of votes is near 0 (given that $v = 0$ is a dominant strategy for a 0 value type). We can thus obtain a reasonable sense for equilibrium behavior with a large number of individuals by solving for exactly optimal voting behavior when g_{-i} is normal.

Figure 1 plots this exact optimal strategy when g_{-i} follows a normal distribution with mean 0 and standard deviation 1. Note that optimal votes are approximately linear for values that are near 0 (roughly from -5 to 5 or five standard deviations of g_{-i} in each direction) and exactly optimal votes level off in magnitude for extreme utilities so that vote distributions will be well-bounded, ensuring the normal approximation will work well. Thus a third-order Taylor approximation is

likely to be quite accurate. It can easily be shown that as the standard deviation of the normal g_{-i} grows, as one would expect to occur, and I show occurs, as the population size grows, this bid function flattens and becomes closer to linear for a wider range of values. These observations are confirmed by the approximate equilibrium logic below. Similar experimentation, omitted for the sake of brevity, with exactly optimal voting functions when facing normal g_{-i} with different means motivates the approximations I use in the next two subsections.

Under the third-order approximation, the distribution g_{-i} , which from now on I abbreviate as g , therefore converges by the central limit theorem to a normal with mean $(n-1)[b(n)\sigma^2 + c(n)\mu_3]$ and standard deviation $a(n)\sqrt{n-1}\sigma$; I neglect the contribution to the standard deviation of the quadratic and cubic terms as they vanish in the limit given the intuitive logic above and as formally shown below. For each individual to be optimizing they must set

$$v = \frac{ug(-v)}{2} = u \frac{g(-[a(n)u + b(n)u^2 + c(n)u^3])}{2} \quad (2)$$

If $\lim_{n \rightarrow \infty} \frac{b(n)\sqrt{n-1}}{a(n)} = \lim_{n \rightarrow \infty} \frac{c(n)\sqrt{n-1}}{a(n)} = 0$, then the variance of g grows large relative the square of its mean and thus, by the same argument as in Morris and Shin (2003), g is well-approximated by a 2nd-order Taylor expansion about 0. That Taylor expansion yields

$$g(-v) = \frac{\frac{(n-1)^2 [b(n)\sigma^2 + c(n)\mu_3]^2}{2a^2(n)(n-1)\sigma^2}}{a(n)\sigma\sqrt{2\pi(n-1)}} \left[1 - \frac{(n-1)[b(n)\sigma^2 + c(n)\mu_3]}{a^2(n)(n-1)\sigma^2} v - \left(\frac{1}{a^2(n)(n-1)\sigma^2} - \frac{(n-1)^2 [b(n)\sigma^2 + c(n)\mu_3]^2}{a^4(n)(n-1)^2\sigma^4} \right) v^2 \right] \approx$$

$$\frac{1}{a(n)\sigma\sqrt{2\pi(n-1)}} \left[1 - \frac{b(n)\sigma^2 + c(n)\mu_3}{a^2(n)\sigma^2} v - \frac{1}{a^2(n)(n-1)\sigma^2} v^2 \right],$$

where the approximation allows terms that die, absolutely or relative to other terms multiplying the same power of v , to vanish. Now, plugging in the conjectured equilibrium voting behavior I drop all but the linear term as the others, again, vanish for large n to obtain

$$g \approx \frac{1}{a(n)\sigma\sqrt{2\pi(n-1)}} \left[1 - \frac{b(n)\sigma^2 + c(n)\mu_3}{a(n)\sigma^2} u - \frac{1}{(n-1)\sigma^2} u^2 \right],$$

so that by Equation (2),

$$v = \frac{1}{2a(n)\sigma\sqrt{2\pi(n-1)}} \left[u - \frac{b(n)\sigma^2 + c(n)\mu_3}{a(n)\sigma^2} u^2 - \frac{1}{(n-1)\sigma^2} u^3 \right].$$

I solve by matching coefficients:

$$a(n) = \frac{1}{2a(n)\sigma\sqrt{2\pi(n-1)}} \implies a(n) = \frac{1}{\sqrt{2}\sigma\sqrt[4]{2\pi(n-1)}}.$$

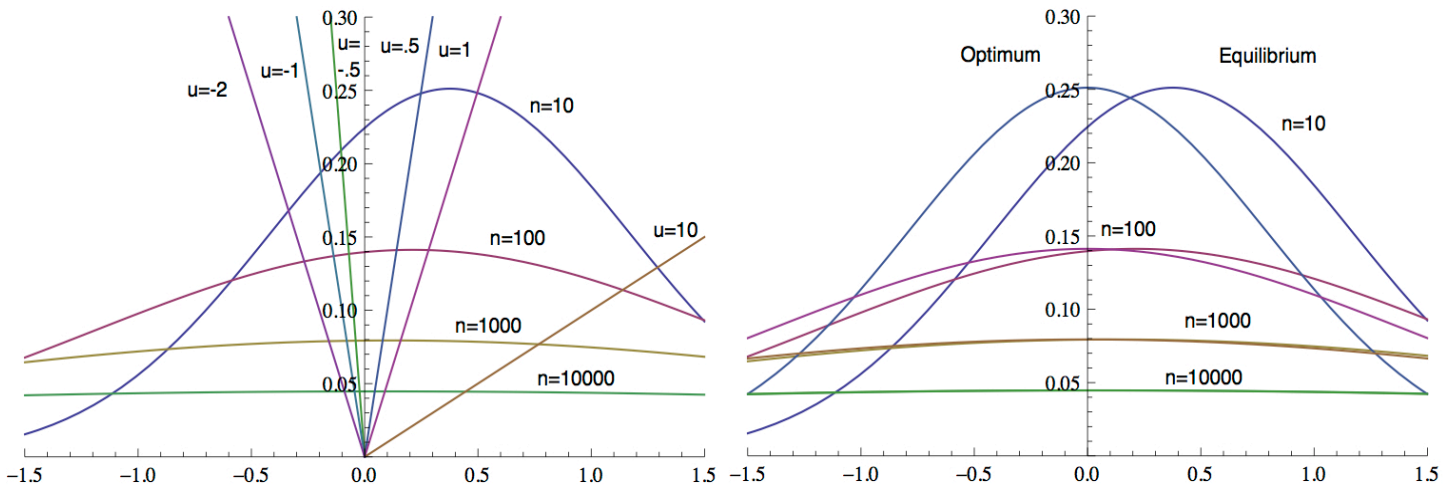


Figure 2: Approximate equilibrium voting when $\mu = 0, \sigma = 1$ and $\mu_3 = 1$. On the left curves represent $\frac{g(-v)}{2}$ for various values of n . The lines represent $\frac{v}{u}$ for various values of u . Intersections are optimal $\frac{v^*}{u}$ ratios. On the right pictured the distribution of votes compared to the scaled-down distribution of values. Limit efficiency can be seen by the convergence of these distributions.

$$c(n) = -\frac{a(n)}{(n-1)\sigma^2} = -\frac{1}{(n-1)^{\frac{5}{4}}\sigma^{\frac{5}{2}}\sqrt{2}\sqrt[4]{2\pi}}.$$

$$b(n) = -\frac{b(n)\sigma^2 + c(n)\mu_3}{\sigma^2} \implies b(n) = -\frac{c(n)\mu_3}{2} = \frac{\mu_3}{(n-1)^{\frac{5}{4}}\sigma^{\frac{5}{2}}\sqrt[4]{\pi}2^{\frac{7}{4}}}.$$

Note that the conjecture $\lim_{n \rightarrow \infty} \frac{b(n)\sqrt{n-1}}{a(n)} = \lim_{n \rightarrow \infty} \frac{c(n)\sqrt{n-1}}{a(n)} = 0$ is amply confirmed. Thus the equilibrium structure desired obtains: at a $\frac{1}{n}$ rate equilibrium vote-buying converges to a linear function of value with coefficient $\frac{1}{\sqrt{2\sigma}\sqrt[4]{2\pi(n-1)}}$.

Figure 2 illustrates the structure of equilibrium and its consequences graphically. It shows $\frac{g(-v)}{2}$ for various values of n in the case when $\mu = 0, \sigma = 1$ and $\mu_3 = 1$ as well as $\frac{v}{u}$ for various values of u . Intersections represent optimal $\frac{v^*}{u}$ ratios for different u types. Because $\mu_3 = 1$ the distribution of values is right-skewed and thus more extreme values of positive value are likely. As shown by $u = 10$ this leads to a lower $\frac{v^*}{u}$ ratio for some positive values, causing the mean of the vote distribution to be slightly negative, as pictured by the peak of the normal distributions being slightly to the right of the vertical axis as it is $\frac{g(-v)}{2}$ not $\frac{g(v)}{2}$ that is pictured. This, in turn, is partial offset by the fact that, for a given value magnitude, $\frac{v^*}{u}$ is larger for positive than negative u given this bias, as positive u 's move towards the peak of the bell and negative u 's move away from its peak. All of these deviations from constant $\frac{v^*}{u}$, however, vanish rapidly as n increases as conjectured. This is shown on the right, where the distributions of votes and of values scaled-down by $\frac{g(0)}{2}$ are shown. As n increases these converge to one another, ensuring limit efficiency. The formal proof will add uniform convergence results to normal distributions to ensure that the result applies in finite samples and will show that any equilibrium must have this structure. It remains only to determine the rate at

which inefficiency dies off given these results.

The welfare of the first best is $E[|U|] \approx \sqrt{\frac{2n}{\pi}}\sigma$ by the central limit theorem. Letting $V \equiv \sum_i v_i$, $U_2 \equiv \sum_i u_i^2$, etc., welfare at equilibrium is

$$E[(2 \cdot 1_{V>0} - 1)U] = E[(2 \cdot 1_{a(n)U+b(n)U_2+c(n)U_3>0} - 1)U] = E\left[\left(2 \cdot 1_{U+\frac{b(n)}{a(n)}U_2+\frac{c(n)}{a(n)}U_3>0} - 1\right)U\right].$$

To calculate this, note that U, U_2 and U_3 are, for large n , jointly normally distributed according to

$$\frac{1}{n} \begin{pmatrix} U \\ U_2 \\ U_3 \end{pmatrix} \sim N \left(\begin{matrix} 0 \\ \sigma^2 \\ \mu_3 \end{matrix}, \frac{1}{n} \begin{bmatrix} \sigma^2 & \mu_3 & \mu_4 \\ \mu_3 & \mu_4 - \sigma^4 & \mu_5 - \sigma^2\mu_3 \\ \mu_4 & \mu_5 - \sigma^2\mu_3 & \mu_6 - \mu_3^2 \end{bmatrix} \right).$$

$\frac{c(n)}{a(n)} = -\frac{1}{(n-1)\sigma^2}$ and $\frac{b(n)}{a(n)} = \frac{\mu_3}{2\sigma^2(n-1)}$. Thus the asymptotic joint distribution of U and $U + \frac{b(n)}{a(n)}U_2 + \frac{c(n)}{a(n)}U_3$ is

$$\frac{1}{n} \begin{pmatrix} U \\ U + \frac{b(n)}{a(n)}U_2 + \frac{c(n)}{a(n)}U_3 \end{pmatrix} \sim N \left(\begin{matrix} 0 \\ -\frac{\mu_3}{2(n-1)} \end{matrix}, \frac{1}{n} \begin{bmatrix} \sigma^2 & \sigma^2 \\ \sigma^2 & \sigma^2 \end{bmatrix} \right),$$

where, as before, I drop asymptotically small terms in any expression. Thus perfect efficiency is achieved whenever U has the same sign as μ_3 because the expressions are perfectly correlated. Without loss of generality, assume $\mu_3 > 0$. Then when $\frac{U}{n} \in \left(0, \frac{\mu_3}{2(n-1)}\right)$ inefficiency will result. The magnitude of this inefficiency is

$$\int_0^{\frac{n\mu_3}{2(n-1)}} \frac{-Ue^{-\frac{U^2}{2n\sigma^2}}}{\sqrt{2\pi}\sqrt{n}\sigma} dU \rightarrow \int_0^{\frac{\mu_3}{2}} \frac{-U}{\sqrt{2\pi}\sqrt{n}\sigma} dU = \frac{\mu_3^2}{8\sqrt{2\pi}\sqrt{n}\sigma}.$$

This is of order $\frac{1}{\sqrt{n}}$ and thus inefficiency (lost welfare relative to the first best) dies off at the rate $\frac{1}{n}$.⁶

2.2 $\mu = \frac{\hat{\mu}}{\sqrt{n}}$

In the case when $\mu = \frac{\hat{\mu}}{\sqrt{n}}$, an approximately linear equilibrium like the one derived in the previous subsection leads to mean values of the of sum of all votes that are close enough to 0 that the chance of a tie does not vanish exponentially, but far enough from 0 so that the second-order term in the strategy expansion is not trivial and may cause inefficiency that does not vanish.

⁶Note the limiting constant also follows from this expression, but I omitted it above because I am interested in a larger constant that provides a finite sample bound derivable from versions of the Berry (1941)-Esseen (1942, 1956) Theorem that I plan to use in the revision of this paper.

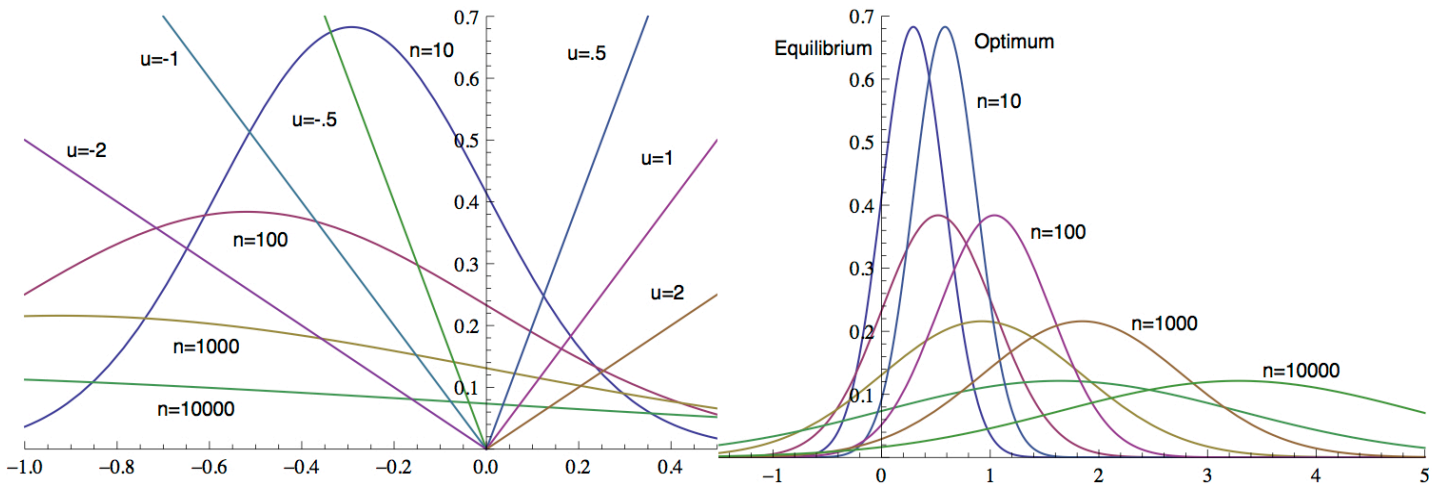


Figure 3: Approximate equilibrium voting when $\mu = \frac{1}{\sqrt{n}}$ and $\sigma = 1$. On the left the curves represent $\frac{g(-v)}{2}$ for various values of n . The lines represent $\frac{v}{u}$ for various values of u . Intersections are optimal $\frac{v^*}{u}$ ratios. On the right is pictured the distribution of votes compared to the scaled-down distribution of values. Limit inefficiency arises from the divergence of the two distribution pairs.

Lemma 2. *In any Bayesian equilibrium of QVB, if $\mu = \frac{\hat{\mu}}{\sqrt{n}}$ where $\hat{\mu} \neq 0$, the efficiency is*

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

where Φ is the cumulative distribution function of the standard normal distribution, ζ is a constant to be calculated in a future draft and $x = \frac{\hat{\mu}}{\sigma}$. If the x -dependent term is maximized over all values of x , it takes on a value $< .045$, implying that if $n > \underline{n}$, another constant yet-to-be-calculated, then the inefficiency ratio in this case is $> .95$. For x that is either very large or small in absolute value, the expression depending on x vanishes.

Whenever the decision should be made in the opposite direction from $\hat{\mu}$ (e.g. against A when $\hat{\mu} > 0$) it is made correctly. All inefficiency arises from cases when the decision should be made in favor of the direction of $\hat{\mu}$ but instead goes against it. That is, the ex-ante “favorite” decision does not receive the full advantage it should and sometimes the “underdog” is inefficiently allowed to win. This is a version of the Taylor and Yildirim (2010) Underdog Principle that will also cause inefficiency in QVB in Section 5.

