Published in final edited form as:

Public Choice. 2017 July ; 172(1-2): 233–263. doi:10.1007/s11127-017-0414-3.

The Relationship Between the Normalized Gradient Addition Mechanism and Quadratic Voting

Daniel Benjamin,

Center for Economic and Social Research, University of Southern California, Los Angeles, CA 90089, USA

Ori Heffetz,

Samuel Curtis Johnson Graduate School of Management, Cornell University, Ithaca, NY 14853, USA. Department of Economics, Hebrew University of Jerusalem, 9190501, IL

Miles Kimball, and

Department of Economics, University of Colorado, Boulder, CO 80309, USA. Department of Economics, University of Michigan, Ann Arbor, MI 48109, USA. Center for Economic and Social Research, University of Southern California, Los Angeles, CA 90089, USA

Derek Lougee

Department of Economics, Cornell University, Ithaca, NY 14853, USA

Abstract

Quadratic Voting and the Normalized Gradient Addition mechanism are both social choice mechanisms that confront individuals with quadratic budget constraints, but they are applicable in different contexts. Adapting one or both to apply to the same context, this paper explores the relationship between these two mechanisms in three contexts: marginal adjustments of continuous policies, simultaneous voting on many public choices, and voting on a single public choice accompanied by private monetary consequences. In the process, we provide some formal analysis of Quadratic Voting when (instead of money) votes are paid for with abstract tokens that are equally distributed by the mechanism designer.

Keywords

Public Goods; Normalized Gradient Addition Mechanism; Quadratic Voting

1. Introduction

Finding a collective choice mechanism that generates the efficient levels of public goods remains a central challenge in economics. A key set of objectives is eliciting and aggregating people's preferences into a Pareto efficient and equitable social choice despite their possible incentives for strategic misreporting of their preferences. A number of impossibility theorems formalize the reality that in general, perfection cannot be achieved in

accomplishing all of these objectives at once (see, for example, Arrow 1950; Gibbard 1973; Satterthwaite 1975). But the practical need for collective choices motivates a continued search for mechanisms that have acceptable properties in particular contexts.

Recent work has analyzed two voting mechanisms that hold some such promise. In the context of choosing whether or not to pay for a (discrete) public good, quadratic voting (hereafter, QV) is a mechanism in which people can buy votes for or against the public good, and the net vote count determines the likelihood that the public good is produced (Weyl 2015, 2016; Lalley and Weyl 2016). In the context of a social choice about how to adjust the levels of a number of (continuous) policy parameters, normalized gradient addition (hereafter, NGA) is a mechanism in which people express a preference for the vector direction of change in the levels of policy parameters—effectively, a multidimensional vote on policy adjustments—and the sum of these multidimensional votes determines the adjustment of the policy parameters (Hylland and Zeckhauser 1980; Benjamin et al. 2013; Chung and Duggan 2015; Benjamin et al. 2016).

Despite the distinct contexts of the two mechanisms, they share two striking similarities. First, both mechanisms make the social choice according to a simple sum of votes, and this property is central to the mechanisms' efficiency properties. Second, both impose a quadratic budget constraint on voters. In QV, individuals can buy as many votes as they want, but the cost of votes is quadratic in the number of votes they buy. In NGA, choosing a vector direction of change is equivalent to choosing a point on a sphere, which implicitly imposes a multidimensional quadratic budget constraint. In both mechanisms, the quadratic budget constraint is central to their strategic properties.

In this paper, we explore the relationship between QV and NGA. Doing so is not straightforward because of their different contexts. Therefore, we first compare NGA with a modification of QV that adapts it to the NGA context and later compare QV with an extension of NGA that adapts it to the QV context. In between, we compare NGA and QV in the context of referenda on multiple public goods where abstract tokens, rather than money, are used to buy votes—a context in which both NGA and QV can be applied after relatively modest modifications. We view these exercises as exploratory: we study modifications of the mechanisms that we believe are instructive rather than necessarily practical; we propose conjectures that we hope will inspire and stimulate more work rather than only those we can prove; and, importantly, we stretch ourselves beyond our comfort zones, driven more by intellectual curiosity than by strongly held convictions regarding the relative strengths and weaknesses of the two mechanisms. More generally, we view the contribution of this paper as asking questions to enrich the discussion rather than providing authoritative answers.

We conclude on the basis of the three exercises above that when adapted to apply to the same context (choosing modifications we argue are reasonable), the two mechanisms either coincide or closely approximate each other.

Since unlike QV, NGA is framed as a direct mechanism in which truthful revelation and misreporting can occur, one feature of our exploration is that it provides a way of discussing truthful revelation in contexts that are relevant to QV, even though truthful revelation is not

directly applicable to QV. On the other hand, our adaptation of NGA to the QV context—deciding on a single public good—is not equivalent to QV, and the differences indicate some of the difficulties in getting a direct mechanism in the spirit of QV, but with a clearly specified notion of truthful revelation, to work in that context.

In addition, we view the adaptations of the QV mechanism to fit the NGA context and the multiple referenda context not only as a means of comparing QV and NGA, but also as a modest contribution that is of interest in itself. Because of concerns that QV favors the rich (Laurence and Sher 2016), that in practice the mechanism will not generate efficient outcomes if people mistakenly spend the wrong amount of money on voting (Kaplow and Kominers 2016), and that implementing it with actual monetary vote buying may often be impractical, there has been interest in implementations of QV in which people buy votes with tokens (instead of money) that are equally allocated by the mechanism designer. Moreover, the experiments with QV to date that we are aware of (Cárdenas et al. 2014; Goeree and Zhang 2016; Quarfoot et al. 2016) have been conducted with such equally allocated (though at times monetarily valuable) tokens. As far as we are aware, however, there has been no prior formal analysis of QV with tokens. Our adaptations of QV constitute initial steps in this direction. While the use of tokens can mitigate some of these concerns, when modified to apply to the same context, NGA and QV are subject to rather similar critiques.

The adaptations of the NGA mechanism to fit the QV context and the multiple referenda context require considering a policy space with boundaries since a discrete choice over a public good, even after allowing for a probabilistic implementation, involves a probability that cannot be smaller than 0 or larger than 1. While NGA was originally defined assuming an unrestricted policy space, Benjamin et al. (2016) define an extension to a convex policy space with boundaries given by inequality restrictions. This paper is the first to study applications of NGA with boundaries.

The rest of the paper proceeds as follows. In sections 2 and 3, we briefly sketch the QV and NGA mechanisms. Section 4 studies the relationship between NGA and an adaptation of QV to the NGA context. Section 5 studies the relationship between NGA and QV when both have been adapted to the multiple referenda context. Section 6 studies the relationship of QV with an adaptation of NGA to the QV context. Section 7 offers concluding remarks.

2. The QV Mechanism

The setup for QV is as follows:

- A single (discrete) public good can be produced or not. Without loss of generality, the cost of producing it can be treated as if it were zero by bundling the decision to produce the public good together with a default set of assessments on individuals to pay for it (before any of the side payments specified by the mechanism proper). (In section 6, this default set of assessments is made explicit and has each individual paying an equal amount for producing the public good.)