Proof. Will appear in next draft. See Appendix A for a sketch. \square

Intuitively, if in the limit everyone buys votes proportional to their value, the ex-ante favored side is expected to be in the lead. WLOG suppose $\hat{\mu} > 0$. Individuals buying negative votes will increase their density of pivotality the more negative votes they buy, as they will move closer to the

peak of the normal distribution that g limits to, while individuals buying positive votes will see their density of pivotality decline as they move down the side of the normal distribution. This depresses the chances of the ex-ante favorite winning. While this effect becomes small with n , it does so at the same rate as μ shrinks with n and thus creates a non-negligible bias so that in cases when the ex-ante favorite should win, but not by much, it ends up losing creating a small but non-vanishing limit inefficiency

This equilibrium and its consequences are pictured in Figure 3 for the case of $\hat{\mu} = \sigma = 1$. On the left is a graph analogous to the left panel of Figure 2. The positive value of $\hat{\mu}$ shifts the mean of the bell curve to the left of the axis, leading negative values to be over-weighted by the underdog effect. As n grows this effect becomes smaller as the curve grows flatter, but note that the peak of the bell-curve remains separated from the origin (actually becomes more so) so that the rate at which $\frac{v^*}{u}$ converges across different u 's is much slower. This leads, as shown in the right panel, to a bias where the peak of the vote distribution is only half as separated from 0 as it should be if it were to represent the mean of the value distribution. The difference between these two curves around 0 (one being to the left of the y-axis when the other is not), which does not vanish as n grows large as in Figure 2, represents the inefficiency created by the Underdog Principle.

2.3 $\mu \neq 0$

While the efficiency results of Subsection 2.2 differ greatly from those of Subsection 2.1, the structure of equilibrium is basically similar. This is not the case when μ is a non-zero constant.

To see this suppose that an equilibrium where the linear term in the voting equation dominated existed. Then equilibrium with many individuals features a mean of votes that grows without bound relative to the variance of votes, and thus $g(0)$ dies off exponentially. Aggregate votes purchased would then die exponentially as well. This would make it attractive for contrarian individuals, those with a value that is the opposite sign to μ , to “buy the whole decision” by purchasing enough votes to guarantee the victory of their side. In fact, the behavior of such *extreme contrarians* (*extremist* for short) is crucial to the nature of the equilibrium with $\mu \neq 0$. In particular, the rate at which inefficiency dies off is determined by the “number” of such extreme contrarians as quantified by the tail index of the tail of the distribution with the opposite sign as μ .

Lemma 3. *In any Bayesian equilibrium of QVB in the case when μ is a non-zero constant, the efficiency is $\geq 1 - \frac{\zeta}{n^{\frac{\alpha-1}{\alpha+1}}}$, where α is the Pareto tail index of the distribution of valuations of sign opposite to μ and ζ is another constant to be determined in a future draft. If the distribution is thin or non-tailed, $\alpha \rightarrow \infty$ and efficiency is $\geq 1 - \frac{\zeta}{n}$; $\alpha > 2$ is imposed by my assumption of finite first and second moments so the rate of die-off is always at least $\frac{1}{\sqrt[3]{n}}$.*

Proof. Will appear in next draft. See Appendix A for a proof sketch. □

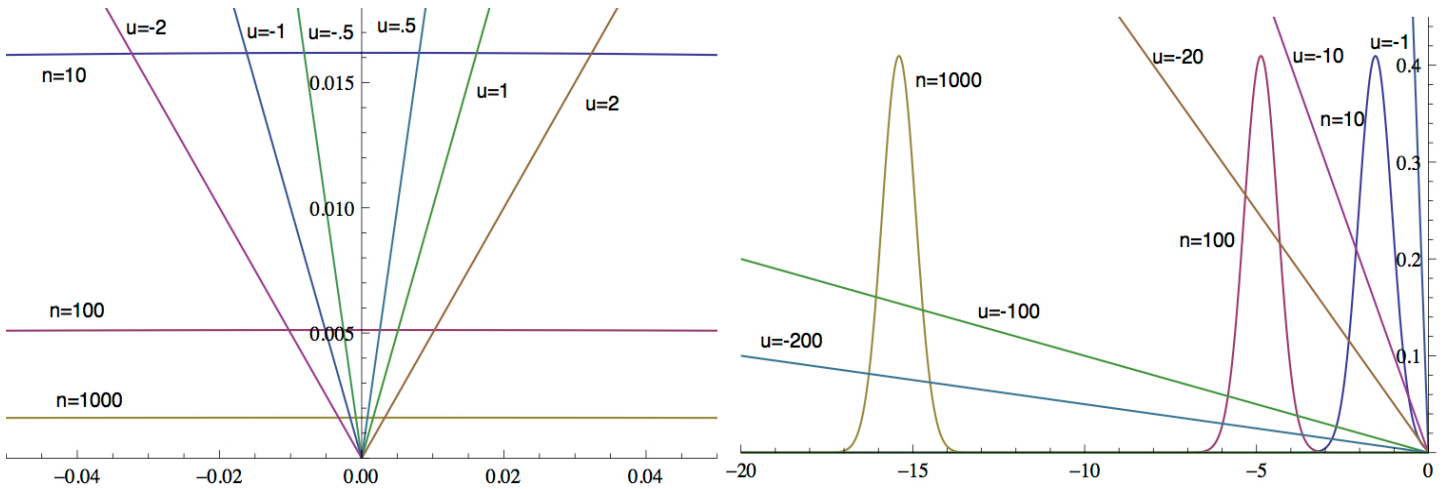


Figure 4: Approximate equilibrium voting when $k = \mu = \sigma = 1$ and $\alpha = 3$. On the left is shown the strategic calculus of the moderates, with $\frac{g(-v)}{2}$ being dominated by the event that an extremist exists. On the right is shown the strategic calculation of extremists. For them the event of no other extremist existing dominates and they choose an extremist strategy if the area above $\frac{g(-v)}{2}$ and below their $\frac{v}{u}$ is smaller than the area below $\frac{g(-v)}{2}$ and above $\frac{v}{u}$ prior to the second intersection of those curves.

The equilibrium features most individuals buying a number of votes that are proportional, in the limit, to their value, following what I dub a *moderate strategy*. However extremists follow an extreme version of the Underdog Principle. They buy a large number of votes sufficient to ensure that their preferred outcome is implemented with probability greater (though not too much greater) than $\frac{1}{2}$. In the case that no such extremist appears, the density of pivotality dies exponentially in n . Thus almost all of the incentive that moderates have to buy votes must arise from the event in which an extremist exists. The more likely is the appearance of an extremist, the more votes moderates buy. But conversely the more votes moderates buy, the more expensive and thus less attractive is extremism. This mixture of strategic substitutability and strategic complementarity between the probability of extremists and number of votes bought by moderates creates an equilibrium. In this equilibrium, both the number of votes bought by moderates and the chance that *any individual* is an extremist must die with n . The rate at which the chance of an extremist dies depends on whether there is a “deep reservoir” of individuals with increasingly extreme value; if not the probability must die quickly to ensure the number of votes bought by moderates dies quickly to maintain incentives for the extremists to be willing to be such. Thus the rate at which the probability of any extremist emerging dies with n increases with the Pareto tail index α . The larger α is, the thinner are tails and thus the less extreme values exist. All inefficiency is created by extremists as the probability of the right decision being opposite to μ goes exponentially to 0. Thus the inefficiency follows precisely the same rate and logic as the chance of an extremist existing.

The determination of this approximate equilibrium is shown in Figure 4 in the case when $\mu =$

$\sigma = k = 1$ and $\alpha = 3$.⁷ On the left are shown the strategies of the moderates, analogous to the left panels in the previous two figures. Note that $\frac{g(-v)}{2}$ is now extremely flat. The reason is that it is scaled-down by pn , the chance of an extremist existing, making the range of votes spanned by reasonable ranges of values very small. Thus moderates converge to proportional strategies extremely rapidly as shown.

However, matters are very different for extremists, as shown in the right panel. Individuals with small magnitudes of negative (contrarian) value do not even touch $\frac{g(-v)}{2}$ and thus will not pursue extremist strategies. On the other hand, those with large negative values do, especially for small n , and thus will consider extremist strategies. It is easy to determine whether such a strategy is optimal and the optimal number of votes to purchase conditional on pursuing extremism. Conditional on extremism, the optimal vote purchases are given by the third (farthest) intersection between the $\frac{v}{u}$ ray and the distribution. It is optimal to pursue such a strategy if the area up to this point above the ray and below the bell curve exceeds the area below the ray and above the bell curve's right tail. Thus an individual with $u = -20$ will pursue an extremist strategy when $n = 10$ but not when $n = 100$. As n grows clearly the set of types willing to pursue an extremist strategy, and thus the chance that any individual will exist who pursue such a strategy, shrinks leading the efficient outcome to be implemented with increasingly high probability.

3 Main Results

In the previous section I considered, separately, welfare calculus in each of three cases for the mean of the distribution of valuations. In practice one is likely interested in something closer to an average over these various cases. Suppose that μ is drawn randomly, and then revealed to all players, from some continuous distribution with support that includes 0. In the notation above this is the case when $\gamma = \mu$ and is commonly known. A natural question is what is the average efficiency of QVB over all the cases captured by this distribution for large n .

Theorem 1. *For any continuous distribution over $\gamma = \mu$, the average efficiency (fraction of average first-best welfare achieved) of QVB is, for large n , is $\geq 1 - \frac{A}{n}$ if the γ -conditional distribution of u is thin-tailed for all γ and is $\geq 1 - \frac{B}{n^{\frac{\bar{\alpha}-1}{\bar{\alpha}+1}}$ where $\bar{\alpha}$ is the supremum over γ of the tail index of the γ conditional distribution of u . A and B are constants to be determined in a future draft of this paper.*

Proof. Will appear in next draft. See Appendix A for more calculations related to this result. \square

First note that the total welfare possible is $E[|U|] \approx nE[|\mu|]$. The only case when welfare gains are not dominated by the mean is when μ is $\frac{1}{\sqrt{n}}$ close to 0. The total measure of these cases is on

⁷Approximations used in this equilibrium, and the meaning of k , are discussed and justified partially using this figure in Appendix A.

the order of $\frac{1}{\sqrt{n}}$, and the efficiency gains from tracking the sign U even when it differs from that of μ is on order \sqrt{n} from the above analysis. Thus the total gains from tracking the sign of U when it is not that of μ are of order 1, which does not change the limiting size $nE[|\mu|]$.

Now to consider the efficiency losses of QVB relative to the first best, first note that the losses in the $\mu = 0$ case are irrelevant because this case is of measure 0. The second sources of losses are in the cases when $\mu = \frac{\hat{\mu}}{\sqrt{n}}$. For large n and any fixed $\hat{\mu}$, these all occur for values of μ close to 0. The density of such an event is thus approximately $\frac{h(0)}{\sqrt{n}}$, where h is the density of $\gamma = \mu$. The efficiency losses in all such cases are of order \sqrt{n} and thus total losses resulting in the limit from such cases are on the order of 1. More detailed calculations are available in Appendix A.

Efficiency losses for the case when μ is separated definitely from 0 depend on the tail index of the distribution of extremist values as indicated in Subsection 2.3. If this has thick tails (finite α) then for large n the inefficiency from this case dominates. On the other hand, if this distribution has thin tails ($\alpha = \infty$) then inefficiency dies at a $\frac{1}{n}$ rate and thus this inefficiency contributes a similar amount to limiting inefficiency to that arising from the $\frac{\hat{\mu}}{\sqrt{n}}$ cases.

4 Collusion

In this section, I explore the incentives for and effects of collusion in QVB. Throughout, I study the incentives for collusive behavior given that all members outside of the collusive group play as in non-cooperative equilibrium and members within the collusive group may perfectly pool information, as in the analysis of Ausubel and Milgrom (2005), though I discuss incentives to unilaterally deviate from such a collusive group in Subsection 4.3 and the response of other individuals in Subsection 4.5.

Subsection 4.1 derives a heuristic framework for studying optimal collusion under QVB. Subsection 4.2 states formal results on the impact of collusion on efficiency and revenue, using the heuristic framework to sketch a proof of these results. Subsection 4.3 provides formal results characterizing the incentives of individuals to unilaterally deviate from collusive agreement and a heuristic proof of these results.

4.1 Optimal collusion

For expositional clarity, in this subsection I make a number of simplifying assumptions and approximations that are not necessary to establish the results of the following subsections. Suppose some subset M of all individuals join a coalition and coordinate their actions to maximize their joint value. Let G_{-M} denote the CDF of the distribution of the sum of all votes outside the coalition.

The total expected utility of the coalition is then

$$\sum_{i \in M} u_i \left[1 - G_{-M} \left(- \sum_{i \in M} v_i \right) \right] - \sum_{i \in M} t_i.$$

Let $x_M \equiv \sum_{i \in M} x_i$ for $x = u, v, t$. For expositional clarity, assume $s_i \equiv \frac{1}{n}$ and thus that $\frac{m-1}{n-1}$ of the revenue collected from the group is remitted back to it, where $m \equiv |M|$. The payoff of the group is given by

$$u_M [1 - G_{-M}(-v_M)] - \left(1 - \frac{m-1}{n-1} \right) (v^2)_M.$$

Notice that this expression depends only on v_M and $(v^2)_M$ and the second term is always negative assuming $m \neq n$. Thus an optimal allocation of votes among the individuals in the coalition always requires minimizing $(v^2)_M$ subject to a given level of v_M . Because the square is convex, this always requires setting the votes of all individual in the group to be the same. I can then rewrite the payoff as

$$u_M [1 - G_{-M}(-v_M)] - \left(1 - \frac{m-1}{n-1} \right) \frac{v_M^2}{m}.$$

The first-order condition is then

$$u_M g_{-M}(-v_M) - \frac{2(n-m)}{m(n-1)} v_M = 0 \iff v_M = \frac{m(n-1)}{2(n-m)} u_M g_{-M}(-v_M).$$

When n is large (both absolutely and relative to m) this is approximately $v_M = m \frac{u_M g_{-M}(-v_M)}{2}$. Section 2 discussed three cases and here I analyze the limiting optimal behavior of a collusive group in each of these cases, assuming n is large and $\gg m$.

Cases when $\mu = 0$ or $\mu = \frac{\hat{\mu}}{\sqrt{n}}$

First, suppose that, as in the $\mu = 0$ and $\mu = \frac{\hat{\mu}}{\sqrt{n}}$ case discussed in Subsections 2.1 and 2.2, g is close to constant at value g_0 for the range of v values chosen by the coalition. Then

$$v_M^* = m u_M \frac{g_0}{2} \approx m \sum_{i \in M} v_i^*,$$

where v_i^* are the unilaterally optimal vote choices of the members of the coalition. Thus the effect of the coalition is two-fold. First, it “evens out” the vote purchases of members so they are equal and thus minimize payments conditional on a level of aggregate, net votes purchased by the coalition. Second, it magnifies this aggregate, net level of votes by a factor m .

Cases when $\mu \neq 0$

In the analysis of Subsection 2.3, some individuals take moderate and some extreme strategies. The analysis of moderate strategies differs little from that when $\mu = 0$ or $\mu = \frac{\hat{\mu}}{\sqrt{n}}$. However matters are different when the collusive group is on average extremist and thus may adopt an extreme strategy. To analyze this case, I use the same approximating assumptions as in Subsection 2.3. The coalition will find it attractive to take the extremist strategy if

$$mu_M < -2a^2(n)n^{\frac{3}{2}}\sqrt{2\pi}\sigma\mu \iff u_M < -2\frac{k^{\frac{1}{1+\alpha}}}{mn^{\frac{3\alpha-1}{2(\alpha+1)}}2^{\frac{\alpha}{\alpha+1}}\sigma\sqrt{2\pi}}n^{\frac{3}{2}}\sqrt{2\pi}\sigma\mu = -\frac{2^{\frac{1}{1+\alpha}}\mu n^{\frac{2}{1+\alpha}}k^{\frac{1}{1+\alpha}}}{m}.$$

4.2 Efficiency cost of collusion

The cost to efficiency of collusion depends on who composes the colluding groups (are they a randomly selected group? Or are they systematically of a particular value type?), as well as on whether collusion aims to implement moderate strategies or extremist strategies. In particular, I consider worst case (when all colluders have large values in the same direction) and average case (when colluders have random values) expected efficiency losses from collusion when collusion may cause extremist strategies in the $\mu \neq 0$ case and when it is used to implement moderate strategies in the $\mu = 0$ case. I focus in the latter situation on the $\mu = 0$ case because moderate strategies cannot cause significant efficiency loss when $\mu \neq 0$, and the efficiency losses when $\mu = \frac{\hat{\mu}}{\sqrt{n}}$ are bounded in the same manner as in the $\mu = 0$ case.

Proposition 1. *In the average case, when collusive groups are composed of randomly selected individuals, the collusion of a single group when $\mu = 0$ causes the relative inefficiency to be $O\left(\frac{m^{\frac{3}{2}}}{\sqrt{n}}\right)$ and thus to decline in n so long as $m = O(\sqrt[3]{n})$. When $\mu \neq 0$ de-merger problems do not affect the rate at which efficiency is achieved as long as $m = O\left(n^{\frac{4}{3(1+\alpha)}}\right)$.*

In that worst case, when the collusive group is composed systematically of the most extreme individuals, if $\mu = 0$ then relative inefficiency is $O\left(m^{\frac{2\alpha-1}{\alpha}}n^{-\frac{2-\alpha}{\alpha}}\right)$ and thus dies for large n as long as $m = O\left(n^{\frac{\alpha-2}{2(2\alpha-1)}}\right)$. When $\mu \neq 0$ and the collusive group aims to implement an extremist strategy, relative inefficiency is $O\left(n^{-\frac{\alpha-1}{\alpha+1}}m^{2\alpha-1}\right)$ and dies off in n so long as $m = O\left(n^{\frac{\alpha-1}{(\alpha+1)(2\alpha-1)}}\right)$.