- There are Θ individuals. The utility function of each individual $\theta = 1, \dots, \Theta$ is $w_\theta + u_\theta$ if the public good is produced and w_θ if it is not produced, where w_θ is θ 's financial wealth and $u_\theta \in [\underline{u}, \bar{u}]$ is her net benefit (positive or negative) from the public good (including her part of its costs).

Producing the public good is efficient if and only if $\sum_\theta u_\theta > 0$.

In the QV mechanism, there is a referendum on whether or not to produce the public good. Each individual can buy any number of votes $v_\theta \in \mathbb{R}$, where positive votes are interpreted as being in favor and negative as opposed. The key feature of QV is that the cost of buying votes, $v_\theta^2/2$, is quadratic in the number purchased. The money collected through vote buying is redistributed in a way that yields a balanced budget for the mechanism while maintaining independence between an individual's own refund and the amount she chooses to spend on buying votes. Therefore, an individual θ 's refund can be written $t_\theta(v_{-\theta})$. The public good is then provided with probability $\Psi(V)$, where $V = \sum_\theta v_\theta$ is the net vote in favor of the public good. By assumption, Ψ is a cumulative distribution function with the property that there is some $\sigma > 0$ such that $\Psi(-\sigma) = 0$ and $\Psi(\sigma) = 1$, and $\Psi(V)$ is strictly increasing for all $V \in (-\sigma, \sigma)$. Lalley and Weyl (2016) show that under reasonable assumptions on $\Psi(V)$ and on the distribution of valuations, as the population gets large, $\Psi(V)$ converges in probability to 1 when producing the public good is efficient and to 0 otherwise.

Under this mechanism, each individual θ 's expected utility is

$$U_\theta = u_\theta [\Psi(V)] - \frac{1}{2} v_\theta^2 + [t_\theta(v_{-\theta})]. \quad (1)$$

Taking the derivative of U_θ with respect to v_θ and setting it equal to 0 indicates that the optimal number of votes to buy is given by $v_\theta = \mathbb{E}[\Psi'(V)] u_\theta$. Lalley and Weyl (2016) show that in a sufficiently large population, $\mathbb{E}[\Psi'(V)]$ is approximately unaffected by an individual's choice of v_θ (that is, it is approximately a constant from an individual's point of view). Therefore, the quadratic cost of votes leads to an important strategic property: most individuals set v_θ to be proportional to u_θ . Formally, in the limit as the number of agents Θ tends to infinity, most agents' equilibrium strategies converge to a linear voting strategy. That is, $v_\theta \approx p_\theta u_\theta$ for all $u_\theta \in [\underline{u} + 1/\Theta, \bar{u}]$, where $p_\theta > 0$ is the constant of proportionality (and is a decreasing function of Θ). Consequently, $\text{sign}(\sum_\theta v_\theta) \approx \text{sign}(\sum_\theta u_\theta)$. Thus, when providing the public good is efficient, the QV mechanism does so with probability approaching 1 as the population size grows to infinity.

3. The NGA Mechanism

The setup for NGA is as follows:

- There are $P > 1$ continuous public policy parameters whose initial levels are given by the status-quo policy vector $x_0 \in \mathbb{R}^P$.

- There are Θ individuals. The utility function of individual $\theta = 1, \dots, \Theta$ is $U_\theta(\mathbf{x})$, where $\mathbf{x}$ is the vector of levels of the policy parameters. Each function $U_\theta(\cdot)$ is continuously differentiable and weakly concave.

A policy vector $\mathbf{x}$ is defined to be Pareto efficient if there is no alternative policy vector $\hat{\mathbf{x}}$ such that $U_\theta(\hat{\mathbf{x}}) \geq U_\theta(\mathbf{x})$ for all individuals θ , with at least one inequality strict.

The NGA mechanism is a voting mechanism for generating a marginal adjustment to the status-quo policy vector $\mathbf{x}_0$. Each individual votes for a direction of policy change—which may or may not be equal to the direction that she actually most prefers. Formally, each individual reports a gradient that is normalized to have length 1, denoted $\mathbf{r}_\theta$. The policy adjustment yielded by the mechanism is proportional to the sum of the reported normalized gradients divided by the number of individuals: $\Delta \mathbf{x} = \frac{\varepsilon}{\Theta} \sum_{\theta} \mathbf{r}_\theta$, where $\varepsilon > 0$ is an exogenous “step size” parameter that describes the largest policy change that the mechanism could produce (if all individuals had the same $\mathbf{r}_\theta$).

The NGA mechanism has a basic one-shot version (described above), and an iterated version in which the above procedure is repeated iteratively. Since the iterated NGA is less closely related to QV, below we discuss it only briefly.

As long as $\varepsilon > 0$ is sufficiently small, the basic one-shot NGA mechanism is approximately strategy-proof: each individual votes for a direction that is approximately the direction she actually most prefers. The reason is that, if ε is small, then the most-preferred direction from the policy vector after everyone else has voted, $\mathbf{x}_0 + \frac{\varepsilon}{\Theta-1} \sum_{-\theta} \mathbf{r}_{-\theta}$, is approximately the same as at $\mathbf{x}_0$. Hence, by truthfully reporting the most-preferred direction from $\mathbf{x}_0$, an individual can push the overall policy change in the approximate direction she prefers regardless of what everyone else reports.

Furthermore, as Benjamin et al. (2013) show, the policy change prescribed by the NGA mechanism has an efficiency property: it has a positive dot product with every direction of Pareto improvement. The positive dot product with all Pareto improvements means that NGA tends to move toward Pareto improvements. More generally, even for changes that are not Pareto improvements, NGA will recommend moving in a direction of policy changes that (despite making some people worse off) makes a large majority of people better off in a way that is a high priority for them. Formally, as shown in Benjamin et al. (2016), let $\alpha_\theta(\mathbf{v})$ be the angle between a given vector $\mathbf{v}$ and the gradient for individual θ ; if $\sum_{\theta} \cos(\alpha_\theta(\mathbf{v})) > 0$, then NGA moves in a direction that has a positive dot product with the vector $\mathbf{v}$. While this inequality can hold under much weaker conditions, it clearly holds if $\mathbf{v}$ makes a large majority of people better off in a way that is a high priority for them.

The normalization of the gradients is central to the approximate strategy-proofness: if instead individuals could have a bigger impact on the policy change by reporting longer vectors, they would have an incentive to exaggerate the strength of their preferences. The vote being constrained to have length 1 imposes a quadratic cost of voting: in NGA, individuals can be thought of as having to choose a point on the unit sphere, so that the sum of squared votes across policy dimensions has to sum to 1.

The reason a small step size ϵ is important for approximate strategy-proofness is that preferences may be nonlinear. However, the applications we discuss in this paper, which involve discrete public goods, have linear preferences: preferences over money are assumed to be quasilinear, and preferences over probabilistically producing a public good are linear due to expected utility theory. For this reason, we do not consider small ϵ in these applications to be a defining feature of NGA per se. But as we will discuss, imposing boundaries on the policy space creates a distinct reason why approximate strategy-proofness requires a small enough step size, as well as a reason why efficiency requires a large enough step size. In section 5, those two requirements are not compatible, while in section 6 they can both be accommodated if the population is large.