Proof. Will appear in next draft. See Appendix A for a proof sketch. □

A natural way to understand the strength or weakness of these limits on collusion's efficacy is to consider their implications for plausible parameter values. In a million person decision, $\sqrt[3]{n} = 100$. This is a large group in which to attempt to sustain collusion and thus, at least in quite large populations, it seems unlikely that average case collusion will pose a significant threat to efficiency. Similarly in the average case collusion to implement extremist outcomes is unlikely to be attractive. Consider a reasonable case when $\alpha = 3$, corresponding to the upper bound of standard tail estimates

for income distribution in the United States (Saez, 2001). Then again a group would have to be of order $\sqrt[3]{n}$ to impact efficiency.

On the other hand, in the worst case, when the collusive group is an extremist oligarchy who can identify one another as the most extreme individuals, collusion is likely to be much more damaging to efficiency. Then, in order for collusion's impact on efficiency to die with n , m must be $O\left(n^{\frac{1}{10}}\right)$ regardless of whether the collusive group aims for an extremist strategy or a local strategy. Even if n is in the billions (viz. the size of world population), this still is little comfort as $n^{\frac{1}{10}}$ is on the order of 1 or 10 at most. Thus even a small, effective and unanticipated oligarchic clique may significantly weaken the efficiency of QVB.

Note that the effects of collusion are most severe, in this case, when α is very close to small or (in the case of implementing extremist strategies) very large. As $\alpha \rightarrow 2$ in the case of moderate strategies or 1 in the case of extremist strategies, a constant-sized clique has a significant efficiency impact. In the case of extremist strategies as $\alpha \rightarrow \infty$ the same is true. However, unlike in VCG, for reasonable values of α a collusive group of very small, constant size does not significantly impact efficiency. For example, a collusive group of two individuals will only raise the probability of an extremist strategy being implemented by a factor of

$$\left[2\left(1 + \frac{1}{2^{\frac{1}{\alpha}}}\right)\right]^{\alpha} < 4^{\alpha},$$

which, for $\alpha = 3$, is only 64 while the order of the inefficiency surviving absent collusion is, with a million individuals, $n^{-\frac{\alpha-1}{\alpha+1}} = \frac{1}{\sqrt{n}} \approx \frac{1}{1000}$. So a very small collusive group of fixed size is unlikely to matter, unlike under VCG.

Results on the revenue consequences of collusion may also be derived, but are omitted here for brevity. They may appear in a future draft. Two things should be noted. The worst case for revenue is when $u_M = 0$ but the values of the group members are highly variable. In this case there is no inefficiency cost to collusion, only a revenue cost. On the other hand, collusion will actually increase revenues in the worst cases for efficiency discussed above.

4.3 Deviation incentives

The previous subsection described the limits under QVB on the efficacy of optimal collusion. Now I discuss the difficulties in achieving optimal collusion. In particular, collusion under QVB often creates large incentives for unilateral deviations by individuals. These unilateral deviation incentives may create the same sort of obstacles to collusion, as emphasized by Laffont and Martimort (1997, 2000), that they do in standard competitive settings such as product market competition or first-price auctions. In particular, if a secret ballot is used, it may be difficult for collusive groups to detect and thereby punish deviations from collusive arrangements. Unilateral deviation incentives

may therefore undermine the sustainability of collusion.

As above, the unilateral incentives created by collusion differ between the the case when the collusive group attempts to implement a moderate strategy and when they attempt to implement an extreme strategy. I consider each in turn. For moderate strategies, I assume that g is constant at g_0 . Every member of the collusive group buys vote $\frac{g_0 u_M}{2}$. When they do, individual i 's marginal incentive to move towards her unilateral optimal choice of $\frac{g_0 u_i}{2}$ is of size $\frac{g_0}{2} |u_i - u_M|$. The average across members of the group of these marginal incentives is $\frac{g_0}{2}$ times the mean absolute deviation of u_i from u_M . In the worst case in terms of efficiency, when all u_i 's are the same, this is $(m - 1)u_i$. These marginal incentives are emphasized by Erdil and Klemperer (2010) and clearly grow the more effective collusion is.

Similarly the gains from defection to the unilateral optimum are

$$\frac{u_i (u_i - u_M) g_0^2}{2} - \frac{g_0^2}{4} (u_i^2 - u_M^2) = \frac{g_0^2}{4} (u_i^2 - 2u_i u_M + u_M^2) = \frac{g_0^2}{4} (u_i - u_M)^2.$$

That is, they are proportional to the squared difference between u_i and u_M . The average size of this equals the variance in u_i plus $(m - 1)^2 u_M^2$. The first term, the variance, measures the gains the group gets from depriving the mechanism of revenue; the second term corresponds to my upper bound on the inefficiency of the collusion in the local deviation case. Thus the more revenue the collusive group “steals” and the more inefficiency it causes, the greater unilateral incentives it creates for deviation if it is pursuing a moderate collusive strategy. Thus QVB is *coalition-proof* in these settings in the sense of Laffont and Martimort (1997, 2000).

When the group pursues an extremist strategy things are somewhat less optimistic. The simplest way to see this is again to consider my earlier approximation that the collusive always chooses to buy just enough votes to tie the election in expectation. Assuming all other member of the group behave this way and that m is large enough so that even if an individual member of the group buys no votes the election will still be close-to-tied in expectation but that n is sufficiently larger than m so that m is a negligible fraction of n , the value of g facing that individual is approximately $\frac{1}{a(n)\sigma\sqrt{2\pi(n-1)}}$, while the number of votes needed to execute an extremist strategy successfully is approximately $na(n)\mu$ in absolute value. Unilaterally individual i will have an incentive to buy approximately $\frac{u_i}{2a(n)\sigma\sqrt{2\pi(n-1)}}$ votes. For individuals to unilaterally have incentives sufficient to support an extremist strategy, therefore, I would need

$$\frac{|u_M|}{2^{\frac{3}{4}} a(n)\sigma\sqrt{\pi(n-1)}} > na(n)\mu \implies |u_M| > 2^{\frac{3}{4}} a^2(n)n^{\frac{3}{2}}\sigma\sqrt{\pi} = \frac{k^{\frac{1}{\alpha+1}}}{2^{\frac{3\alpha-1}{4(\alpha+1)}}} n^{\frac{2}{\alpha+1}}.$$

This condition is precisely the same as that in Subsection 4.1 for the collusive group to find it optimal to follow an extremist strategy *except that the expression is not divided by m* . In the worst case when the members of the group have the most extreme values, this makes the size of the collusive group m

needed to sustain unilaterally incentive compatible collusion for an extremist strategy of order $n^{\frac{2}{\alpha+1}}$ while such a strategy is attractive when m is of order only $n^{\frac{1}{\alpha+1}}$. Repeating the analysis in Appendix A for this case shows that collusion that does not yield unilateral deviation incentives makes relative inefficiency $O\left(n^{-\frac{\alpha-1}{\alpha+1}}m^{\alpha-1}\right)$ rather than $O\left(n^{-\frac{\alpha-1}{\alpha+1}}m^{2\alpha-1}\right)$, meaning Laffont and Martimort-style incentive compatible coalitions that are of size $O\left(\frac{1}{\alpha+1}\right)$ will prevent relative efficiency from dying in n .

This does not make a large difference when α is large (tails are thin or close to thin). However for reasonable values of α , such as $\alpha = 3$, and large populations, these unilateral deviations incentives will be just enough to restore the robustness of QVB to moderate-sized coalitions as $\sqrt[4]{n}$ when n is on the order of the number of voting-age individuals in the United States is on the order of 100. Thus in my view the concerns about the potential for collusion to implement extremist strategies are significantly but not fully allayed by incentives for unilateral deviation.

4.4 The de-merger problem

Another challenge to the efficiency of a mechanisms that is particularly worrisome under VCG is the de-merger problem identified Ausubel and Milgrom (2005). This refers to gains individuals might achieve by pretending to be multiple individuals. The de-merger problem arises for reasons similar to those that create incentives for collusion and its analysis is similar to that of collusion. Clearly de-mergers can create problems for many mechanisms throughout economics and particularly in social choice. “Vote early, vote often” is a classic refrain in democratic politics. The concern, therefore, is not so much that an individual could gain some advantage by taking on multiple identities. Rather, the fear is that an individual might, just by acquiring one or a small number of additional identities, completely change the behavior of the mechanism. The following proposition gives results closely related to those in Subsection 4.2 that bound the efficiency cost created by an individual taking on a small number $l > 1$ of identities.

Proposition 2. *In the average case, when the de-merging individual is random and $\mu = 0$, a de-merger of an individual into l individuals causes relative inefficiency to be $O\left(\frac{l}{\sqrt{n}}\right)$ and thus to decline so long as $l = O(\sqrt{n})$. When $\mu \neq 0$ de-merger problems do not affect the rate at which efficiency is achieved as long as $l = O\left(n^{\frac{2}{1+\alpha}}\right)$.*

In the worst case, when the de-merging individual is the most extreme individual in the population and $\mu = 0$, relative inefficiency is $O\left(\ln \frac{2-\alpha}{2\alpha}\right)$ and declines in n so long as $l = O\left(n^{\frac{\alpha-2}{2\alpha}}\right)$. In the worst case when $\mu \neq 0$, relative inefficiency is $O\left(l^\alpha n^{-\frac{\alpha-1}{\alpha+1}}\right)$ and declines in n so long as $l = O\left(n^{\frac{\alpha-1}{\alpha(\alpha+1)}}\right)$.

The results follow from essentially the same logic as in Subsections 4.1 and 4.2 above, with one crucial difference: a de-merging individual operates like a collusive group of size l but with $u_M = u_i$. She does not have larger scale of value than a single individual does. This reduces the dangers created by de-mergers relative to those created by collusion for the same-sized group ($l = m$).

Proof. Will appear in next draft. See Appendix A for a proof sketch. □

Again one can interpret this result by considering realistic values of n and α . The average cases are of no concern; for $\alpha = 3$ problems only arise if $l = \Omega(\sqrt{n})$, which is implausible for even moderate-sized n . In this worst-cases for $\alpha = 3$ both results indicate that inefficiency will fail to die if $l = \Omega(\sqrt[6]{n})$. This requires that l be on the order of 10 in a population of a million individuals. This seems quite implausible and is certainly less of a concern than is collusion. Similar calculations to those above also show that a de-merger into two individuals, which is probably all that is feasible in most cases, only raises the chance of an extremist strategy by a factor of 16 if $\alpha = 3$ which would have a negligible impact on efficiency.

4.5 Other challenges to collusion and de-mergers

While unilateral deviations do not fully preclude collusion to implement an extremist strategy, collusion to implement extremist strategies are still not an equilibrium. The reason is that the same increase in the density of pivotality caused by the collusive extremism that makes collusion potentially incentive compatible also raises all individuals' incentives to purchase votes, not just those within the collusive group. In fact, if individuals believe there is a chance that a collusive clique will attempt to pursue an extremist strategy, this will raise their optimal vote purchases, reducing the attractiveness of such a collusive strategy. The possibility that an individual will de-merge will have the same effect. Collusion, then, effectively reduces the tail coefficient (increases the chance of extreme values) but does not fundamentally change the logic of equilibrium from Subsection 2.3.

In fact, in accompanying work (Posner and Weyl, Forthcoming), we plan to show that in equilibrium being an asymmetrically "large" contrarian actually biases equilibrium inefficiently *against* the interest of this contrarian, because "small" players anticipate a larger initial chance of pivotality than does the contrarian who must buy her way to this pivotality. Of course collusion and de-mergers also multiply the optimal votes of the group by a factor of m compared to the simple asymmetric case, and I hope in a future draft of this paper to analyze the equilibrium with a collusive group. However, it seems clear that even if unilateral incentives do not undermine collusion for extremism as strongly as they undermine collusion for moderate strategies, if others anticipate collusion its effects will be significantly diminished. It even seems possible that the threat of collusion or a de-merger could raise efficiency overall by increasing the density of pivotality and thus the number of votes that moderates buy, thereby deterring unilateral extremists.

5 Aggregate Uncertainty

All of the analysis above assumes that γ is commonly known and thus that all individuals values are drawn independently and identically from a known distribution. When γ is not known, individuals' values are implicitly correlated through γ .

Suppose that γ is a scalar in \mathbb{R} and that the distribution of values is increasing in γ in the sense of first-order stochastic dominance. Let γ have a continuous density function h and suppose that its support is an open interval including a point for which $\mu(\gamma) = 0$. In this case, I now sketch an analysis of the large population *limiting* behavior of QVB. I do not consider rates of convergence here, as above, though I hope to do so in a future draft.

It seems reasonable to conjecture, as shown in a closely related setting (discussed more extensively in Subsection 6.7 below) by Krishna and Morgan (2012) and in less closely related settings by other authors they cite, that there exists a critical value γ^* such that, with a large population size, if $\gamma > \gamma^*$ the action is undertaken with probability near unity and if $\gamma < \gamma^*$ the action is not undertaken with probability near zero. If this is the case, then standard arguments show that the density of pivotality for an individual with value u is approximately proportional to $h(\gamma^*|u)$ because all of the events in which individuals are pivotal concentrate around $\gamma = \gamma^*$. Thus each individual buys $v_i \approx kh(\gamma^*|u_i)u_i$, where k is a (limiting, n -dependent) constant *across individuals*.

I can then derive an equilibrium equation governing γ^* based on the consistency condition that the mean of votes purchased must equal 0 when $\gamma = \gamma^*$, as otherwise the point of pivotality would not be γ^* . Thus

$$\bar{v}(\gamma^*|\gamma^*) = E[h(\gamma^*|u)u|\gamma^*] = 0, \quad (3)$$

where $\bar{v}(\gamma|\gamma') \equiv E[h(\gamma'|u)u|\gamma]$, a number proportional to the expected number of votes given that γ is realized and individuals believe the pivotal event is γ' . My conjecture that the decision is made in favor of A and against A respectively with very high probability when γ is below and above γ^* will hold if $\bar{v}(\gamma|\gamma^*)$ is below 0 for $\gamma < \gamma^*$ and above 0 for $\gamma > \gamma^*$. There is a unique solution for γ^* if $\bar{v}(\gamma^*|\gamma^*)$ crosses the x-axis once and only once.

Note that limiting efficiency requires that $\mu(\gamma^*) = 0$. Thus equation (3) provide a method for characterizing when efficiency that depends only on properties of the joint distribution of u and γ . In particular let γ_0 be the value of γ such that $E[u|\gamma_0] = 0$. Limiting efficiency is equivalent to $\gamma^* = \gamma_0$ because this guarantees that the action is taken if and only if it is efficient ($E[u|\gamma] > 0$) except in the (zero-measure) case of γ_0 .

$$\bar{v}(\gamma^*|\gamma^*) = E[h(\gamma^*|u)u|\gamma^*] = \int_u \frac{[f(u|\gamma^*)]^2 h(\gamma^*)}{f(u)} u du = \frac{E[ux^2(u|\gamma^*)]}{h(\gamma^*)},$$

where $x(u|\gamma^*) = \frac{f(u|\gamma^*)h(\gamma^*)}{f(u)}$. Thus the condition defining γ^* is that $E[ux^2(u|\gamma^*)] = 0$. On the other

hand, the condition defining γ_0 is that

$$E[u|\gamma_0] = \int_u u f(u|\gamma_0) du = \frac{\int_u u \frac{f(u|\gamma_0)h(\gamma_0)}{f(u)} f(u) du}{h(\gamma_0)} = 0 \iff E[ux(u|\gamma_0)] = 0.$$

Limiting efficiency occurs if $\gamma_0 = \gamma^*$, and thus if there is a γ^* such that $E[ux(u|\gamma^*)] = E[ux^2(u|\gamma^*)] = 0$. In general, this does not imply that $E[u] = 0$; a counter-example is shown in Appendix B. However, I believe that there are some reasonable conditions one can place on the joint distribution of γ and u such that $E[ux(u|\gamma^*)] = E[ux^2(u|\gamma^*)] = 0$ implies that $E[u] = 0$ and conversely will also yield that $E[u] = E[ux(u|\gamma^*)] = 0$ implies $E[ux^2(u|\gamma^*)] = 0$, so that limit efficiency is equivalent to $E[u] = 0$. If this is the case, then limiting efficiency occurs if and only if the decision should be an ex-ante expected tie in the sense that the unconditional mean of value is 0 and, under similar conditions, I believe that I could show that QVB is “biased” in favor of the ex-ante expected utilitarian losing outcome (the “underdog”) in the sense that whenever $E[u|\gamma]$ favors this outcome it wins, but there are cases when $E[u|\gamma]$ disfavors this outcome (and thus it should lose) but still wins. Again, as in Subsection 2.2, this is a manifestation of the Taylor and Yildirim (2010) Underdog Principle.

While I have yet to obtain any general conditions under which this logic holds, I have been able to analyze three examples. Analysis of these three involves calculations that are of limited interest and I thus simply describe the results obtained in the text and defer all calculations justifying these results to Appendix B.

Example 1. *Suppose that $f(u)$ and $f(u|\gamma_0)$ are both symmetric about 0. Then QVB is limit-efficient, though so is voting.*

In fact if my conjecture is correct then in the cases when $f(u|\gamma)$ is symmetric but majority rule is not expected to produce a tie (one side is favored), majority voting will outperform QVB; the next example illustrates this in the case of the normal distribution.