Benjamin et al. (2013) emphasize the potential application of the NGA mechanism in connection with the expressed desire by a growing number of governmental agencies to integrate subjective well-being data (such as data on happiness, life satisfaction and mental health) into policy formation (see Helliwell et al. 2015; OECD 2016). A growing literature estimates the effects of various policies on different measures of subjective well-being (Finkelstein et al. 2012; Ludwig et al. 2012). If economists had a trusted method of measuring well-being through surveys (see Benjamin et al. 2014, for one potential direction), and if estimates of the effects of a set of different policy changes on the well-being measure were collected, then the NGA mechanism could constitute a method of using such estimates for recommending policy changes. From econometric estimates of the effects of various policy changes on the welfare of different groups in the population, one could construct an estimate of the utility gradient for each group. These gradients can then be normalized and summed across groups (using NGA) to yield a recommended policy change. Because the would-be “votes” produced by this procedure to replace actual votes are constructed from estimates of policy effects, they may generate desirable social outcomes even when individuals are not well informed about the effects of the policies. However, we would not advocate replacing actual voting unless and until safeguards are in place to prevent potential abuses and the procedure has passed extensive empirical testing. One approach is to limit this would-be-voting approach to providing recommendations as inputs into traditional political processes. From that point of view, the hope is that NGA in conjunction with subjective well-being research (and perhaps, in the future, also with research on objective health markers, bio-markers, and other well-being markers) could identify non-obvious combinations of policy changes that are jointly attractive.

In practice, it is likely that the NGA mechanism would be used iteratively. Benjamin et al. (2016) study the properties of the dynamic path of the iterated NGA mechanism. Hylland and Zeckhauser (1980), which was considerably prior to Benjamin et al. (2013), and Chung and Duggan (2015), which was contemporaneous with and independent of Benjamin et al. (2013), can be understood as studying the properties of stationary points of the iterated NGA mechanism—that is, points for which the normalized gradients sum to 0. The most important two properties of the set of stationary points—both proved by Hylland and Zeckhauser (1980)—are the guaranteed existence of at least one such stationary point when each function $U_{\theta}(\cdot)$ has a global maximum and the Pareto optimality of each stationary point.

Although the NGA mechanism is approximately strategy-proof at each iteration when the step size ϵ is small, the dynamic path of the iterated mechanism could be strategically manipulated if individuals are non-myopic and have nonlinear preferences. With linear preferences, the unique subgame perfect equilibrium of the iterated NGA mechanism is equivalent to the one-shot mechanism with a larger step size. However, extending the mechanism to have boundaries, as we will do in subsequent sections, introduces new possibilities for strategic manipulation despite linear preferences.

4. Comparing NGA and QV Adapted to the NGA Context

In this section, we consider how QV could naturally be adapted to the NGA context, and argue that the most natural adaptation may be NGA itself.

The NGA context differs from the QV context in two main ways. First, unlike in QV, where individuals buy votes with money and the initial allocation of money may be unequal, in NGA individuals are given an equal number of votes. Second, QV as described by Lally and Weyl (2016) focuses on whether or not to produce a single public good, while NGA studies how to adjust the levels of multiple, continuous public policy parameters (which could be levels of public goods).

To adapt QV to resemble NGA on the first point, it is natural to give individuals an equal number of abstract tokens (that exist only within the mechanism) as their vote-buying budget. But if there were only one policy dimension, then it would be optimal for each individual (unless she is indifferent to the result) to spend all her tokens voting either for or against the single policy change, in which case QV with tokens is simply equivalent to majority voting. Switching from money to tokens is therefore complementary with modifying QV on the second point: adapting it from one to multiple policy dimensions, so that participants face a tradeoff between expending points to vote on one dimension versus another.

Formally, if casting $v_{\theta,i} \in \mathbb{R}$ votes on policy dimension i has a quadratic cost in tokens for each policy dimension $i = 1, 2, \dots, P$, then the participant's vote budget constraint can be written as

$$v_{\theta,1}^2 + v_{\theta,2}^2 + \dots + v_{\theta,P}^2 \leq 1,$$

where the number of tokens given to each participant is normalized to 1. Assuming participants remain unsatiated on at least one dimension, the constraint holds with equality. In that case, the choice of a vote vector is equivalent to choosing a point on the unit sphere.

Moreover, in the NGA context in which the public policy parameters are continuous and participants cannot affect policy by spending money, it is natural to model utility as a direct function of the policy levels (a reduced-form representation of the effects of policies on the arguments of utility):

$$U_{\theta}(x_1, x_2, \dots, x_P).$$

Finally, consider the rule for how votes affect policy. In QV, the probability of producing the public good is a function of the net vote. To adapt the mechanism so that its output is an adjustment of the levels of the public policy parameters, a natural modification of the QV rule is to set the adjustment in each policy dimension proportional to the net vote on that dimension:

$$\Delta x_i \propto \sum_{\theta} v_{\theta,i}.$$

This variant of QV, however, now fully coincides with NGA as described in section 3: the vector $\mathbf{v}_{\theta}$ takes the place of $\mathbf{r}_{\theta}$ and the proportionality constant is $\frac{\epsilon}{\theta}$. From this point of view, we argue that NGA is the natural version of QV for the NGA context.

5. Comparing NGA and QV With Tokens in the Context of Multiple Referenda

In the previous section, we argued that if QV with tokens were modified in a natural way to the full context of NGA by applying it to votes about continuous adjustments of a policy parameter vector, it would coincide with NGA. However, the most obvious application of QV with tokens based on the internal logic of the QV literature itself is to the case of multiple referenda over discrete public goods choices. We consider that context in this section. To keep things simple, we consider here the case in which preferences have an expected utility representation that is additively separable across the public goods, and the referenda are conducted simultaneously.

As discussed in the previous section, a quadratic cost of votes expressed in tokens is meaningful only if using the tokens has an opportunity cost. For tokens that only exist within a collective choice mechanism, the opportunity cost of using them for voting on one choice dimension is their foregone use on other dimensions. The simplest case is when each individual is exogenously given an equal number of tokens (normalized to 1) that can be distributed to vote for or against each of P public goods, $i = 1, 2, \dots, P$. Formally, each individual buys votes to maximize her utility subject to her budget constraint,

$v_{\theta,1}^2 + v_{\theta,2}^2 + \dots + v_{\theta,P}^2 = 1$. Votes are then aggregated on each dimension, $V_i = \sum_{\theta} v_{\theta,i}$. For the purposes of discussing the QV mechanism in this context, we simplify it (while retaining its essential structure) by assuming that as long as the net vote is nonzero, the outcome is deterministic: $1(V_i > 0)$, meaning that public good i is built if the net vote for i is positive. If the net vote is exactly 0, then the public good is produced with probability 50%.