Example 2. *Suppose that u and γ are jointly normally distributed, with the mean of u being γ , the variance of u given γ being σ_1^2 , the variance of γ being σ_2^2 and the mean of γ being μ . Then voting is limit-efficient, but QVB is not. $\gamma^* = \frac{\sigma_1^2}{2(\sigma_1^2 + \sigma_2^2)}\mu$, favoring the underdog given that $\mu_0 = 0$. Maximal inefficiency occurs when $\sigma_1^2 \rightarrow \infty$, in which case the analysis is exactly as in Subsection 2.2 and limiting relative inefficiency never worse than .045 and is far smaller in all but extreme cases.⁸*

Thus while, in this case, voting is clearly more efficient than QVB, QVB does not perform badly. This case was, by conditional symmetry, as biased in favor of voting as it could be. I now consider a simple example proposed by Krishna and Morgan (2012), though the results I derive are both different from theirs and in contrast to them in some ways; see footnote 9 below.

⁸The many striking similarities between the analysis here and in the μ local-to-zero case explored in Subsection 2.2 remain mysterious to me, though hopefully I will be able to explain them more clearly in a future draft.

Example 3. Suppose that γ is the fraction of individuals who have a positive value, the others having negative value, but that the distribution of value intensity conditional on its sign is known and possibly different for each sign. Then the Underdog Principle holds in the sense that $\gamma^* > (<)\gamma_0$ if and only if $E[u] > (<)0$.

Suppose, additionally, that γ has a uniform distribution on $[0, 1]$. Voting and QVB are both limit-efficient if the average intensity of preference for individuals with positive value, μ_+ , is the same as the average intensity of preference for individuals with negative value, μ_- . Both are inefficient otherwise. However, except in the knife-edge case of equality, QVB is always more efficient than voting and often substantially so. The relative efficiency of QVB is always greater than .94 while the relative efficiency of voting may be as low as .5. QVB is most inefficient when $\frac{\mu_+}{\mu_-}$ is close to .15 (or $\frac{1}{15}$) while voting is most inefficient when $\frac{\mu_+}{\mu_-} \rightarrow 0$ or ∞ . Qualitatively and quantitatively the same characterization applies when γ has a symmetric Beta distribution (with parameters $\alpha = \beta$).

More generally, if γ has a Beta distribution QVB or voting may be more efficient. However, despite extensive experimentation with the parameters of the Beta distribution (α, β) and with (μ_-, μ_+) , I have yet to find an example where the relative efficiency is ever below .9. Voting, on the other hand, can be worse than a coin toss under a reasonable range of cases, such as when $\alpha = 10, \beta = 1$ and $\frac{\mu_+}{\mu_-} < .1$, and in fact can be arbitrarily worse. Generally I found that for most parameter ranges QVB outperforms voting, sometimes very greatly. This was reversed only when $\frac{\mu_+}{\mu_-}$ was close to 1. Plots are shown in Appendix B.

Krishna and Morgan’s case is clearly very special. But it gives a flavor of a broader conjecture I have that while voting may outperform QVB when the conditional distribution of value is very close to symmetry, it both does not do so typically and has the potential to do very badly when symmetry is significantly broken while QVB rarely loses much efficiency. Voting, on the other hand, may easily be arbitrarily inefficient. For example, under aggregate certainty, which is a special case of aggregate uncertainty, voting can always make the wrong choice and thus be as much worse than a coin toss as the first best is better than a coin toss. In a future draft I hope to show the robustness of QVB’s efficiency compared to that of voting more generally.

Finally, I relate the conditions for efficiency to whether the election is an ex-ante expected tie.⁹ The ex-ante average number of votes in the election is

$$E[h(\gamma^*|u)u] = \int_{\gamma} \int_u u f(u|\gamma) h(\gamma) \frac{f(u|\gamma^*) h(\gamma^*)}{f(u)} du d\gamma = \int_u u \frac{f(u|\gamma^*) h(\gamma^*)}{f(u)} \int_{\gamma} f(u|\gamma) h(\gamma) du =$$

⁹Krishna and Morgan (2012) refer to ex-ante expectation of votes aligning with the ex-ante expected welfare optimal choice “utilitarian optimality”. Given the result here, QVB is “utilitarian optimal” in this sense as long as my conditions for the Underdog Principle hold as the ex-ante expected winner of the election is always that individual favored under QVB, who in turn is the ex-ante expected welfare-optimal choice if the Underdog Principle holds. This generalizes the Krishna and Morgan analysis under aggregate uncertainty.

$$h(\gamma^*) \int_u u f(u|\gamma^*) du,$$

which equals 0 if and only if $\gamma^* = \gamma_0$. Thus the election is an ex-ante expected tie under QVB if and only if QVB is limiting efficient. Whichever side is ex-ante favored will be discriminated against (there will be some γ values for which it should win but does not).

6 Relationship and Comparison to Other Mechanisms

This section discusses the conceptual and normative comparisons and contrasts between QVB and other mechanisms for binary collective decisions. I mostly draw on my formal analysis from above, but some informal new results and one formal new result are also established.

6.1 Convex power vote buying, democracy and dictatorship

QVB is one member of a broader class of *convex power vote buying* (CPVB) mechanisms in which the rules are identical to QVB except that (gross, pre-refund) payments by individuals are determined by v_i^x for some $x > 1$. For $x = 2$, this mechanism is exactly QVB. To analyze the mechanism for other values of x , assume that the density of pivotality, g , is constant at g_0 as a function of the number of votes purchased and symmetric across individuals, as in the heuristic analysis that begins Section 2. Then equilibrium is given by

$$g_0 u_i = x v_i^{x-1} \implies v_i^* = \text{sign}(u_i) \left(\frac{g_0}{x}\right)^{\frac{1}{x-1}} |u_i|^{\frac{1}{x-1}}.$$

Two limits are useful to consider. First, if $x \rightarrow \infty$ every individual purchases exactly the same absolute value of votes regardless of her value, but the sign of votes purchased is still determined by the sign of value. Thus the decision is made based on which alternative has more individuals who weakly favor it. This system is precisely majority rules, one-man-one-vote democracy.

Second, consider the behavior as $x \rightarrow 1$. Suppose there are two individuals, i and j with $\frac{|u_i|}{|u_j|} = r > 1$.

$$\frac{|v_i^*|}{|v_j^*|} = r^{\frac{1}{x-1}} \rightarrow \infty,$$

as $x \rightarrow 1$ from above. Thus the individual with the highest absolute value of value (or if tails are thin or non-existent, the few individuals with the highest magnitude of value) will purchase infinitely more votes in this limit than individuals with slightly lower value.¹⁰ This is consistent with the analysis of standard vote buying by Casella et al. (2012) and Casella and Turban (2012).¹¹

¹⁰This effect is even more extreme if one takes into account the possibility of extremist strategies as in Subsection 2.3, as extreme extremists will typically want to “buy the whole election.”

¹¹They, respectively, propose a novel solution concept for the typically ill-defined equilibrium of linear vote-buying

In the case of thin tails and x small but not quite 1, one should expect the few individuals with the largest magnitude of value to essentially determine the outcome regardless of the preferences of other individuals.

Thus CPVB essentially nests democracy, dictatorship (of the individual with the largest magnitude of value) and passes through something like an oligarchy of the few individuals with close to the largest magnitude of value on the way.¹² It also reveals a range of alternative collective choice procedures in between, including the one I claim is optimal, $x = 2$ (QVB). These systems could be analyzed much as above and compared for various properties other than those related to collusion, budget balance and efficiency considered here.

For example, one might study which of these reduces inequality of wealth by the most in equilibrium. A detailed analysis of this is available on request, but shows that under calibrated numerical assumptions a value of x around 1.5 – 1.7 appears to be maximally redistributive, though $x = 2$ is not far off. Thus in some cases democracy may not be in the interests of those with lower incomes, as they are more than compensated for their reduced influence on the outcome by the redistribution that occurs more heavily with lower values of x . Goeree and Zhang (2012) find that a mechanism closely related to QVB along these dimensions is preferred by nearly all individuals to democracy in a laboratory experiment for this reason.

6.2 Uniqueness

The previous subsection suggests that QVB is the unique rule achieving limit efficiency under common knowledge of γ . Of course this claim is only true within some class of mechanisms: clearly the VCG and EE mechanisms, discussed in greater detail in the next subsection, are always efficient. And the uniqueness only holds if one insists a mechanism be efficient with sufficient robustness: under symmetry conditional on γ , majority voting is efficient in large populations.

To limit the class of mechanism under consideration, I use four criteria. These are defined more formally in the proof of uniqueness in the appendix and are stated informally here for the sake of clarity:

1. *Symmetry*: I call a mechanism *symmetric* if all individuals are treated symmetrically, even at the interim stage. Thus any mechanism involving randomization (such as randomized dictatorship) would not be symmetric in this sense.

models based on a centralized clearinghouse analogous to the one in CPVB and, under that concept, find that the dictatorship of the individual with the largest value in magnitude is the unique equilibrium whenever such an individual exists.

¹²While it may seem a bit unnatural to identify these latter two outcomes with “dictatorship” or “oligarchy” as traditionally conceived, note that in standard discrete choice models (Shaked and Sutton, 1982; Berry et al., 1995) quasi-linear values (approximating log-utility models) are the product of income and idiosyncratic tastes. Thus, especially if the distribution of income is fat-tailed while idiosyncratic tastes are thin tailed, the relevant dictatorship or oligarchy is likely to be of the wealthiest or a wealthy caste, as is typically the case in the classical sample of constitutions described by Aristotle (350 B.C.E.).

2. *Detail-freedom*: I call a mechanism *detail-free* if its rules depend only on the number of agents (and thus not on any other property, such as the distribution of valuations, the value of γ , etc.). Detail-freedom has no bite unless one requires that a mechanism be applicable across a range of settings, but has a strong bite in this case. For example, the EE mechanism is not detail-free if it is required to apply to more than a single distribution of valuations; similarly Goeree and Zhang (2012)'s mechanism is not detail-free if it applies to distributions that do not have the same standard deviation. On the other hand, the VCG mechanism is detail-free.
3. *Separability*: I call a mechanism *separable* if each individual's transfer may be written as the sum of two terms, one of which depends only on her own action/report and the other of which depends only on the reports of other individuals. While, as we saw in Section 4, separability does not prevent all collusion. However, many of the extreme incentives for collusion in VCG are (partly) a result of its non-separability. By contrast, the EE mechanism is separable.
4. *Scale-invariance*: I call a separable mechanism *scale-invariant* if the component of the individual's transfer depending on her own action is independent of population size. Clearly the notion of scale-invariance makes sense only for a separable mechanism, as any component depending on other individuals' actions must (implicitly) depend on the population size. But some separable mechanisms are not scale-invariant. For example, EE and the Goeree and Zhang mechanism are not scale-invariant, even though they are separable.

Note that most standard mechanism used in practice in a wide range of settings satisfy most or all of these properties. All standard, symmetric, reserve-free auction formats are detail-free, and all but the second-price and English auction are separable and scale-invariant (at least conditional on the allocation). Voting is symmetric, detail-free, separable and scale-invariant.

Theorem 2. *A mechanism is symmetric, detail-free, separable, scale-invariant and achieves efficiency approaching 1 in any Bayesian equilibrium for large n given any fixed, i.i.d. value distribution whose first two moments exist if and only if it is equivalent (up to relabeling of actions) to a mechanism where individuals report a scalar v_i , the action is taken if $\sum_i v_i > 0$ and only if $\sum_i v_i \geq 0$ and have transfer $t_i = f(v_i) + T(v_{-i}, n)$ where $f'(0) = 0$ and $f''(0) > 0$.*

Thus any detail-free, scale-invariant mechanism where different players' reports do not interact to produce an individuals' transfer and which is limit-efficient regardless of the details of the distribution in question must be approximately quadratic about 0. This is not precisely QVB, or even a rescaled version of QVB, as higher order terms might exist; for example, an inverted normal probability density function would have this property. But it seems reasonable to say that QVB is the simplest of rules that are symmetric, detail-free, separable, scale-invariant and robustly limit-efficient. QVB is clearly the only CPVB rule satisfying this criterion.

Proof. Will appear in a future draft. □

The argument will proceed in two steps. The first step will show that any symmetric, detail-free, separable, scale-invariant rule robustly achieving limit-efficiency must have a “vote-buying form” where the decision is taken based on the sign of the sum of votes and votes are costly. Only such a rule can aggregate the additive function of individuals’ values necessary to achieve efficiency.

Once the class of possible rules is narrowed to such vote-buying scheme, the argument will formalize the logic of the previous subsection using the mechanics developed in Section 2. Any rule other than one that is quadratic near the origin fails to linearly aggregate values and thus may be inefficient in the limit when the distribution of valuations is such that the divergence between linear aggregation and whatever function of value the rule aggregates leads to a switch in signs, which will occur under some distribution with its first two moments existing. Only behavior near the origin is critical as, in large populations, only the behavior of the cost function near the origin is relevant.

To see why each of the properties is necessary for the theorem to hold, note that if separability were dropped, VCG would fit the bill. If scale-invariance were dropped, mechanisms whose payment rule was not quadratic for any n but converged to being quadratic for large n would work. If the requirement of efficiency for *all* (reasonable) value distributions were dropped, some form of voting might work; for example Ledyard and Palfrey (1994, 2002) show that if attention is restricted to distributions with a quantile that is known to correspond to the mean, voting with an appropriate super-majority rule is efficient as $n \rightarrow \infty$. If symmetry is dropped, randomization can be used to imitate QVB as discussed in Subsection 6.7 below. The one redundancy is detail-freeness; there is no separable, scale-invariant, robustly limit-efficient rule that is not detail-free. This property is included for emphasis and because it has a clearer economic motivation than scale-invariance.

6.3 Comparison to voting

The most canonical and common system of collective decision-making in the modern world is voting, most commonly by majority rule and sometimes by super-majority rule. In this subsection I compare the merits of voting and QVB along many dimensions. I assume voting is free and all individuals vote; a discussion of the case when voting is costly is deferred to Subsection 6.7.

- Dependence on details and breadth of applicability: Both QVB and voting (with a fixed majority or super-majority rule) are detail-free and may be applied to essentially arbitrary settings, though QVB requires the availability of transfers while voting does not.
- Efficiency with aggregate certainty: With aggregate certainty, voting is never efficient in small populations (Bowen, 1943) and is efficient in the large population limit if and only if the supermajority quantile always has the same sign as the mean of the distribution of values (Ledyard and Palfrey, 1994, 2002). QVB is efficient, under some regularity conditions, in the

limit for any distribution of valuations, and I hope to show in future work that even in small populations it is almost always more efficient than voting. Goeree and Zhang (2012) have done so for some simple examples both experimentally and theoretically, and I hope to extend their analysis to consider a richer set of environments.

- **Efficiency with aggregate uncertainty:** Voting is efficient in large populations if and only if, conditional on γ , the appropriate quantile of the distribution of valuations has the same sign as its mean. QVB, by contrast, is conjectured to be efficient, under some regularity conditions, in large populations if and only if the ex-ante expectation of utility is 0. Numerical results suggest that voting has much greater potential for inefficiency than does QVB and that QVB is typically more efficient than voting, though results thus far are extremely specific.
- **Strategic complexity:** QVB requires individuals to form beliefs about the actions of other players. Voting dominant strategy and direct and thus very easy for participants.
- **Susceptibility to collusion and de-mergers:** Collusion and de-mergers are ineffective at reducing efficiency for very small groups or de-mergers (of fixed size) under both QVB and voting under plausible assumptions, though this is more robustly true of voting than of QVB. Moderately-sized collusive groups or de-mergers can be effective in some settings under QVB and will not be under voting. Large collusive groups and large de-mergers are effective under both mechanisms. Unilateral deviation incentives discipline collusion in both contexts. Unilateral deviation incentives discipline more strongly under voting, but the reactions of others discipline collusion and de-mergers more under QVB than under voting (which is dominant strategy and thus does not provoke equilibrium reactions by other players).
- **Budget balance:** Both mechanisms are fully budget-balanced.

Overall, there are fairly strong reasons to believe that when intensity of preference is important (when the quantile of the distribution of valuations corresponding to its mean is not known ex-ante), QVB is likely to be more attractive than voting along many dimensions. However voting has a number of stronger robustness properties than QVB satisfies and thus is still likely to be more attractive in contexts where strategic simplicity and robustness to collusion are highly valued, as emphasized by Bierbrauer and Hellwig (2011).¹³

6.4 Comparison to the Vickrey-Clarke-Groves mechanism

The most canonical mechanism for collective decisions proposed by economists is the Vickrey (1961)-Clarke (1971)-Groves (1973) mechanism. This mechanism has rarely been used in practice but is, in some sense, the archetypical economist's alternative to voting.

¹³Given that, as discussed in Subsection 6.1 above, CPVB buying nests both voting and QVB it is possible that in some settings a convex vote buying scheme with $2 < x$ would be desirable to balance these concerns.