5.1. NGA with Boundaries

To study the case of $P > 1$ discrete public goods, we need to define an extension of the NGA mechanism that applies to that context. We do so in two steps. First, we interpret the

continuous policy parameter in dimension i to be the probability of producing the i^{th} public good. Second, to guarantee that the policy parameters lie inside $[0, 1]$, as probabilities must, we follow Benjamin et al. (2016) in extending the NGA mechanism to the case where the policy space has a boundary in at least one direction.

Formally, the *restricted policy space* R is defined as $R = \{x \in \mathbb{R}^P: G_j(x) \leq 0, j = 1, \dots, M\}$, where each G_j is a quasi-concave, continuously differentiable function, $G_j: \mathbb{R}^P \rightarrow \mathbb{R}$, and M is the number of constraints (which may be larger or smaller than P). We then define the restricted NGA mechanism as follows.

Definition 1—The restricted NGA mechanism policy adjustment, x_R , is defined by

$$\Delta x_R = \text{proj}_R(x_0 + \Delta x) - x_0,$$

where x is the unrestricted (“latent”) NGA policy adjustment. That is, the outcome of the restricted NGA mechanism is simply the projection of the latent NGA mechanism outcome to the nearest point in the restricted policy space R .

In the multiple referenda context, the NGA policy space is the hypercube defined by $x_i \in [0, 1]$ for all policy dimensions $i = 1, 2, \dots, P$.

5.2. Strategic Equivalence Between QV and NGA in the Context of Multiple Referenda

Under NGA, individuals report a normalized gradient $r_\theta = (r_{\theta,1}, r_{\theta,2}, \dots, r_{\theta,P})'$ to maximize their utility. The normalization means that individuals are choosing a point on the unit sphere. Before the restriction to the hypercube $[0, 1]^P$ is applied, the latent policy adjustment is the sum of the reported normalized gradients: $\Delta x = \frac{\varepsilon}{\theta} \sum_{\theta} r_\theta$. For each policy dimension i , the latent policy change is $\Delta x_i = \frac{\varepsilon}{\theta} \sum_{\theta} r_{\theta,i}$.

The unit-sphere budget set for the vote vector r_θ in NGA is identical to the budget set for the vote vector v_θ in QV, $v_{\theta,1}^2 + v_{\theta,2}^2 + \dots + v_{\theta,P}^2 \leq 1$, as long as the QV budget constraint is satisfied with equality. Since we assume that the utility functions are the same, the votes will be the same if the mapping from votes to outcomes is the same in the two mechanisms.

To make the mapping from votes to outcomes the same in NGA as in QV requires taking the limit of restricted NGA as the step size goes to infinity. As shown in Figure 3, if the latent policy change $x_i \notin [0, 1]$ for every i , then x will be outside of the shaded cross (depicted as the solid arrow), and the closest point onto which the mechanism will project is a vertex. As a result, the outcome is deterministic. On the other hand, if $x_i \in [0, 1]$ for some but not all i (as could result from an intermediate step size), then the latent policy adjustment vector x will stay within the shaded cross region. The restricted NGA mechanism will then project the endpoint of that vector onto the closest point within the restricted region, depicted as the middle square, resulting in a probabilistic outcome for at least one referendum. It is clear from the figure that as long as the sum of votes is nonzero in every dimension, there exists a sufficiently large step size ε to ensure that the latent policy-change

vector x will be outside of the shaded region. Thus, in the large- ϵ limit, restricted NGA gives exactly the same mapping from votes to outcomes as the simplified version of QV defined at the beginning of the present section, and the two mechanisms become strategically equivalent.

5.3. Lack of Strategy-Proofness

Because QV is an indirect mechanism, the concept of strategy-proofness does not directly apply to it. But the concept of strategy-proofness does apply to the strategically equivalent NGA mechanism. In this section's case of $P > 1$ discrete public goods, NGA is not strategy-proof. Suppose there are three individuals, denoted A, B, and C, and they are voting on two referenda, 1 and 2. Let the agents have preferences over the referenda as follows: agent A only wants referendum 1 to pass, agent B wants both referenda to fail and cares equally about the two, and agent C only wants referendum 2 to fail. Clearly the dominant strategy for A is to vote (1, 0), and for C is to vote (0, -1). The truthful strategy for B would be to vote $(\frac{-1}{\sqrt{2}}, \frac{-1}{\sqrt{2}})$, which would result in referendum 1 passing and referendum 2 failing. However, the summed votes without B are $V_{-B} = (1, -1)$. Note that even if B buys no votes for referendum 2, it fails. Therefore, B will use all of her influence to change the outcome of referendum 1. That is, B maximizes her utility by voting (-1, 0). The resulting vote sums are (0, -1) so that referendum 1 passes with probability 50% and referendum 2 fails.

This example points to a rather general complication that would afflict a wide range of mechanisms: with multiple referenda, individuals have an incentive to concentrate their votes on the referenda for which they are more likely to be pivotal, rather than wasting votes on the referenda that are foregone conclusions. Therefore, individuals will devote their voting power more to the close referenda than would be predicted by a simple ratio of individual valuations of the outcomes of the referenda. Such (possibly complex) strategic incentives are identical for the simplified version of QV considered in this section and NGA restricted to the probability hypercube in the large- ϵ limit.

5.4. Conjecture: Asymptotic Efficiency with a Large Number of Referenda

Suppose that the number of referenda goes to infinity, with the separable preferences over each referendum drawn identically and independently from a well-behaved distribution. Then the opportunity cost of tokens must become effectively linear as they are optimally distributed to equalize the marginal value of using tokens in different referenda. The linearity of the opportunity cost of tokens in the limit as the number of referenda goes to infinity should make the analysis essentially equivalent to that of QV with money. The formal proof of this would be nontrivial, but we put forward the conjecture that (under these i.i.d. regularity conditions) as the number of separable referenda goes to infinity, tokens become a quasilinear currency to which the logic of quasilinear money applies—including the logic pointing to asymptotic efficiency of QV with money in large populations. Thus, we conjecture that with many referenda and a large population, QV with tokens and (because of its equivalence to QV in this context) restricted NGA over referenda with large ϵ would both become asymptotically efficient. (There is an extension of this conjecture that allows correlations across referenda by allowing preferences over small groups of referenda to be

drawn independently and identically from a well-behaved distribution. The same basic argument applies to that case.)

The argument for efficiency with many referenda and a large population is that (a) with a large number of referenda, the situation for QV with tokens for each referendum becomes more and more similar to an instance of QV with money for that referendum; (b) the basic logic of the Lalley and Weyl result about efficiency with large populations will then carry over. Note that conceptually, the Kaldor-Hicks compensation criterion for efficiency can be applied with tokens instead of money. The money-like nature of tokens is clearest when the referenda are conducted sequentially, since then the compensation could conceptually be conducted in what is ex post relative to a particular referendum. But ex ante Kaldor-Hicks compensation based on expected values for a given referendum can also be considered. To avoid income effects from ex ante Kaldor-Hicks compensation piling up for certain individuals, it may only be possible to carry through the notion of ex ante Kaldor-Hicks compensation when, in addition to the assumptions mentioned already, the distribution of preferences over referenda is symmetric between different individuals so that the ex ante Kaldor-Hicks compensation for each individual would tend to zero as the number of referenda goes to infinity. This corresponds to the idea that the Kaldor-Hicks compensation criterion is appropriate when there are a large number of decisions for which the distributional effects on any given individual tend to cancel out.

To clarify the conjecture, with a large number of referenda and a large population, there should be “internal” efficiency in the sense of the Kaldor-Hicks compensation criterion with tokens in place of money. If, in addition, there are symmetric draws of values across people (so no one in aggregate values the public goods more than anyone else on average), then the logic of ex ante Kaldor-Hicks compensation should guarantee global Pareto efficiency.

6. Comparing QV to NGA Adapted to the QV Context

In this section, we first extend NGA so that it can be adapted to the QV setting introduced in section 2, and then we compare the two mechanisms in this setting. The bottom line is that the extension of NGA is approximately strategically equivalent to QV. We then show two limit results: (i) assuming truthful reporting, votes for or against the public good proportional to net valuations are obtained exactly in the limit case of the NGA mechanism that closely approximates the QV mechanism; and (ii) as the population gets large, a careful ex ante choice of the step size ε in NGA can simultaneously make the incentive to misreport one’s gradient vanishingly small and make the probability of an inefficient outcome vanishingly small.

6.1. The Setup of NGA in the QV Context

As we have assumed throughout, there are Θ individuals, indexed by $\theta = 1, \dots, \Theta$. Their valuations for the public good, u_θ , are independently and identically drawn from a known continuous probability distribution F with finite support $[\underline{u}, \bar{u}]$, $\underline{u} < 0 < \bar{u}$.

The “policy dimensions” in this context are expected spending on the public good, denoted z , and the amount of money held by each of the Θ individuals, denoted by the vector $\mathbf{w} =$

$(w_1, w_2, \dots, w_\theta)'$. Thus, the policy vector is $x = (z, w) \in \mathbb{R}^{1+\theta}$. Its initial value is $(c/2, w_0)$, where c is the cost of the public good, and starting at $z = c/2$ implies that there is a 50% probability of producing the public good when the net vote total is exactly 0. (This is meant to match the fact that under QV, if the net vote is 0, then there is a non-trivial probability of the public good being provided and a nontrivial probability of the public good not being provided.)

To make the notation and the form of the utility function comparable to QV, we assume that individuals' von Neumann-Morgenstern utility functions are quasilinear in money:

$$U_\theta(x) = u_\theta \Psi \left(\frac{2z - c}{c} \right) + w_\theta, \quad (2)$$

where $\Psi(\cdot)$ is the probability of the public good being produced (discussed further below), and $w_\theta = w_{\theta,0} + t_\theta$ is individual θ 's total wealth after the transfer t_θ (possibly negative) resulting from the mechanism itself and from the lump-sum tax (discussed below) to finance the public good if it is produced. The function Ψ is assumed to be a smooth cumulative distribution function with $\Psi(c) = 1$ and $\Psi(0) = 0$, and $\Psi(\cdot)$ strictly increasing on the interval $(0, c)$. In this adaptation of NGA to the QV context, we further assume that $\Psi(\cdot)$ is a linear function on that interval. Thus, the probability of producing the public good is $\Psi([2z - c]/c) = z/c$ for $z \in [0, c]$, with the probability equal to 0 when $z = 0$ and equal to 1 when $z = c$. For both QV and this adaptation of NGA, we think of the assumption of quasilinearity in money as an approximation to a more realistic utility function with some curvature in money when the amount of money involved in the public-good assessment for any one individual is small compared to each individual's wealth. In the notation here, that can be thought of as the assumption that a change in wealth by amount $(c/\theta) > 0$ —the amount that each individual would have to pay if the cost of the public good were shared equally—would cause only a negligible change in the marginal utility of wealth.

In order to have an adaptation of NGA that is closely comparable to QV, we need to ensure that the mechanism does not involve expected spending on the public good outside the range $[0, c]$ (to keep the probability of producing the public good within $[0, 1]$), and we need to ensure budget balance. Both of these are accomplished by an appropriate specification of the policy space with boundaries. The restricted policy space is $R = \{(z, w) \in \mathbb{R}^{1+\theta} : z \in [0, c], z + \sum_\theta w_\theta = 1\}$, where, without loss of generality, we have normalized total wealth in the population to 1. The next subsection shows how the budget constraint can be built into the specification of the policy dimensions over which the individuals vote.

6.2. Two-Person Graphical Example for NGA Adapted to the QV Context

In this subsection, we walk through graphical examples that illustrate how NGA works in the QV context. Our two-person examples will demonstrate how restrictions of probabilities to the $[0, 1]$ interval are incorporated into the policy space, and how we embed budget balance into the specification of the policy dimensions. We will also use this graphical apparatus and its θ -person extension to provide intuition for formal results on efficiency and strategy-proofness given in the subsequent subsections.

We start without the public good, and then add it. Because the total amount of wealth in the population is held fixed at 1, we can describe society's budget constraint by $\sum_{\theta} w_{\theta} = 1$, whose boundary is a hyperplane of dimension $\Theta - 1$. We define the policy space by $\Theta - 1$ "redistributive dimensions" parallel to this hyperplane. Within these $\Theta - 1$ redistributive dimensions are Θ directions that correspond to transferring wealth to one individual, financed by an equal lump-sum tax on everyone else. By construction, any direction of motion orthogonal to the redistributive direction that transfers money to individual θ does not affect the amount of money she has, and so she is indifferent to motion in such direction. (By assumption, redistribution between other individuals is not in individual θ 's utility function.)

In our two-person example, Figure 4 graphs person 1's wealth, w_1 , vs. person 2's wealth, w_2 , with the horizontal line indicating the budget-constraint hyperplane. The redistributive dimension is left-right. Moving left transfers money to person 1, and moving right transfers money to person 2. Each person seeks to maximize their wealth by moving in a direction parallel to their wealth vector: from any starting point in the mechanism, we can imagine each person pulling the vector in their direction of wealth. The NGA mechanism then adds the vectors from the two people to get a direction, which results in an equal distribution of the remainder of the pie if the step size is made large enough to reach (or cross) the boundary representing the limit on the total amount of wealth. (The step size is an important part of the mechanism that we will discuss in further detail in subsections 6.4 and 6.5, but in this subsection we will only depict the effects of varying step sizes ε graphically.) If the starting point is already on this boundary, these same vector contributions result in no movement from that initial point: equal division of zero remaining pie changes nothing.