- Dependence on details and breadth of applicability: QVB and VCG are both detail-free and applicable to essentially arbitrary problems, though both require transfers and VCG also requires individuals to have unlimited budgets (be able to make arbitrarily large payments, up to their cardinal value).
- Efficiency with aggregate certainty: VCG is efficient (at the unique equilibrium in weakly-dominant strategies) regardless of the number of individuals. QVB is efficient when the number of individuals is large.
- Efficiency with aggregate uncertainty: VCG is always efficient, in large or small populations. QVB is sometimes efficient in large populations and rarely very inefficient.
- Strategic complexity: VCG is dominant strategy and individuals simply report their valuations. QVB requires individuals to estimate the density of pivotality.
- Susceptibility to collusion and de-mergers: As emphasized by Ausubel and Milgrom (2005), VCG is highly susceptible to collusion and de-mergers in that any two individuals (or a single individual de-merged into two) with the same sign of value can achieve their desired outcome at zero cost and in equilibrium. It is thus not even *coalition-proof* for a two-person coalition in the sense of Laffont and Martimort (1997, 2000). By contrast, QVB is moderately robust against collusion and de-mergers as discussed above.
- Budget balance: QVB has a balanced budget while VCG, if it is self-financing, runs a surplus in many cases. In fact, when $\mu = 0$ the surplus is typically on the same order as the efficiency gains as compared to making a random decision.¹⁴
- Another problem of VCG is that it “requires individuals to have budgets “in reserve” much larger than the payments they are expected to make with any substantial probability, as otherwise they will be forced to under-report the intensity of their value to stay within their budget (Ausubel and Milgrom, 2005). QVB and the other mechanisms only require budgets commensurate with anticipated payments.

While there are many dimensions on which VCG outperforms QVB, its problems with collusion and budget balance are so severe that many (Ausubel and Milgrom, 2005), including myself, consider it unworkable in practice. While QVB has some issues with collusion, they are far less severe, and it entirely solves the budget balance problem of VCG.

¹⁴Some have suggested mitigating this problem by dividing the group of individual into two, running VCG in each sub-group, randomly implementing the outcome chosen by one of the two groups and refunding the revenue to the other group. While this makes the mechanism budget-balanced, it introduces other problems: when $\mu = 0$ or $\mu = \frac{\hat{\mu}}{\sqrt{n}}$ the mechanism is no longer efficient even asymptotically. When $\mu \neq 0$ this does solve the budget balance problem, but the budget balance problem is extremely small in this case (0 with probability near 1). Nor does it address the collusion problem. Thus I do not view this alternative as serious.

6.5 Relation to Expected Externality and Goeree-Zhang mechanisms

Another canonical mechanism proposed by economists is the *Expected Externality* (EE) mechanism of Arrow (1979) and d'Aspremont and Gérard-Varet (1979). In this mechanism individuals report their valuations but pay the expectation of their VCG payments. More recently a variant on this mechanism that is closely related to QVB was proposed by Goeree and Zhang (2012). They note that in a large population if $\mu = 0$, EE payments converge to $\frac{1}{2\sqrt{2n\pi}\sigma}u_i^2$ for individuals with value u_i .¹⁵ They advocate replacing EE payments with this approximation. I refer to this mechanism as the GZ mechanism and discuss the comparison of QVB to EE and GZ. I then discuss in greater detail the relationship between QVB and GZ.

- Dependence on details and breadth of applicability: QVB is detail-free and can be applied in any setting with transfers. EE requires common knowledge of the distribution of valuations (it is not well-defined otherwise) and its rules depend on this distribution. GZ requires that the distribution not only be commonly known but also that $\mu = 0$ (it is not well-defined otherwise, as it is not clear whether it should be the EE limit in this other case or something closer to QVB), but in this case does not depend on the whole shape of the distribution as EE does but only on the standard deviation of the distribution σ .
- Efficiency with aggregate certainty: EE is fully efficient, while QVB and GZ have identical properties when $\mu = 0$ (both limit-efficient). GZ is ill-defined when $\mu \neq 0$ or $\mu = \frac{\hat{\mu}}{\sqrt{n}}$.
- Efficiency with aggregate uncertainty: EE and GZ are ill-defined in this case. QVB is efficient under conditions described above.
- Strategic complexity: EE and GZ are direct mechanisms and thus strategically simple, though neither has a dominant strategy. QVB requires individuals to estimate the density of pivotality in equilibrium.
- Susceptibility to collusion and de-mergers: In the $\mu = 0$ case where the two are essentially identical, QVB and GZ's efficiency is robust to collusion by moderate-sized groups; while I am not aware of a similar analysis for EE, I suspect the same is true. All of these mechanisms are coalition-proof. However, when $\mu \neq 0$ (and GZ is thus ill-defined), collusive strategies nearly as devastating as those possible against VCG exist in EE in large populations, though I am not aware of any previous observation that this is the case. Suppose WLOG that $\mu > 0$ and that two individuals i and j with negative value both report $u_i = u_j = -\frac{3n\mu}{4}$. Note that their Expected Externality payments will be extremely small as the chance that the outcome

¹⁵There is a slight divergence between this formula and what appears in their paper because their normalization is such that my u_i is equivalent to twice their u_i (which they label v_i).

changes as a result of their report is approximately

$$\Phi\left(\frac{\sqrt{n}\mu}{\sigma}\right) - \Phi\left(\frac{\sqrt{n}\mu}{4\sigma}\right) < \frac{3\mu e^{-\frac{n\mu^2}{16\sigma^2}}}{4},$$

which is tiny for even moderately-sized n . Despite making such *de minimis* payments that would likely be easy to enforce internally, the group ensures that with probability near 1, A is not implemented. Thus EE is, in the more common case that $\mu \neq 0$, nearly as problematic from a collusion perspective as is VCG. Because this strategy could and would be attractive to a de-merging individual who would face no constraint from unilateral deviation, from a de-merger perspective EE is equally problematic as VCG. Thus the discussion in the previous subsection about the advantages of QVB over VCG on this dimension largely carry over to the comparison to EE outside the $\mu = 0$ case. This emphasizes the sense in which QVB, unlike GZ, is *not simply the large population limit of EE*. The large population limit of EE outside the $\mu = 0$ behaves very differently to QVB both qualitatively and quantitatively.

- Budget balance: When they are well-defined, EE and GZ are budget-balanced, as is QVB.

I consider EE and GZ to be unworkable as general purpose mechanisms. Both apply only to cases of mostly theoretical rather than practical interest, especially in the case of GZ. EE, which applies somewhat more broadly, is nearly as vulnerable to collusion as VCG in those additional contexts. As literally written in their paper, the GZ mechanism does not offer any significant advantages over EE that I am aware of and is less broadly applicable. I thus do not see any context where one would be unwilling to use EE but would be willing to use GZ in practice, and I find it unlikely that EE is a mechanism of practical interest in many settings.

However, GZ is based on a fundamental insight that is closely related to that which motivates QVB: the large population limit of EE is quadratic payments in the case when $\mu = 0$. I begin by describing the intuition behind this insight and then discuss how QVB and GZ, despite both being based on this insight, split ways following it.

The basic intuition behind the GZ result is that in a large population the distribution of the sum of all other individuals' valuations is approximately uniform because the distribution is "stretched" by the entry of more individuals leading it to flatten over the range of valuations any individual is likely to draw. This makes the Expected Externality of any individual approximately quadratic because both the probability of being pivotal ($P(U \cdot U_{-i} < 0)$) and the expected harm to all other individuals by being pivotal ($E[|U_{-i}| | U \cdot U_{-i} < 0]$) grow linearly with reported value, and the expected externality of an individual is the product of these two. This point is essentially equivalent to the classic Dupuit (1844)-Jenkin (1871–1872)-Harberger (1964) argument that small distortions to prices create welfare losses that are triangles and thus have a size quadratic in the extent of the distortion. In a large population, the preferences of any individuals create a small distortion to the

maximization of all other individuals' preferences and thus create a triangular loss, quadratic in the intensity of the preferences of that individual.

While QVB and GZ are both based on this idea, QVB splits ways by noticing that the scaling coefficient in front of value is irrelevant to the set of equilibria of QVB and GZ in the case of $\mu = 0$. This is because any equilibrium of GZ is an equilibrium of QVB where individuals scale down their reports by a factor of $\sqrt{\frac{1}{2\sqrt{2n\pi}\sigma}}$. This delivers a detail-free mechanism which may then be applied outside the case of $\mu = 0$, a broader set of cases on which most of my analysis focuses. The cost of this detail-free approach is that it is not a direct mechanisms and thus relies on participants to estimate properties of the distribution of equilibrium play as do the all-pay and first-price auction. But the benefit is that it is far more robust along other dimensions.

6.6 Comparison to other mechanism

Many other mechanisms for collective decision-making have been proposed that have received relatively less attention in the literature. I believe this limited attention is a result of these mechanisms being clearly unrealistic or of severely limited applicability for various reasons. In the interests of brevity I do not devote an extended discussion to each of these. However, some of the reasons that these mechanisms are unrealistic illustrate the “practicality” and “simplicity” of QVB. As a result, I briefly discuss several of these mechanisms and for each describe what I view as its clearest limitation and how QVB avoids this. The list is far from exhaustive, but I do believe it is representative of most potential strategies one might take to alternative mechanisms for this problem. An exception is a class of mechanism proposed by Casella (2005), Jackson and Sonnenschein (2007) and Hortala-Vallve (Forthcoming) that I discuss in Subsection 8.2.

1. *Implement the ex-ante efficient choice*: One simple mechanism that applies only under aggregate certainty is to simply implement the ex-ante efficient choice. While this sacrifices some efficiency when $\mu = 0$ and $\mu = \frac{\hat{\mu}}{\sqrt{n}}$ it can be shown that for large n in the average case this always out-performs QVB. I see there being two fundamental weaknesses of this mechanism. First, it does not apply (is or at least is not efficient) with aggregate uncertainty. Second, and more importantly, it is hard to imagine this being implemented in practice. It would require some benevolent agent to make a determination of the sign in practice. This individual would become a dictator as there would be no ability to check her actions. By contrast QVB's clearly defined rules avoid such a possibility.
2. *Maskin (1999)'s mechanism*: Concerned by precisely this problem, Maskin (1999) proposes a rather intricate mechanism that I will not describe here in which individuals who commonly know the state are induced to report even if there is no benevolent agent who can act on it directly. Unfortunately this mechanism, too, seems implausible. Again it fails with even a

	Detail-dependence	Aggregate certainty	Aggregate uncertainty	Strategic complexity	Collusion and de-mergers	Budget	Other
QVB	Detail-free	Limit-efficient	Efficient when $E[v] = 0$ and inefficiency small otherwise	Non-direct	Moderately robust	Balanced	
Voting	Detail-free	Limit-efficient when mean and median have same sign, often quite inefficient	Limit-efficient when mean and median have same sign, often quite inefficient	Strategy-proof	Highly robust	Transfer-free	
VCG	Detail-free	Efficient	Efficient	Strategy-proof	Extremely sensitive	Surplus	Requires large "reserve budgets"
Expected Externality	Detail-based, requires aggregate certainty	Efficient	Ill-defined	Direct but Bayesian	Highly sensitive	Balanced	
Goeree and Zhang (2012)	Detail-based, requires $\mu = 0$	Limit-efficient when $\mu = 0$, otherwise ill-defined	Ill-defined	Direct but Bayesian	Highly robust when $\mu = 0$	Balanced	

Table 1: Comparison of binary collective choice mechanisms

small amount of aggregate uncertainty. Second and more importantly, it has many equilibria, not just those that are efficient. I have never heard anyone seriously propose using this mechanism in practice. By contrast QVB can be used with aggregate uncertainty, and its efficiency properties hold at all of its equilibria in the limit.

3. *Thompson (1966)'s insurance mechanism*: Thompson (1966) proposes an mechanism in which individuals purchase insurance against the collective decision being made equal and opposite to their preferences. If individuals are risk-averse, the insurance market is frictionless and, crucially, *the utility derived from the public good has the same impact on the marginal utility of wealth as does a change in money* then in large populations this leads to efficiency. However, the assumption that changes in the public good are commensurate with money, and therefore are the sort of thing that consumers would like to perfect insure, seems implausible in most settings. For example a victory by a kleptocratic candidate in an election might well lower all agents' values but also lower the marginal utility of (an insurance payout of) money for them, in which case Thompson's mechanism would pick out precisely the wrong outcome. By contrast, so long as the stakes in the decision are not too high, the quasi-linear approximation that QVB relies on is always valid (Willig, 1976; Bewley, 1977) to achieve efficiency in the sense of Kaldor (1939) and Hicks (1939).
4. *Crémer and McLean (1988)-McAfee and Reny (1992)-style mechanisms*: Crémer and McLean (1988) propose an auction mechanism exploiting the correlation structure of beliefs and preferences to extract information from agents. McAfee and Reny (1992) show that a similar structure works to achieve efficiency and budget balance with aggregate uncertainty (though not without it) in a collective choice problem. Just as the auctions literature views the Crémer and McLean result largely as a *reductio ad absurdum* on the standard ways of modeling correlation rather than a serious proposal, I do not believe the McAfee and Reny mechanism is plausible for social choices for many of the same reasons. First, unlike QVB (which is detail-free) the strategy is highly sensitive to the precise structure of correlation, which is unlikely to be available to the planner. Second and more importantly, it is highly sensitive to the perfect relationship between beliefs and preferences (Heifetz and Neeman, 2006). While I have not studied here the performance of QVB when beliefs and preferences are not in perfect relationship to one another, I suspect that QVB is more robust. This is an interesting direction for future research. Third, introducing a small amount of risk-aversion ruins the McAfee and Reny strategy in most settings. QVB is not sensitive to a small amount of risk-aversion as payments are small.

While theoretically provocative, few would seriously consider the mechanisms mentioned above or various related mechanisms as proposals for practical social choice problems. The comparison to

these mechanisms illustrates the robustness of QVB and why I believe it should be considered as a serious alternative for binary collective decisions.

6.7 Relationship to existing practices

If QVB is such a compelling mechanism, why don't we see it in use? The closest analog I have been able to identify from anecdotes in my own life is committee work, such as hiring committees or departmental votes on hiring. While hiring decisions are usually formally decided by some sort of vote, the decision threshold is usually ambiguous and votes are not usually fully sincere but instead reflect a process of bargaining among the faculty. In my experience, it is usually easier for an individual faculty member, in this bargaining process, to just push over (or under) the bar several candidates who are very close to the bar than it is for them to take a single candidate who others judge to be well below or well above the bar and switch them (though this is also possible, just requiring much more effort/political capital expenditure). This convex cost of changing an individual decision is roughly consistent with the spirit of QVB, or at least convex vote buying as in Subsection 6.1. It may even be that, under careful scrutiny, informal rules roughly approximate quadratic costs in political capital/effort. My analysis suggests a more formalized system with quadratic costs that might aid efficiency.

Ledyard (1984) proposed another potential analog between QVB and practice; see also Myerson (2000) and Krishna and Morgan (2012). He considers a model of standard majority voting where voting is costly, voting costs are heterogeneous and distributed independently of value where the distribution of values has a bounded support and the distribution of voting costs has a continuous density on a weakly positive support (with no mass point at 0). In such a model, in large populations, there is effectively a representative individual with each value whose cost of buying (a mass of) votes v is v times the average cost of voting for the v fraction of the population with the lowest voting cost. If voting costs are uniformly distributed this leads to a quadratic cost and thus QVB.

Ledyard goes on to argue that in a large population (because of the boundedness of the value distribution) it is only the individuals with lowest voting cost that will vote. By much the same arguments I used in Section 2, any continuous distribution is approximately uniform in a local neighborhood about a single point. Individuals with very small voting costs are located local to 0, and thus in large populations the distribution of voting costs that matter are approximately uniform. Therefore, in the large population limit each representative individual faces quadratic costs of voting, implying that costly voting converges to QVB and thus, as Ledyard emphasizes, efficiency.

While I appreciate the elegance of this argument, I agree with the Ledyard that the conditions required to ensure it holds are unlikely to obtain in many applications. The result requires that in large populations a vanishingly small fraction of the population votes. This is rarely observed in

reality. It is also questionable whether a bounded value distribution and independence of value and voting costs is plausible; both are likely to rise close to multiplicatively with income, and income distributions are known to be well-approximated for the purposes of mechanism design by fat-tailed distributions rather than ones with finite upper-bounds (Saez, 2001). I instead view Ledyard's argument as a suggestion of a potentially useful mechanism (assigning random costs of voting to different individuals) and an alternative to QVB. However this implementation is not symmetric, which I believe would limit its appeal in practice, and I also hope to show in a future draft of this paper that randomization slows the rate of convergence of the Ledyard mechanism to efficiency compared to QVB, even when value distributions are bounded.

7 Applications

I have tried to make a case that QVB is a plausible, practically-applicable mechanism for collective decision making without yet mentioning any context in which I imagine it being applied. In this section I therefore discuss several applications, though only briefly. As I elaborate below, other papers I am working on with various co-authors explore some of these applications in greater detail. I begin with perhaps the least ambitious and most short-term plausible applications and end with more ambitious and challenging applications.

7.1 Committees

As discussed in Subsection 6.7 above, many committees currently use formal voting procedures but in equilibrium take decisions that, in one way or another, incorporate preference intensities. Formalizing such systems with vote buying would be a natural way to improve on their efficiency. In some contexts, using money for vote buying would be acceptable. But more often some non-pecuniary currency would be more attractive. This could take the form of additional responsibilities for the firm or organization, reduced resources for the committee member's division or simply spending down of storable votes (Casella, 2005) as discussed in Subsection 8.2 below. Experimentation with QVB at such a level would be low cost and hiring committees in economics departments would be a natural starting point, given the openness of economists to novel mechanisms. I have experimented on a small scale with a system like this, have been in discussions with an organization I work with about trying it more broadly, and hope members of other economics departments will consider experimenting with it.