We now introduce the public good. Expected spending on the public good constitutes an additional dimension of the policy space, which we call "the public good dimension." Budget balance is maintained because in the ultimate outcome, producing the public good is bundled with an arrangement for people collectively paying for it. Thus, graphically, "money for self" is actually "expected money for self." In Figure 4, the public good dimension is up-down; up corresponds to a higher probability of producing the public good and correspondingly higher expected assessments on everyone to pay for it. Since producing the public good costs money, doing so decreases the total expected wealth in the system available for redistribution from the thicker solid line to the dashed line. Given the right angle at the top, this distance is $\frac{\sqrt{2}}{2}c$ (recall that c is the cost of the public good). Since we can consider the benchmark case to be one where that vertical distance is c , what is depicted is actually that benchmark case "squished" by the factor $\sqrt{2}$ in the vertical direction. The thin solid line halfway between the two total wealth lines is the line at which there is a 50% chance of building the public good.¹

¹Using this setup, it is straightforward to calculate the gradients and indifference curves for each individual. Setting the right angle of the triangle as the origin, the horizontal movement of the mechanism is $x_1 = \frac{\sqrt{2}}{2}(w_2 - w_1)$ and the vertical movement is $x_2 = -\frac{\sqrt{2}}{2}(w_1 + w_2)$, where w_1 and w_2 are the wealth for person 1 and person 2, respectively. Then, we have the money-metric utility for person 1 as $U_1 = \frac{u_1}{c} - \frac{\sqrt{2}}{2}x_1 + \frac{\sqrt{2}}{2}\left(\frac{2u_1}{c} - 1\right)x_2$, where u_1 is her valuation of the public good, x_1 is the horizontal distance,

The gradients for person 1 and person 2 are shown in Figure 5. Each gradient is labeled by the person's valuation of the public good. We can see that as each person's valuation of the public good increases, her gradient points more upward, and vice versa. Before normalization, the gradients depend in a linear way on the valuation of the public good. The semicircles in Figure 5 mark the endpoint of each vector in the normalized gradients, in which each person's vector is of the same length.

We will now discuss four particular cases in our two-person scenario and the resulting outcome for the public good. In each case, we consider a range of possible step sizes, but to keep the figures readable, step sizes that would go beyond the boundary of the policy space are not shown.

Figure 6 depicts four scenarios in which person 1 and person 2 have different valuations of the public good. In each case, the sum of vectors is depicted with varying step sizes, ϵ , corresponding to different arrow heads. In case 1, the sum of each person's valuation of the public good is greater than the cost, so the public good will be built. This is depicted by gradients ∇U_1 and ∇U_2 and their sum which, with a large enough step size, hits the public-good line. In case 2, the sum of person 1's and person 2's vectors with a large enough step size hits the no-public-good line because, although they both value the public good at more than nothing, they do not collectively value it enough to pay for it. In case 3, person 2 values the public good at less than nothing, meaning that even if the good were free, she still would not want it. Nonetheless, person 1 wants it enough relative to how much person 2 dislikes it that the public good will still be built, as indicated by the sum of their vectors pointing up towards the public-good line. Case 4 shows a knife-edge scenario in which there is a stalemate. Because person 1's and person 2's gradients are equal and opposite—reflecting valuations that barely add up to the cost of the public good with nothing to spare—the mechanism has no direction to move in, regardless of the step size. Thus, the mechanism is stuck at the starting point, on the halfway line between building and not building the public good. As a result, the public good has a 50% chance of being built.

6.3. Voting Approximately in Proportion to Net Valuation in NGA

We now argue that the implicit votes for or against the public good in NGA are approximately in proportion to an individual's net valuation if a particular individual's vote is small and if the individual can ignore the limitation represented by the boundaries of the restricted region. (By net valuation we mean the valuation of the bundle of having the public good produced and each individual paying an equal share of the cost.) We begin by discussing the policy space when the height of the triangle shown in the previous subsection is 1 instead of $\frac{\sqrt{2}}{2}$, and continue on to consider stretching or squishing the triangle to an arbitrary height h .

and x_2 is the vertical distance. From this utility function, the gradient is $\nabla U_1 = \frac{\sqrt{2}}{2}(-1, \frac{2u_1}{c} - 1)$, and the indifference curve is

$$x_2 = \frac{\sqrt{2}U_1 - \frac{2u_1}{c} + x_1}{\frac{2u_1}{c} - 1}.$$

As per the discussion above, each individual θ only cares about the public-good dimension and the redistribution direction toward (or against) herself. Therefore, we can describe the decision problem in this way: each individual θ chooses votes for the public good, z_θ and transfer to self, s_θ subject to the voting constraint, $s_\theta^2 + z_\theta^2 \leq 1$ (with inequality rather than equality because the individual could instead vote along other dimensions). Nonsatiation in money implies that the constraint will hold with equality: $s_\theta^2 + z_\theta^2 = 1$.

Assuming that the mechanism remains inside the boundaries of the restricted region, individual θ 's indirect utility function H_θ can be written as:

$$H_\theta(s_\theta, z_\theta; s_{-\theta}, z_{-\theta}) = u_\theta \frac{z_\theta(z_\theta, z_{-\theta})}{c} + w_\theta(s_\theta, z_\theta; s_{-\theta}, z_{-\theta}) \\ = u_\theta z(0, z_{-\theta}) + w_\theta(0, 0; s_{-\theta}, z_{-\theta}) + \frac{\varepsilon}{\Theta} \left\{ \frac{z_\theta}{c} [u_\theta - (c/\Theta)] + s_\theta \right\}. \quad (3)$$

The first equality simply reflects the direct utility function, $U_\theta = u_\theta z/c + w_\theta$. The second equality decomposes individual θ 's utility into the effect of everyone else's choices of $z_{-\theta}$ and $s_{-\theta}$ plus the effect of individual θ 's choices of z_θ and s_θ . To understand the effect of individual θ 's choices, recall that individual θ 's impact on the overall policy-change vector is of length $\frac{\varepsilon}{\Theta}$, and thus allocating z_θ votes to the public good changes the probability of producing it by $\frac{\varepsilon}{\Theta} \frac{z_\theta}{c}$; the net benefit of the public good is $u_\theta - (c/\Theta)$ (individual θ 's valuation minus her share of the cost); and allocating s_θ votes to transfer to self increases her wealth by $\frac{\varepsilon}{\Theta} s_\theta$.

Now consider the difference in utility between individual θ 's optimal choices and the alternative of allocating all votes to transfer to self:

$H_\theta(s_\theta, z_\theta) - H_\theta(1, 0) = \frac{\varepsilon}{\Theta} \left\{ \frac{z_\theta}{c} [u_\theta - (c/\Theta)] + s_\theta - 1 \right\}$. Rearranging the budget constraint and taking a third-order Taylor approximation, $s_\theta = \sqrt{1 - z_\theta^2} = 1 - \frac{1}{2} z_\theta^2 + O(z_\theta^4)$. Finally, substituting into the utility difference yields:

$$H_\theta(s_\theta, z_\theta) - H_\theta(1, 0) \approx \frac{\varepsilon}{\Theta} \left\{ \frac{z_\theta}{c} \left(u_\theta - \frac{c}{\Theta} \right) - \frac{1}{2} z_\theta^2 \right\}. \quad (4)$$

This equation means that the benefit of voting for or against the public good is linear, while the cost is approximately quadratic. The quality of the approximation becomes better when the vote for or against the public good is small. Thus, the individual's strategic incentives are approximately the same as in QV for sufficiently small values of z_θ .