7.2 Holdout and complement assembly

As we discuss in Kominers and Weyl (2012b), in many settings a buyer views a collection of goods owned by different sellers as perfect or near-perfect complements. This creates a familiar holdout

problem that prevents assembly in the absence of coercion, motivating the “eminent domain” procedures that governments often use to address these problems. But, as we argue in Kominers and Weyl (2012a), such coercive procedures can be abused to expropriate sellers. We therefore propose that any buyer be allowed to make an offer for a collection of land but that sellers collectively decide whether to accept the offer using QVB. This system limits inefficiency from both the failure of efficient sales and from abusive sales, while preserving some, but not all, of the property rights of sellers.

A natural starting point for implementing such a system would be in spectrum reassembly efforts underway in many countries. One solution to holdout problems has been to abrogate property rights entirely (Kominers and Weyl, 2012b), analogous to eminent domain. A natural alternative would be the system we propose, and given that the use of sophisticated economic mechanisms is common in the spectrum context both government officials and industry participants are more likely to be comfortable participating in such a system. Recent controversies over eminent domain, and the fact that money changes hands regardless in such transactions, makes land assembly another natural application. All of these are discussed more extensively in Kominers and Weyl (2012a).

7.3 Corporate governance

Corporations are typically governed by share-weighted majority rule of shareholders or, in bankruptcy, by a more complicated voting arrangement among different classes of debt-holders. Such arrangements are not only likely to be inefficient for the reasons I discussed above, but as discussed extensively in the literature on corporate governance (Shleifer and Vishny, 1997), given that shares may be purchased (linearly) they create opportunities for expropriation of minority shareholders by the owner of a bare-majority. Scope for such manipulation of effectively linear vote buying for corporate control has been greatly expanded by so-called “empty voting”, where derivatives positions can be used to separate beneficial ownership from controlling ownership (Hu and Black, 2006).

In Posner and Weyl (Forthcoming), we propose a solution to this problem not by tying ownership and control more closely together but by separating them fully, allowing the purchase of votes direct by QVB. While such a system would not maximize value to the initial shareholders (who off-load shares in an initial public offering), or bond-holders, as would likely be the goal in designing corporate governance, it would avoid many of the problems of inefficient expropriation that corporate governance reforms have been designed to solve. We also show that in many cases ex-post efficient decision making based on QVB is more effective than either simple share-weighted majority voting or linear empty voting in promoting ex-ante share-holder value. Because money is already clearly involved in such decisions, hostility based on the use of money is likely to be minimal, and the capacity to simplify rules for protecting minority shareholders is likely to bring substantial value on its own.

7.4 Public choice

Most ambitiously, QVB or a variant on it could be applied to innumerable public decisions, on passing bills or resolutions in an international body or national legislature, public referenda, election of candidates for public office, etc. A natural hostility to the use of money exists in these settings, but I believe this could eventually be overcome in a combination of five ways.

First, as described in Subsection 8.2 below, if more than one decision is made, at one time or over time, a currency can be allocated to each individual to be spread over different decisions and spent down using a quadratic rule. Properly defined this rule should be internally Pareto-efficient unlike issue-by-issue voting. Of course, it does not allow the exchange of influence over the collective for influence over the allocation of private goods that using money would. At this unavoidable cost, such a system should neutralize any concern over “money in politics”.

Second, if the cause of concern about “money in politics” is the placing of greater weights on the welfare of the rich, this can easily be undone by changing the implicit equal weights on dollar willingness-to-pay in QVB. For example if, as is natural under logarithmic utility and assumed in many industrial organization models since Berry et al. (1995) willingness-to-pay is the product of income and some idiosyncratic component, then making individuals pay a quadratic *share of their income* rather than quadratically in dollars would achieve non-income-weighted efficiency. Other changes in coefficients individuals face in front of their cost of vote buying would similarly shift Pareto-weights in any desired manner.

Third, the concerns that the poor may have about QVB will be partly offset by the redistribution it creates. Goeree and Zhang (2012) show that in some cases individuals with moderate preference (i.e. the poor) actually gain more than those with extreme preferences from QVB because of the resulting redistribution, at least in small populations.

Fourth, some of the concerns among the public about vote buying may be motivated by the fact that most proposals for it have been for linear vote buying. As suggested in Subsection 6.1 and shown by Casella et al. (2012) and Casella (2005), this leads to dictatorship of the singled individual with the most intense preference which is not generally more efficient and often less efficient than majority rule. So there are good efficiency reasons for the public to be wary of standard vote buying. While this reasonable skepticism has hardened into an ideological aversion to money in politics, there may be some scope to persuade the public, at least over the long-term, that quadratic vote buying is “different” than linear vote buying and would be beneficial.

Fifth, some hostility to a system of vote buying may simply arise from a lack of familiarity, as suggested by the experimental work of Goeree and Zhang. Field experimentation in the other areas discussed above may help familiarize the public with the efficacy of QVB and thus, if successful, eventually increase the openness to it as an alternative institution. A natural place to start, rather than directly with choices by voters, would be decisions by representative bodies where the relevant currency could be the net allocation of funds to the regions (e.g. states/provinces/electoral districts

in a federal system or countries in an international body) each delegate represents, funds that are already used for log-rolling in many contexts.

8 Plans for Future Work

This paper is very much in-progress, and the broader project it represents a part of is even more so. In this section I discuss my plans for revising this piece (in the first subsection) and for expanding on its themes in future work (in the second subsection).

8.1 Plans for future drafts

In future drafts of this paper, I hope to accomplish several things in addition to those described in specific sections above. First, I hope to calculate the constant in all asymptotic results above, allowing bounds that would be non-trivial for large but finite populations. Second, I plan to include full formal proofs of all results. Third, I plan to have a much more complete characterization of efficiency under aggregate uncertainty including comparisons of the efficiency of QVB to that of voting under a variety of realistic joint distributions of γ and u , possibly with multidimensional γ . Finally, I plan to include detailed finite population results for particular value distributions, both for QVB on its own and in comparison to voting.

8.2 Plans for other papers

In this paper I attempted to focus as tightly as possible on the core mechanism in a quasi-linear, binary decision, Bayesian equilibrium environment. However, I also hope to explore a range of other issues related to the mechanism in follow-on work (Posner and Weyl, Forthcoming; Weyl, 2013).

1. I hope to extend the mechanism to the case when there are several alternatives. I believe this should be possible by allowing individuals to purchase votes in the decision for each pairwise comparison between alternatives, as individuals will agree on (and view as locally uniform) the probability of that binary comparison being close and binding. I also hope to explore whether there is a way to simplify the mechanism in cases of a large number of alternatives, to avoid individuals having to consider the large set of binary comparisons.
2. I hope to show how the mechanism can work in settings without transfers. The first and simpler of these would be the dynamic contexts considered by Jackson and Sonnenschein (2007) or the simultaneous multi-issues setting considered by Hortala-Vallve (Forthcoming). Jackson and Sonnenschein argue that repeated collective decision problems without transfers are easier than one-shot problems because individuals can be forced to trade-off reports of their types in one period against others. This requires a small type-space and suffers from many

of the critiques discussed in Subsection 6.6. A simpler way to create inter-temporal trade-offs may be to give each individual a number of tokens that may be spread across periods, similar to the storable votes of Casella (2005), or used across many, simultaneously decided issues as proposed by Hortala-Vallve. While Casella and Hortala-Vallve achieve efficiency only under highly restricted conditions (and may perform much worse than standard voting under equally or more general conditions), my results imply that if influence on any issue costs a quadratic quantity of tokens efficiency should result more broadly.¹⁶ Such a mechanism is effectively the same as that proposed by Hylland and Zeckhauser (1979), except for a series of binary rather than continuous public goods. In particular, I conjecture and hope to prove, combining arguments analogous to those employed by Hylland and Zeckhauser (1979) (in the continuous public goods, perfect information setting) with those I use above, that such a rule guarantees Pareto efficiency when *only the number of participants* is large. While the efficiency gains over simple majority rule on each issues are largest when the number of issues are large, Pareto efficiency does not rely on the number of issues being large as it does in Jackson and Sonnenschein (2007).

A more challenging extension would be to determine conditions under which the arguments could be applied to a one-shot decision with several alternatives. Intuitively if the one-shot decision is just a composition of many independent decisions the logic is equivalent to the case above. Conditions under which such a decomposition is possible would be interesting to explore. This might provide conditions under which a Bayesian mechanism can soften the conclusions of the Gibbard (1973)-Satterthwaite (1975) Theorem and mechanisms where individuals do more than report their rankings can soften Arrow (1951)'s Impossibility Theorem.

3. The mechanism in its current form relies on individuals reaching an estimate of the probability of their pivotally and all sharing the same estimate. This makes it potentially vulnerable to the Wilson (1987) critique and makes the mechanism more strategically complex than are direct and especially strategy-proof mechanisms. It may be possible to address this concerns by exploiting Azevedo and Budish (2012)'s "Strategyproofness in the Large" construction. Individuals would be asked to directly report their valuations and a continuous valuation distribution would be constructed. Individuals would then play the equilibrium strategies from QVB given that value distribution. Azevedo and Budish's existing results only apply in private goods (or what they call "market design") contexts not in public goods (or what they call "social choice") contexts, so proving that such a construction provided good limiting incentives would require new theoretical results in addition to an appropriate statistical procedure for fitting smooth value distributions.

¹⁶For example, if the number of issues is large, storable or qualitative voting behaves like linear vote buying, converging to dictatorship.

9 Conclusion

In this paper I propose a novel, detail-free, budget balanced mechanism, Quadratic Vote Buying, for binary collective decision-making with transfers. I show that, under aggregate certainty, in large populations the mechanism approaches first-best efficiency at a $n^{-\frac{\alpha-1}{\alpha+1}}$ rate, where α is the Pareto-tail index of the distribution, and that it is not overly sensitive to collusion, especially when compared to existing efficient mechanisms. Unfortunately QVB's limit-efficiency is not as robust to aggregate uncertainty as are other efficient mechanisms like VCG and it requires strategic calculations by agents participating in it. But it is robust along several dimensions that make most other mechanisms for collective decisions proposed by economists impractical.

In this sense, I view QVB as being close in spirit to the first-price or all-pay auctions, formats that are well-known to be imperfect along many exacting theoretical dimensions but that are robust and thus have been used extensively throughout history. Given this robustness, I believe QVB merits further exploration to determine whether important practical limitations were missed in my analysis. If not, QVB appears to be a promising alternative to democracy.

One natural direction for exploring QVB's practicability, already taken up to some extent by Goeree and Zhang (2012), is in laboratory experiments. Goeree and Zhang consider, as discussed extensively in Subsection 6.5, a variant under which truthful revelation of values is approximately optimal for individuals (given that $\mu = 0$). Equally important is considering whether in the detail-free, non-direct form of the mechanism advocated here, individuals manage to find equilibrium and under what conditions they manage to do so. It would also be interesting to consider cases when $\mu \neq 0$, where the structure of equilibrium is more subtle than that in the $\mu = 0$ case.

Other interesting directions would be extensions of the theoretical analysis. One natural example would be considering a continuous level-of-public-good choice model, an explicitly game-theoretic version of the Hylland and Zeckhauser (1979) set-up. In such a setting it is harder to see how a parsimonious and detail-free mechanism could be constructed that allows use of external transfers, though the Hylland and Zeckhauser (1979) mechanism should create internal Pareto-efficiency if there are multiple dimensions of public goods to be adjusted. Another classical question of particular interest in this setting would be Green and Laffont (1977)'s query about the incentives to acquire information about preferences. The analysis of the model in a context of partially common values with or without heterogeneous priors would be valuable. Another important question is how QVB would perform if, as appears to be the case in practice, agents overestimate their chance of being pivotal or have a civic, non-instrumental motivation for vote buying.

Finally and most broadly, the strategy used here and by Goeree and Zhang (2012) of seeking mechanisms inspired by large-population limits of the Expected Externality mechanism seems promising and might be applied in more complicated settings, such as the combinatorial problems considered by Kominers and Weyl (2012b).

References

Aristotle, *Politics*, Athens, 350 B.C.E.

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

– , “The Property Rights Doctrine and Demand Revelation under Incomplete Information,” in Michael Boskin, ed., *Economics and Human Welfare*, New York: Academic Press, 1979, pp. 23–39.

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

Azevedo, Eduardo M. and Eric Budish, “Strategyproofness in the Large,” 2012. <http://faculty.chicagobooth.edu/eric.budish/research/Azevedo-Budish-SPL.pdf>.

Benjamin, Daniel J., Ori Heffetz, Miles S. Kimball, and Nichole Szembrot, “Aggregating Local Preferences to Guide Marginal Policy Adjustments,” *American Economic Association Papers & Proceedings*, 2013, 103 (3).

Berry, Andrew C., “The Accuracy of the Gaussian Approximation to the Sum of Independent Variates,” *Transactions of the American Mathematical Society*, 1941, 49 (1), 122–136.

Berry, Steven T., James Levinsohn, and Ariel Pakes, “Automobile Prices in Market Equilibrium,” *Econometrica*, 1995, 63 (4), 841–890.

Bewley, Truman F., “The Permanent Income Hypothesis: A Theoretical Formulation,” *Journal of Economic Theory*, 1977, 16 (2), 252–292.

Bierbrauer, Felix J. and Martin F. Hellwig, “Mechanism Design and Voting for Public-Good Provision,” 2011. Mimeo, Max Planck Institute for Research on Collective Goods.

Bowen, Howard R., “The Interpretation of Voting in the Allocation of Economic Resources,” *Quarterly Journal of Economics*, 1943, 58 (1), 27–48.

Casella, Alessandra, “Storable Votes,” *Games and Economic Behavior*, 2005, 51 (2), 391–419.

– **and Sebastien Turban**, “Democracy Undone - Systematic Minority Advantage in Competitive Vote Markets,” 2012. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2183033.

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

- Clarke, Edward H.**, “Multipart Pricing of Public Goods,” *Public Choice*, 1971, 11 (1), 17–33.
- Crémer, Jacques and Richard P. McLean**, “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 1988, 56 (6), 1247–1257.
- d’Aspremont, Claude and Louis-André Gérard-Varet**, “Incentives and Incomplete Information,” *Journal of Public Economics*, 1979, 11 (1), 25–45.
- Dupuit, Arsène Jules Étienne Juvénal**, *De la Mesure de L’utilité des Travaux Publics*, Paris, 1844.
- Erdil, Aytok and Paul Klemperer**, “A New Payment Rule for Core-Selecting Package Auctions,” *Journal of European Economic Association*, 2010, 8 (2–3), 537–547.
- Esseen, Carl-Gustav**, “On the Liapunoff Limit of Error in the Theory of Probability,” *Arkiv för Matematik, Astronomi och Fysik*, 1942, (A28), 1–19.
- , “A Moment Inequality with an Application to the Central Limit Theorem,” *Skandinavisk Aktuarietidskrift*, 1956, (39), 160–170.
- Gabaix, Xavier**, “Power Laws in Economics and Finance,” *Annual Review of Economics*, 2009, 1 (1), 255–294.
- Gibbard, Alan**, “Manipulation of Voting Schemes: A General Result,” *Econometrica*, 1973, 41 (4), 587–602.
- Goeree, Jacob K. and Jingjing Zhang**, “Electoral Engineering: One Man, One Vote Bid,” 2012. <http://www.econ.uzh.ch/esei/Events/prague/Zhang.pdf>.
- Green, Jerry and Jean-Jacques Laffont**, “Imperfect Personal Information and the Demand Revealing Process: A Sampling Approach,” *Public Choice*, 1977, 29 (2), 79–94.
- Groves, Theodore**, “Incentives in Teams,” *Econometrica*, 1973, 41 (4), 617–631.
- Harberger, Arnold C.**, “The Measurement of Waste,” *American Economic Review*, 1964, 54 (3), 58–76.
- Heifetz, Aviad and Zvika Neeman**, “On the Generic (Im)Possibility of Full Surplus Extraction in Mechanism Design,” *Econometrica*, 2006, 74 (1), 213–233.
- Hicks, J. R.**, “The Foundations of Welfare Economics,” *Economic Journal*, 1939, 49 (196), 696–712.
- Hortala-Vallve, Rafael**, “Qualitative Voting,” *Journal of Theoretical Politics*, Forthcoming.

- Hu, Henry T. C. and Bernard S. Black**, “The New Vote Buying: Empty Voting and Hidden (Morphable) Ownership,” *Southern California Law Review*, 2006, 79, 811–908.
- Hylland, Aanund and Richard Zeckhauser**, “A mechanism for selecting public goods when preferences must be elicited,” *Kennedy School of Government Discussion Paper D*, 1979, 70.
- Jackson, Matthew O. and Hugo F. Sonnenschein**, “Overcoming Incentive Constraints by Linking Decisions,” *Econometrica*, 2007, 75 (1), 241–257.
- Jenkin, Henry Charles Fleeming**, “On the Principles which Regulate the Incidence of Taxes,” *Proceedings of the Royal Society of Edinburgh*, 1871–1872, pp. 618–631.
- Kaldor, Nicholas**, “Welfare Propositions in Economics and Interpersonal Comparisons of Utility,” *Economic Journal*, 1939, 49 (145), 549–552.
- Kominers, Scott Duke and E. Glen Weyl**, “Concordance among Holdouts,” 2012. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1591466.
- and –, “Holdout in the Assembly of Complements: A Problem for Market Design,” *American Economic Review Papers & Proceedings*, 2012, 102 (3), 360–365.
- Krishna, Vijay and John Morgan**, “Majority Rule and Utilitarian Welfare,” 2012. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2083248.
- Laffont, Jean-Jacques and David Martimort**, “Collusion under Asymmetric Information,” *Econometrica*, 1997, 65 (4), 875–911.
- and –, “Mechanism Design with Collusion and Correlation,” *Econometrica*, 2000, 68 (2), 309–342.
- Ledyard, John**, “The Pure Theory of Large Two-Candidate Elections,” *Public Choice*, 1984, 44 (1), 7–41.
- , “Public Goods: A Survey of Experimental Research,” in John Kagel and Alvin Roth, eds., *Handbook of Experimental Economics*, Princeton: Princeton University Press, 1995.
- Ledyard, John O. and Thomas R. Palfrey**, “Voting and Lottery Drafts as Efficient Public Goods Mechanisms,” *Review of Economic Studies*, 1994, 61 (2), 327–355.
- and –, “The Approximation of Efficient Public Good Mechanisms by Simple Voting Schemes,” *Journal of Public Economics*, 2002, 83 (2), 153–171.
- Maskin, Eric**, “Nash Equilibrium and Welfare Optimality,” *Review of Economic Studies*, 1999, 66 (1), 23–38.

- McAfee, R. Preston and Philip J. Reny**, “Correlated Information and Mechanism Design,” *Econometrica*, 1992, 60 (2), 395–421.
- Morris, Stephen and Hyun Song Shin**, “Global Games: Theory and Applications,” in Mathias Dewatripont, Lars Hansen, and Stephen Turnovsky, eds., *Advances in Economics and Econometrics, the Eight World Congress*, Cambridge, UK: Cambridge University Press, 2003.
- Myerson, Roger B.**, “Large Poisson Games,” *Journal of Economic Theory*, 2000, 94 (1), 7–45.
- Posner, Eric and E. Glen Weyl**, “Efficient Corporate Governance,” *University of Chicago Law Review*, Forthcoming.
- Saez, Emmanuel**, “Using Elasticities to Derive Optimal Income Tax Rates,” *The Review of Economic Studies*, 2001, 68 (1), 205–229.
- Satterthwaite, Mark Allen**, “Strategy-Proofness and Arrow’s Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions,” *Journal of Economic Theory*, 1975, 10 (2), 187–217.
- Shaked, Avner and John Sutton**, “Relaxing Price Competition Through Product Differentiation,” *Review of Economic Studies*, 1982, 49 (1), 3–13.
- Shleifer, Andrei and Robert W. Vishny**, “A Survey of Corporate Governance,” *A Survey of Corporate Governance*, 1997, 52 (2), 737–783.
- Taylor, Curtis R. and Huseyin Yildirim**, “A Unified Analysis of Rational Voting with Private Values and Group-Specific Costs,” *Games and Economic Behavior*, 2010, 70 (2), 457–471.
- Thompson, Earl A.**, “A Pareto-Efficient Group Decision Process,” *Papers on Non-Market Decision-Making*, 1966.
- Vickrey, William**, “Counterspeculation, Auctions and Competitive Sealed Tenders,” *Journal of Finance*, 1961, 16 (1), 8–37.
- Weyl, E. Glen**, “An Alternative to Democracy,” 2013. This paper is under preparation. Email me at weyl@uchicago.edu for notes.
- Willig, Robert D.**, “Consumer’s Surplus Without Apology,” *American Economic Review*, 1976, 66 (4), 589–597.
- Wilson, Robert**, “Game Theoretic Analysis of Trading Processes,” in Truman F. Bewley, ed., *Advances in Economic Theory*, Cambridge, UK: Cambridge University Press, 1987.

Appendix

A Proofs

Sketch of Proof of Lemma 2. Here I again provide a proof sketch, following closely the logic of the sketch of Lemma 1, and thus I include many fewer steps. I now approximate $v = a(n)u + b(n)u^2$, dropping the cubic term as its contribution is small relative to that of the quadratic term. For the equilibrium not to be dominated by the quadratic term, I would need that $\frac{b(n)}{a(n)} = O\left(\frac{1}{\sqrt{n}}\right)$. The distribution g now has mean

$$(n-1) \left[a(n) \frac{\hat{\mu}}{\sqrt{n}} + b(n) \left(\sigma^2 + \frac{\hat{\mu}^2}{n} \right) \right] \approx \sqrt{n}a(n)\hat{\mu} + nb(n)\sigma^2$$

and variance

$$(n-1) \left[a^2(n) \left(\sigma^2 + \frac{\hat{\mu}^2}{n} \right) + 2a(n)b(n)\mu_3 + b^2(n)\mu_4 - \frac{a^2(n)\hat{\mu}^2}{n} - \frac{2a(n)b(n)\hat{\mu}\sigma^2}{\sqrt{n}} - b^2(n)\sigma^2 \right] \approx na^2(n)\sigma^2,$$

using the fact that $\frac{b(n)}{a(n)} = O\left(\frac{1}{\sqrt{n}}\right)$. As above, I take a Taylor approximation to g , though only the first-order approximation is needed now:

$$\begin{aligned} g(-v) &= \frac{e^{-\frac{[\sqrt{n}a(n)\hat{\mu} + nb(n)\sigma^2]^2}{2a^2(n)(n-1)\sigma^2}}}{a(n)\sigma\sqrt{2\pi(n-1)}} \left[1 - \frac{\sqrt{n}a(n)\hat{\mu} + nb(n)\sigma^2}{2a^2(n)(n-1)\sigma^2} v \right] \implies \\ v &= \frac{e^{-\frac{[\sqrt{n}a(n)\hat{\mu} + nb(n)\sigma^2]^2}{2a^2(n)(n-1)\sigma^2}}}{2a(n)\sigma\sqrt{2\pi(n-1)}} \left[1 - \frac{\sqrt{n}a(n)\hat{\mu} + nb(n)\sigma^2}{a^2(n)(n-1)\sigma^2} (a(n)u + b(n)u^2) \right] u \approx \\ &\quad \frac{e^{-\frac{[\sqrt{n}a(n)\hat{\mu} + nb(n)\sigma^2]^2}{2a^2(n)(n-1)\sigma^2}}}{2a(n)\sigma\sqrt{2\pi(n-1)}} \left[u - \frac{\sqrt{n}a(n)\hat{\mu} + nb(n)\sigma^2}{a(n)n\sigma^2} u^2 \right]. \end{aligned}$$

Again matching coefficients I have

$$\frac{b(n)}{a(n)} = -\frac{\hat{\mu}}{\sqrt{n}\sigma^2} - \frac{b(n)}{a(n)} \implies \frac{b(n)}{a(n)} = -\frac{\hat{\mu}}{2\sqrt{n}\sigma^2}.$$

This simplifies the expression for the mean of g to $\frac{\sqrt{n}}{2}a(n)\hat{\mu}$ and allows us to more easily match coefficients to solve for $a(n)$:

$$a(n) = \frac{e^{-\frac{\hat{\mu}^2}{2\sigma^2}}}{2a(n)\sigma\sqrt{2\pi(n-1)}} \implies a(n) = \frac{e^{-\frac{\hat{\mu}^2}{4\sigma^2}}}{\sqrt{2\sigma}\sqrt[4]{2\pi(n-1)}}.$$

As in the previous section, the asymptotic perfect correlation between U and $U + \frac{b(n)}{a(n)}U_2$ holds. The mean of $U + \frac{b(n)}{a(n)}U_2$ is $\sqrt{n}\frac{\hat{\mu}}{2}$ while the mean of U is $\sqrt{n}\hat{\mu}$. First-best welfare is

$$E[|U|] = \sigma\sqrt{\frac{2n}{\pi}}e^{-\frac{\hat{\mu}^2}{2\sigma^2}} + \sqrt{n}\hat{\mu}\left[1 - 2\Phi\left(-\frac{\hat{\mu}}{\sigma}\right)\right],$$

where Φ is the cumulative distribution function of the standard normal distribution. On the other hand the loss of welfare relative to the first best in the equilibrium of QVB is

$$\begin{aligned} \int_0^{\frac{\sqrt{n}\hat{\mu}}{2}} \frac{U e^{-\frac{(U-\sqrt{n}\hat{\mu})^2}{2n\sigma^2}}}{\sqrt{n}\sigma\sqrt{2\pi}} dU &= -\sigma\sqrt{\frac{n}{2\pi}} \int_{-\frac{\hat{\mu}^2}{2\sigma^2}}^{-\frac{\hat{\mu}^2}{8\sigma^2}} e^{\hat{U}_2} d\hat{U}_2 + \hat{\mu} \int_0^{\frac{\sqrt{n}\hat{\mu}}{2}} \frac{e^{-\frac{(U-\sqrt{n}\hat{\mu})^2}{2n\sigma^2}}}{\sigma\sqrt{2\pi}} dU = \\ &\hat{\mu}\sqrt{n}\left[\Phi\left(\frac{\hat{\mu}}{\sigma}\right) - \Phi\left(\frac{\hat{\mu}}{2\sigma}\right)\right] - \sigma\sqrt{\frac{n}{2\pi}}\left(e^{-\frac{\hat{\mu}^2}{8\sigma^2}} - e^{-\frac{\hat{\mu}^2}{2\sigma^2}}\right). \end{aligned}$$

Thus, as indicated in the proposition statement, asymptotic inefficiency is

$$\frac{\hat{\mu}\sqrt{n}\left[\Phi\left(\frac{\hat{\mu}}{\sigma}\right) - \Phi\left(\frac{\hat{\mu}}{2\sigma}\right)\right] - \sigma\sqrt{\frac{n}{2\pi}}\left(e^{-\frac{\hat{\mu}^2}{8\sigma^2}} - e^{-\frac{\hat{\mu}^2}{2\sigma^2}}\right)}{\sigma\sqrt{\frac{2n}{\pi}}e^{-\frac{\hat{\mu}^2}{2\sigma^2}} + \sqrt{n}\hat{\mu}\left[1 - 2\Phi\left(-\frac{\hat{\mu}}{\sigma}\right)\right]}.$$

This expression is maximized when $x = \frac{\hat{\mu}}{\sigma} \approx 1.6$ at a value of .044. □

Sketch of the Proof of Lemma 3. The structure of equilibrium is very different in this case than in the above two, and thus I sketch its construction in greater detail here. Again WLOG suppose that $\mu > 0$. The equilibrium structure involves all individuals except for extremists, those with large negative values, following the *moderate strategy* of buying votes approximately linear in their value $v = a(n)u$, while individuals with large negative value buy a “large” number of votes, enough to ensure that A is not taken with probability greater than 50%. The complications in solving for the equilibrium arise for two reasons:

1. One must identify the set of extremists, which requires comparing the total utility of the extremist strategy to that of the moderate strategy and determining the type with sufficiently extreme value so that they are just indifferent between these two strategies. Note that the shape of the normal distribution resulting from the sum of all other individuals’ votes is such that there can only be two local maxima, so only these two need be compared.
2. One must identify the distribution of optimal votes bought by extremists and incorporate this into the calculations of optima for individuals following the moderate strategy.

To simplify this for purposes of exposition, here I use two approximating assumptions that do not alter the limiting behavior but simplify the calculations in the exposition below:

1. I assume the set of extremists is exactly those individuals who, if they bought enough votes to exactly tie the election (in expectation), would have a local incentive to purchase more votes. From Figure 4 one can see both that this approximation is incorrect, but also that it will not alter the order of magnitude in calculations. While clearly any individual whose $\frac{u}{v}$ ray just touches the top of the relevant curve will not want to pursue an extremist strategy, it can also easily be seen that the individual who will pursue such a strategy has a value magnitude that is a constant multiple larger. Such a constant multiple will not alter the limiting order of magnitude calculations here.

2. I assume that every individual following the extreme strategy buys precisely enough votes to to tie the election in expectation. That is, they buy enough to exactly cancel the mean number of votes purchased by all other individuals. Again this approximation is clearly not perfectly accurate and is inconsistent with the first approximation; every extremist will have an incentive to buy a greater number of votes than this by definition, again as shown in Figure 4. However, the typical extremist will bring the election to a density of pivotality that is only a constant multiple less than the peak of the normal distribution and the dramatic reduction in the chance of pivotality by the low probability of such an extremist existing makes the effective values of $\frac{g(-v)}{2}$ so much lower for most moderates that proportionality is an extremely accurate approximation, as shown in the left panel of 4. Thus neither the constant multiplicative change in the level of g nor the perfectly constant approximation used are likely to affect the order of magnitude of my estimates below. If anything, my results slightly overstate limiting inefficiency because the typical extremist buys enough votes to make the favorite side the underdog in the case of an extremist participating, raising the votes of moderates and thereby reducing the chance of extremism.

Under these approximations, the equilibrium is straightforward to compute; in a future draft I hope to improve or replace these approximations with precise calculations. Let p be the probability that any individual adopts the extremist strategy. Assuming np is small and thus that $n^2p^2 \ll np$ and $1 - (1 - p)^n \approx np$, g is a $np, 1 - np$ mixture of normal distributions with mean given by the moderate behavior of most individuals in the $1 - np$ case and mean 0 in the np case. The standard deviation conditional on each case is given by the moderates' behavior. These mean and standard deviation are respectively, ignoring the negligible subtraction by extremists, $a(n)\mu(n - 1)$ and $a(n)\sigma\sqrt{n - 1}$.

An extremist will not be motivated to buy the decision based on the case when it is already tied, as in this case the density decreases as she purchases more votes; this event is low probability as well. Instead, her motivation arises from the $1 - np$ cases in which the distribution is approximately normal with mean $a(n)\mu(n - 1)$ and $a(n)\sigma\sqrt{n - 1}$. In this case, the number of votes that must be purchased to, as discussed above, cause the election to again be tied in expectation is $a(n)\mu(n - 1)$.

The value of g when this is achieved is $\frac{1}{a(n)\sqrt{2\pi(n-1)}\sigma}$. Thus the magnitude of value such that an individual would have a local incentive to purchase more votes given that she is playing the extreme strategy is given by

$$\frac{\underline{u}}{a(n)\sqrt{2\pi(n-1)}\sigma} = -2a(n)\mu(n-1) \implies \underline{u} = -2a^2(n)(n-1)^{\frac{3}{2}}\sqrt{2\pi}\sigma\mu.$$

Let us assume that the lower tail of the distribution of values is well-approximated by a Pareto tail; the cases with finite first and second moments that are not so approximable have thin tails and behave like the Pareto case with the tail index $\alpha \rightarrow \infty$. In this case the probability of an individual with a value $u < -\underline{u}$ is, for any $-\underline{u}$, $\frac{k}{|\underline{u}|^\alpha}$ for some $k > 0$ and $\alpha > 2$; the latter requirement stems from the assumption of finite second-moment. Then

$$p = \frac{k}{\left(2a^2(n)(n-1)^{\frac{3}{2}}\sqrt{2\pi}\sigma\mu\right)^\alpha}. \quad (4)$$

On the other hand, those playing moderate strategies have almost no contribution to g near 0 from the $1 - np$ scenario, as in this case the density of a tie dies off with an exponential square. Thus the density of a tie is approximately

$$\frac{np}{a(n)\sqrt{2\pi(n-1)}\sigma} = \frac{nk}{\left(2a^2(n)(n-1)^{\frac{3}{2}}\sqrt{2\pi}\sigma\mu\right)^\alpha a(n)\sqrt{2\pi(n-1)}\sigma} \approx \frac{k}{a^{2\alpha+1}(n)n^{\frac{3\alpha-1}{2}}2^\alpha\sigma^{\alpha+1}(2\pi)^{\frac{\alpha+1}{2}}}.$$

Matching coefficients as above then requires

$$\frac{k}{2a^{2\alpha+1}(n)n^{\frac{3\alpha-1}{2}}2^\alpha\sigma^{\alpha+1}(2\pi)^{\frac{\alpha+1}{2}}} = a(n) \implies a^{2(\alpha+1)}(n) = \frac{k}{n^{\frac{3\alpha-1}{2}}2^\alpha\sigma^{\alpha+1}(2\pi)^{\frac{\alpha+1}{2}}} \implies$$

$$a(n) = \frac{k^{\frac{1}{2(\alpha+1)}}}{n^{\frac{3\alpha-1}{4(\alpha+1)}2^{\frac{\alpha}{2(\alpha+1)}}\sqrt{\sigma}\sqrt[4]{2\pi}}.$$

Plugging this into Equation (4) yields

$$p \approx \frac{k}{\left(2\frac{k^{\frac{1}{\alpha+1}}}{n^{\frac{3\alpha-1}{2(\alpha+1)}}2^{\frac{\alpha}{\alpha+1}}\sigma\sqrt{2\pi}}n^{\frac{3}{2}}\sqrt{2\pi}\sigma\mu\right)^\alpha} = \frac{k^{\frac{1}{1+\alpha}}}{\left(2\mu n^{\frac{2}{1+\alpha}}\right)^\alpha} = \frac{k^{\frac{1}{1+\alpha}}}{(2\mu)^\alpha n^{\frac{2\alpha}{1+\alpha}}}.$$