In the next subsection, we will transform the policy space by stretching the public-good dimension by the factor h . Doing so reduces the effect of a vote for or against the public good on the probability that the public good is produced, changing the equation above to

$$H_\theta(s_\theta, z_\theta) - H_\theta(1, 0) = \frac{\varepsilon}{\Theta} \left\{ \frac{z_\theta}{ch} \left(u_\theta - \frac{c}{\Theta} \right) - 1 + \sqrt{1 - z_\theta^2} \right\} \approx \frac{\varepsilon}{\Theta} \left\{ \frac{z_\theta}{ch} \left(u_\theta - \frac{c}{\Theta} \right) - \frac{1}{2} z_\theta^2 \right\}.$$

By reducing the benefit of voting for or against the public good, this transformation reduces the optimal size of votes on the public good, making the quadratic approximation for the cost of such votes more exact.

6.4. A Limit Result on Efficiency

In subsection 6.3, we showed that if an individual does not need to worry about the boundaries of the restricted region and has a small implicit vote, that vote will be approximately proportional to the individual's net valuation of the public good. In this subsection, we show that a limit within a particular family of extended NGA mechanisms can guarantee that everyone's votes for or against the public good become very small, so that in the limit the voting in proportion to net valuation result of subsection 6.3 goes through for everyone in a way that becomes exact in the limit. The reason is that the cost of voting on the public good becomes exactly quadratic in that limit. This proportionality result of votes to net valuations is part of what is needed for efficiency. (Subsection 6.5 discusses the other dimension of efficiency: whether the mechanism decisively gives a probability of 1 of producing the public good if the Samuelson criterion of a positive sum of net valuations for the public good is satisfied and a probability of 0 of producing if that criterion is not satisfied.)

The relevant family of NGA mechanisms is defined by the choice of units for the public good dimension. Although it is not immediately obvious, we consider the need to choose the units for each policy dimension—or more generally the coordinatization of the policy space—to be the analog in NGA to the “agenda-setting” issue that appears in any voting mechanism. (In QV, the agenda-setting issue arises in deciding which votes to conduct.) The reason we consider the choice of units to be an agenda-setting issue is that in Benjamin et al. (2016), we show that for any individual who wants to stay at the status quo, there is a sizable advantage to gaining control over the choice of units and the orientation of the axes for the policy space. Taking this view of the choice of units and orientation of the axes as a type of agenda setting, part of getting NGA to look as much as possible like QV is to set the agenda for NGA appropriately. Somewhat paradoxically, the way to get the adaptation of NGA to the QV context to look most like QV is to have the public good decision *not* be the main thing the NGA mechanism is focused on. The reason is that the less of the implicit vote budget (given by the limitation to a unit sphere) an individual devotes to voting for or against the public good, the more exact the approximation will be to a quadratic financial cost of voting for or against the public good.

Below we show graphically how “stretching” the policy space in the vertical (public-good) dimension makes gradients for any finite positive or negative valuation of the public good more nearly horizontal, implying that the components of the gradient in the public-good direction become smaller. Smaller votes in the public good dimension lead in the limit to an exact quadratic cost for voting on the public good and therefore to a close analog to QV. Conversely, “squishing” the policy space in the vertical (public-good) dimension makes gradients more nearly vertical, and in the limit yields majority voting over the public good. Thus, NGA adapted to the QV context provides both QV-like and majority voting limits, as

well as members of the family of mechanisms that fall in between the QV-like and majority voting limits.

Figure 7 illustrates graphically how the choice of units affects NGA adapted to the QV context. An individual's gradient is orthogonal to the indifference curves shown. Stretching the policy space vertically makes the indifference curves more nearly vertical, and thus makes the gradients orthogonal to those indifference curves more nearly horizontal—implying little voting in the public good dimensions. Conversely, squishing the policy space vertically makes the indifference curves more nearly horizontal and thus makes the gradients orthogonal to those indifference curves more nearly vertical, implying that NGA votes focus on the public good dimension.

To show what this stretching implies when there are θ individuals, Figure 8 shows a two-dimensional cross-section of the policy space with θ individuals and vertical stretching by the factor h . The base of what looks like a triangle is actually a regular simplex with $\theta - 1$ redistributive dimensions, but only the redistributive direction the individual cares about is drawn. That is, the x-axis is individual θ 's wealth. The y-axis, as before, represents the public-good dimension: votes for or against producing the public good and charging everyone an equal amount to pay for it. The apex of the triangle is directly above the point on the base where each individual has equal wealth $1/\theta$. Because of the stretching by factor h , the height of the triangle is equal to h , and the vertical distance between producing and not producing the public good is now ch .

Now, using individual θ 's indirect utility function

$$H_{\theta}(s_{\theta}, z_{\theta}; s_{-\theta}, z_{-\theta}) = u_{\theta}z(0, z_{-\theta}) + w_{\theta}(0, 0; s_{-\theta}, z_{-\theta}) + \frac{\varepsilon}{\theta} \left\{ \frac{z_{\theta}}{ch} [u_{\theta} - (c/\theta)] + s_{\theta} \right\},$$

we can calculate the gradient as $\nabla H_{\theta} = \frac{\varepsilon}{\theta} (1, (\frac{1}{ch}(u_{\theta} - \frac{c}{\theta}))$. The limit as h decreases to 0 is majority voting, because gradients will become very close to vertical, and everyone will expend all their votes in the public-good dimension. In the other direction, in the limit as h approaches infinity, the mechanism becomes very close to QV because gradients become so close to horizontal that the cost of the vertical component becomes exactly quadratic.

For efficiency, the key implication of having a cost of voting that becomes exactly quadratic in the limit is that the vertical components of the normalized gradient vectors become exactly proportional to individuals' valuations in the limit. We prove this formally in Theorem 1 and its corollary. Proofs are relegated to the Appendix. In the statement and proof of theorem, we use the fact that the vertical transformations we consider above are equivalent to keeping the original parameterization of the policy space but making the set from which an individual can choose a vector vote an ellipsoid rather than a sphere.

Theorem 1—Denote by S_n an ellipsoid whose axis in the z dimension is squished by a factor of $\sqrt{n}$: $S_n = \{(z, w) \in \mathbb{R}^{1+\theta} : nz^2 + w_1^2 + \dots + w_{\theta}^2 \leq 1\}$. Let $z_{\theta}(n)$ denote θ 's optimal

contribution to the public good when the choice set is S_n . Then, $z_{\theta}(n) \rightarrow 0$ as $n \rightarrow \infty$ and $z_{\theta}(n)/z_{\theta'}(n) \rightarrow u_{\theta}/u_{\theta'}$ as $n \rightarrow \infty$.