Thus

$$pn \approx \frac{k^{\frac{1}{1+\alpha}}}{(2\mu)^\alpha n^{\frac{\alpha-1}{\alpha+1}}}.$$

Inefficiency essentially occurs only in the pn cases, as the probability that the correct action is not

of the same sign of μ vanishes with a negative exponential. Thus inefficiency dies with $\frac{1}{n^{\frac{\alpha-1}{\alpha+1}}}$. Again, constants, uniformity of the results across possibly multiple equilibria, as well as corrections for the approximation error of the two approximation assumptions, are established in the appendix proof.¹⁷

□

Additional calculations related to the proof of Theorem 1. From Subsection 2.2 I have that welfare loss relative to the first best is approximately

$$\hat{\mu}\sqrt{n} \left[\Phi \left(\frac{\hat{\mu}}{\sigma} \right) - \Phi \left(\frac{\hat{\mu}}{2\sigma} \right) \right] - \sigma \sqrt{\frac{n}{2\pi}} \left(e^{-\frac{\hat{\mu}^2}{8\sigma^2}} - e^{-\frac{\hat{\mu}^2}{2\sigma^2}} \right)$$

for positive $\hat{\mu}$ and an equivalent size *mutatis mutandis* for negative $\hat{\mu}$. Thus total inefficiency from all values of $\hat{\mu}$ are approximately

$$2h(0) \int_{\hat{\mu}=0}^{\infty} \hat{\mu} \left[\Phi \left(\frac{\hat{\mu}}{\sigma} \right) - \Phi \left(\frac{\hat{\mu}}{2\sigma} \right) \right] - \sigma \sqrt{\frac{1}{2\pi}} \left(e^{-\frac{\hat{\mu}^2}{8\sigma^2}} - e^{-\frac{\hat{\mu}^2}{2\sigma^2}} \right) d\hat{\mu} =$$

$$2h(0) \int_{x=0}^{\infty} x \left[\Phi(x) - \Phi \left(\frac{x}{2} \right) \right] - \frac{e^{-\frac{x^2}{8}} - e^{-\frac{x^2}{2}}}{\sqrt{2\pi}} dx = \frac{h(0)}{2},$$

where $x = \frac{\hat{\mu}}{\sigma}$. Thus for large n the efficiency losses from this case are approximately $\frac{h(0)}{2}$, which is a constant and thus causes relative efficiency losses that die off at a rate of $\frac{1}{n}$. □

Sketch of proof of Proposition 1. First I consider the approximate distribution of u_M in the worst and average cases. In the average case, u_M is (for large m) approximately normal with mean $m\mu$ and variance $m\sigma^2$. In the worst case, when all of the collusive members are extremists in one direction, u_M is the sum of the first or last m order statistics of the value distribution. Without loss of generality, assume that the collusive group has negative value. Then for large n and small \underline{u} , the probability that the lowest order statistic is less than \underline{u} is less than $nk\underline{u}^{-\alpha}$, for some constants $\alpha > 1, k > 0$ using the Pareto tail approximation. By standard results from extreme value theory, as surveyed by Gabaix (2009), the j th-to-last order statistic is approximately $\frac{1}{j^{\frac{1}{\alpha}}}$ times the last order statistic. Thus, using the continuous approximation, the sum of the last m order statistics is approximately

$$\int_1^{m+1} \frac{1}{x^{\frac{1}{\alpha}}} dx = \frac{\alpha}{\alpha-1} \left[(m+1)^{\frac{\alpha-1}{\alpha}} - 1 \right]$$

¹⁷At one level, this equilibrium may appear a bit odd. There is a vanishingly small chance that some extremist buys the election, yet it is precisely this vanishing prospect that persuades most voters to buy a number of votes that vanishes only at a moderate rate. While this is certainly a bit odd, one can imagine a pattern not unlike this one emerging with behavioral individuals who systematically overestimate g : a few extremists emerge in each election and are tempered by a few extremists on the side of the mean. Both rationally spurred and irrationally angered by these extremists, most individuals buy more votes than rationally they should, ensuring that with high probability in a large population the outcome is efficient. Studying such a behavioral equilibrium is well beyond the scope of this paper.

times the first-order statistic. Thus for moderately large m and α not too close to 1, the probability that u_M is larger in absolute value than u is approximately

$$nk \left[\frac{u(\alpha - 1)}{\alpha m^{\frac{\alpha-1}{\alpha}}} \right]^{-\alpha} = nk \left[\frac{u(\alpha - 1)}{\alpha} \right]^{-\alpha} m^{\alpha-1}.$$

That is, a collusive group raises the probability of extreme values by a factor of order approximately $m^{\alpha-1}$.

On the other hand the typical value of u_m is $\frac{\alpha}{\alpha-1} m^{\frac{\alpha-1}{\alpha}}$ times the last order statistic. This last order statistic is of order $n^{\frac{1}{\alpha}}$ in n . Thus the typical size of u_M is of order $m^{\frac{\alpha-1}{\alpha}} n^{\frac{1}{\alpha}}$.

Now I consider the efficiency impact of groups with such value when they aim to implement extremist strategies and $\mu \neq 0$. In the average case, the size of u_M is approximately \sqrt{m} . The probability that such a group finds it attractive to pursue extremism is the probability that u_M is less than $\frac{\underline{u}}{m}$, where \underline{u} is the threshold for extremism calculated in Subsection 2.3 above. From the calculations above, this probability is of order (neglecting constant factors) $n\underline{u}^{-\alpha} m^{2\alpha-1}$. The order of \underline{u} calculated in Subsection 2.3 is that of

$$a^2(n)n^{\frac{3}{2}} = n^{-\frac{3\alpha-1}{2(\alpha+1)}} n^{\frac{3}{2}} = n^{\frac{3\alpha+3-3\alpha+1}{2(\alpha+1)}} = n^{\frac{2}{\alpha+1}}.$$

In the average case, therefore, for collusion to have any chance of leading to an extremist strategy, \sqrt{m} must be of the order $\frac{n^{\frac{2}{\alpha+1}}}{m}$ or m must be of order $n^{\frac{4}{3(\alpha+1)}}$.

In the worst case, the order of the probability of the collusive group finding an extremist strategy attractive is

$$n^{\frac{\alpha+1-2\alpha}{\alpha+1}} m^{2\alpha-1} = n^{-\frac{\alpha-1}{\alpha+1}} m^{2\alpha-1}.$$

Thus, depending on the size of m relative to n , the rate of decay of inefficiency changes from $O\left(n^{-\frac{\alpha-1}{\alpha+1}}\right)$ to $O\left(n^{-\frac{\alpha-1}{\alpha+1}} m^{2\alpha-1}\right)$. Inefficiency still dies in the limit as long as m is $O\left(n^{\frac{\alpha-1}{(\alpha+1)(2\alpha-1)}}\right)$. I interpret this result more fully below and here move onto analyzing the non-extremist case.

In the average case, u_M is of order \sqrt{m} when $\mu = 0$. Collusion overweights the values of this group by a factor of m and the efficiency cost associated with this overweighting is bounded above by $(m-1)\sqrt{m}$. Given that the total welfare from the first-best is on the order of \sqrt{n} in this case, the relative efficiency losses from collusion will be small as long as $m^{\frac{3}{2}}$ is $O(\sqrt{n})$ or $m = O(\sqrt[3]{n})$. More generally collusion in the average case will slow the rate of decay of inefficiency to $\frac{m^{\frac{3}{2}}}{\sqrt{n}}$.

On the other hand, in the worst case, u_M is of order $m^{\frac{\alpha-1}{\alpha}} n^{\frac{1}{\alpha}}$. The inefficiency this creates is $O\left(m^{\frac{2\alpha-1}{\alpha}} n^{\frac{1}{\alpha}}\right)$ by the logic above or relative inefficiency that is $O\left(m^{\frac{2\alpha-1}{\alpha}} n^{\frac{2-\alpha}{2\alpha}}\right)$. This declines in n if $m = O\left(n^{\frac{\alpha-2}{2(2\alpha-1)}}\right)$. \square

Sketch of proof of Proposition 2. These results follow directly from the logic of the preceding proof (of Proposition 4.2) except it is assumed that the size of the value u_M is, in the average case, a

constant and in the worst case is exactly that of the most extreme individual (rather than larger than this). Then l is substituted for m and the derivation follows precisely the same argument. \square

B Calculations for Examples of Aggregate Uncertainty

I first provide a counter example to the claim that $E[ux(u|\gamma)] = E[ux^2(u|\gamma)] = 0$ implies $E[u] = 0$, the identification of efficiency with ex-ante expected welfare ties. Suppose that u can take on three values ($-18, 3$ and 24) and that $x(u|\gamma)$ takes on values $(1, 2, \frac{1}{2})$ in perfect correlation with these values of u respectively. Each of these events occurs with probability $\frac{1}{3}$. $x(u|\gamma)$ is strictly positive, the only restriction that can be placed on it ex-ante.

$$E[ux(u|\gamma)] = \frac{-18 + \frac{24}{2} + 3 \cdot 2}{3} = 0 = \frac{-18 + \frac{24}{2^2} + 3 \cdot 2^2}{3} \neq \frac{-18 + 24 + 3}{3} = 3.$$

This example can easily be smoothed by making the distribution of u a normal about these points; this clearly contradicts the general basis of the Underdog Principle, though by creating a bit of an odd joint structure on u and x that appears to leave room for a fairly general set of conditions under which this principle would hold.

Proof of claims in Example 1. In this case $x(u|\gamma_0)$ is symmetric about 0 and thus $\gamma^* = \gamma_0$. Unfortunately this is not a particularly interesting case as in it majoritarian voting also yields efficiency because whenever $\gamma > (<)\gamma_0$ both the mean and median of the conditional distribution of u are positive (negative). \square

Proof of claims in Example 2. In the case described u has marginal distribution $N(\mu, \sigma_1^2 + \sigma_2^2)$. Thus $f(u)$ is, up to multiplicative constants, $e^{-\frac{(u-\mu)^2}{2(\sigma_1^2 + \sigma_2^2)}}$. Conditional on γ , u has distribution $N(\gamma, \sigma_1^2)$ so $f(u|\gamma)$ is (again up to multiplicative constants) $e^{-\frac{(u-\gamma)^2}{2\sigma_1^2}}$. γ^* is defined by satisfying

$$\int_u u \frac{f^2(u|\gamma^*)}{f(u)} du = 0 \iff \int_u u e^{\frac{(u-\mu)^2}{2(\sigma_1^2 + \sigma_2^2)} - \frac{(u-\gamma^*)^2}{\sigma_1^2}} = 0 \iff \frac{(u-\mu)^2}{2(\sigma_1^2 + \sigma_2^2)} - \frac{(u-\gamma^*)^2}{\sigma_1^2} = au^2 + b$$

for some constants a and b that are independent of u , as this is the only time the quadratic form gives the symmetry about 0 necessary and sufficient for the expression to have 0 mean (integral).

$$\frac{(u-\mu)^2}{2(\sigma_1^2 + \sigma_2^2)} - \frac{(u-\gamma^*)^2}{\sigma_1^2} = \frac{\sigma_1^2(u-\mu)^2 - 2(\sigma_1^2 + \sigma_2^2)(u-\gamma^*)^2}{2(\sigma_1^2 + \sigma_2^2)\sigma_1^2} = au^2 + b - 2\frac{\sigma_1^2\mu - 2(\sigma_1^2 + \sigma_2^2)\gamma^*}{2(\sigma_1^2 + \sigma_2^2)\sigma_1^2}u.$$

Thus γ^* must solve

$$2\frac{\sigma_1^2\mu - 2(\sigma_1^2 + \sigma_2^2)\gamma^*}{2(\sigma_1^2 + \sigma_2^2)\sigma_1^2} = 0 \iff \gamma^* = \frac{\sigma_1^2}{2(\sigma_1^2 + \sigma_2^2)}\mu.$$

Potential efficiency in the limit all arises from making the decision in line with the sign of γ . The welfare of the first best is thus

$$nE[|\gamma|] = n \left(\sigma_2 \sqrt{\frac{2}{\pi}} e^{-\frac{\mu^2}{2\sigma_2^2}} + \mu \left[1 - 2\Phi \left(-\frac{\mu}{\sigma_2} \right) \right] \right).$$

Welfare is clearly higher the closer γ^* is to 0 and the farthest γ^* can be from 0 is $\frac{\mu}{2}$ which is achieved as $\sigma_1^2 \rightarrow \infty$. The inefficiency created in this case is exactly as in Lemma 2 replacing $\hat{\mu}$ with μ and σ with σ_2 . \square

Proof of claims in Example 3. Average per-person welfare for taking the in this case is

$$\gamma\mu_+ - (1 - \gamma)\mu_-.$$

γ_0 is thus $\frac{\mu_-}{\mu_+ + \mu_-}$. However $x(u|\gamma) = \frac{\gamma h(\gamma)}{E[\gamma]}$ if u is positive, and $\frac{(1-\gamma)h(\gamma)}{1-E[\gamma]}$ if u is negative. Thus γ^* is defined by

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

where $k = \frac{\mu_+(1-E[\gamma])}{\mu_-E[\gamma]}$. This in turn implies

$$\gamma^2(k-1) + 2\gamma - 1 = 0 \iff \gamma = \frac{-2 \pm \sqrt{4 + 4(k-1)}}{2(k-1)} = \frac{-1 \pm \sqrt{k}}{k-1}.$$

The solution for γ^* must be between 0 and 1 and thus must be positive; the $\sqrt{k} - 1$ has the same sign as $k - 1$ and $-1 - \sqrt{k}$ does not for $k > 1$ so the positive solution must be correct and

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

Efficiency results if $\gamma^* = \gamma_0$, that is if

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

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

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

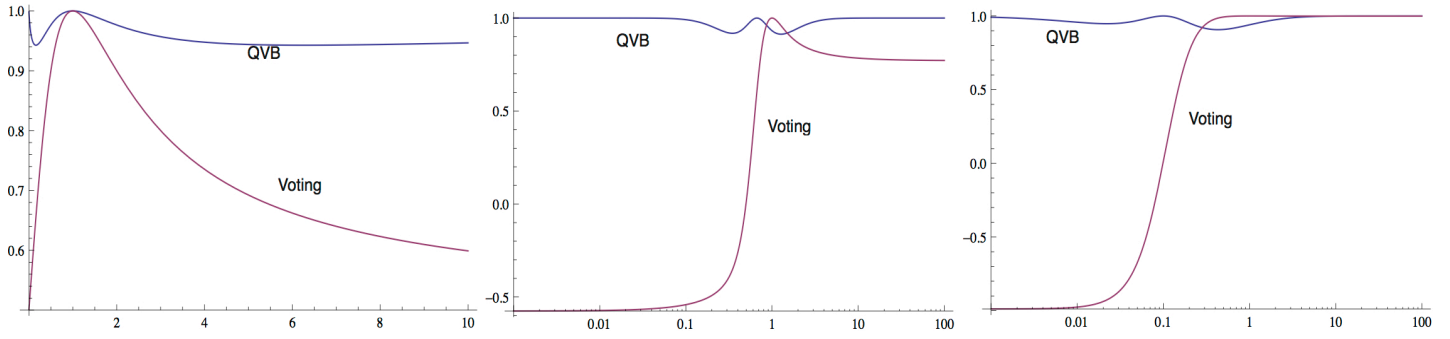


Figure 5: Relative-to-first-best efficiency of QVB compared to voting in the Krishna and Morgan (2012) example with γ distributed uniform (left), $\beta(15, 10)$ (center) and $\beta(10, 1)$ (right). The x-axis, measuring $\nu = \frac{\mu_+}{\mu_-}$ is on a log-scale (though labeled linearly) for the center and right graphs.

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

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

the Underdog Principle.

Under voting, the threshold value of γ for implementing the action is $\frac{1}{2}$. For each regime (the first-best, QVB and voting) I compute welfare as

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

where γ_t is the appropriate threshold value of γ . Using this method and explicit integration on Mathematica, I computed the relative (to the first best) efficiency of QVB and voting assuming h is a Beta distribution. Note that, if one divides the numerator and denominator by μ_- ,

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

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

so that the results for relative efficiency depend only on the ratio $\nu \equiv \frac{\mu_+}{\mu_-}$ and not on the level of μ_- . Furthermore, by symmetry of the Beta distribution in switching signs and simultaneously interchanging α and β , as long as I consider the full range of values for ν , it is WLOG to assume $\beta < \alpha$.

Figure 5 shows three examples that are representative of the more than 100 cases I experimented with. Whenever $\alpha = \beta$ (the distribution of γ is symmetric), QVB always dominates Voting as it does in the left panel shown, which is $\alpha = \beta = 1$, the uniform distribution. Voting obviously

performs best when ν , shown on the horizontal axis, is near to unity. When α is larger than β voting may out-perform QVB near $\nu = 1$. 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 QVB. However, it is precisely in these cases where, if ν is very small, voting is most dramatically inefficient. Intuitively voting may outperform QVB by blindly favoring the majority which is almost always in favor of taking the action for $\alpha \gg \beta$, while QVB may be a bit too conservative in favoring the action because of the Underdog Principle. However this blind favoritism towards the majority view can be highly destructive under voting, but not under QVB, 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 I have sampled) QVB's efficiency is above 90% and usually it is well above this.

□