Corollary 1—*The quadratic approximation in (4) is asymptotically valid as $n \rightarrow \infty$.*

6.5. Approximate Strategy-proofness in Large Populations

A key result for QV is that it becomes efficient asymptotically as the number of individuals becomes large. There is also an asymptotic efficiency result for NGA adapted to the QV context. The efficiency result for NGA in large populations involves a result about the incentives for truth-telling. It turns out that when adapted to the QV context, NGA is approximately strategy-proof in large populations in the sense that the incentive to report anything other than one's true normalized gradient can be made vanishingly small, even faster than an individual's arrow gets small as that individual becomes a smaller fraction of the population. We assume that there is some very small direct utility cost $\kappa > 0$ of lying (i.e., misreporting) beyond what is shown in the calculations below, so that if the benefit of strategic misreporting shown in the calculations below is small enough, an agent will report truthfully. In assuming this, it is important to remember that we would implement NGA by surveys in which the framing of truthful reporting is quite salient. In the QV context, that boils down to a survey asking people how much they value the public good. We would then use the gradient implied by that answer. The issue in this subsection is trying to guarantee that, given the way those answers are used, individuals can answer that survey question truthfully without any substantial harm to their own interests.

What follows is a summary of an argument spelled out in greater detail in the appendix. The basic thread of the argument is as follows: (1) within the interior of the restricted policy space, truth-telling is optimal, and if the probability that an agent gets close to the boundary becomes small, the temptation to misreport becomes very small. We will describe this as “approximate strategy-proofness.” (2) We discussed in the previous subsection, that as the stretching parameter h goes to infinity, the NGA mechanism goes to a higher or lower probability of producing the public good according to the Samuelson criterion of whether the sum of net valuations of the public good is positive or negative (or equivalently, whether the sum of gross valuations is greater or less than the cost). But there is another dimension of efficiency: how close the mechanism gets to the “produce the public good with probability 1” boundary or the “produce the public good with probability 0” boundary. (3) As the number of agents goes to infinity, the law of averages makes the distribution of where everyone else's actions will put the mechanism tight. Indeed, the distribution of where everyone else's actions put the mechanism becomes tight enough that it is possible to choose a step-size ε that both (a) makes the probability an agent's effect on the mechanism will be influenced by the boundary very small and (b) makes the distance from the boundary likely to be quite small in comparison to the total distance from the “don't produce” to the “produce” boundary. The formal argument shows (i) that the temptation to misreport is small if the chance of interacting with the boundary is small, (ii) that the chance of interacting with the boundary can be made small, and (iii) that the chance of interacting with the boundary can be made small without actually being that far from the boundary. What follows is a more detailed description of what each formal result does and a bit of the reasoning behind each

piece. Some readers may wish to skip directly to “**Remarks**” below while others may prefer to read the appendix rather than the intervening paragraphs, but we suspect that some readers will find it useful to see an argument intermediate in formality between the current paragraph and what is in the appendix.

Define an agent “not interacting with the boundary” as having the collective choices of other agents resulting in a movement $\mathbf{z}_{-\theta} \equiv (\varepsilon/\Theta)\Sigma_{\hat{\theta}} \mathbf{z}_{\hat{\theta}}$ that leaves the starting point for θ 's arrow at least ε/Θ away from the edge, and an agent “interacting with the boundary” as having the collective choices of other agents put the mechanism (before θ 's arrow is added) closer to the boundary than ε/Θ . Then let ξ be an upper bound on the probability any agent perceives of interacting with the boundary. We can establish the following bound on the benefit to strategic misreporting:

Theorem 2—Let $\mathbf{x}_{\theta}(h)$ denote θ 's truthful response given stretching parameter h , let $\mathbf{x}_{\theta}^*(h, \Theta)$ denote θ 's optimal response given stretching parameter h , when there are Θ agents in the economy, and suppose the probability each agent perceives that the outcome of the mechanism from others' actions, $\mathbf{z}_{0+} - \mathbf{z}_{-\theta}$, will be at least ε/Θ from the boundary is at least $1 - \xi$. Then $u_{\theta}(\mathbf{x}_{-\theta}, \mathbf{x}_{\theta}(h, \Theta)) - u_{\theta}(\mathbf{x}_{-\theta}, \mathbf{x}_{\theta}(h)) \leq \xi \frac{\varepsilon}{\Theta} \|\nabla \mathbf{u}_{\theta}(\mathbf{x}_{-\theta}, \mathbf{x}_{\theta}(h))\|$ where, with some abuse of notation, the last factor is the constant value the gradient takes on in the interior of the restricted region.

That is, the size of the benefit from strategic misreporting as opposed to reporting the true normalized gradient is bounded above by $\xi(\varepsilon/\Theta)\|\nabla \mathbf{u}_{\theta}(\mathbf{x}_{-\theta}, \mathbf{x}_{\theta}(h))\|$. To parse this upper bound on the benefit of strategic misreporting, first note that since each agent's preferences are linear, the (constant) magnitude of the true gradient of an agent's indirect utility function H in the interior of the restricted region indicates the most that an agent can gain by moving the final output of the mechanism by one unit. In fact, a single agent can move the latent outcome of the mechanism by at most ε/Θ , and lemma 2 (reproduced in the Appendix from Benjamin et al. (2016)) guarantees that the outcome of the restricted NGA mechanism is moved by an individual's vote no more than the amount that the latent outcome is moved. Lemma 2 also guarantees that voting reporting one's true normalized gradient leads to a final outcome that the individual prefers to the final outcome from reporting 0 (which is equivalent to not voting at all). So the benefit from strategic misreporting instead of reporting the truth must be less than the benefit from misreporting instead of not voting at all. Finally, truth-telling is optimal if the agent is not interacting with the boundary, which occurs with at least a $1 - \xi$ probability. So any positive benefit of strategic misreporting instead of reporting the true normalized gradient is limited to the ξ of the time at most that the agent is interacting with the boundary. As long as we can choose a step size $\varepsilon(h, \Theta)$ as a function of the stretching parameter and the number of agents that guarantees that the probability ξ of interacting with the boundary goes to 0 as the population size increases, the benefit of strategic reporting will go to 0 faster than ε/Θ , which is the weight of each agent in decision-making. (Note that $\varepsilon(h, \Theta)$ must be an apriori ex ante function of the stretching parameter and number of agents, determined in advance of seeing anyone's strategic decisions, in order to avoid interfering with strategy-proofness.)

The next key step is, given a common-knowledge distribution of preferences, to choose the step size function $\varepsilon(h, \Theta)$ so that anything within $1/\sqrt{\xi(\Theta)}$ standard deviations of the mean of $z_0 + z_{-\Theta}$ is at least ε/Θ away from the boundary. Then by Chebyshev's inequality, the probability any agent will interact with the boundary is less than ξ . In particular, one can choose $\varepsilon(h, \Theta)$ as follows.

Lemma 1—Fix h, Θ , θ and $k = \xi(\Theta)^{-1/2} > 0$. Suppose agents (except perhaps θ) are truthfully reporting. Define $\varepsilon(h, \Theta)$ by $|m(h)| + kd(h) = \frac{ch}{2} - \frac{\varepsilon}{\Theta}$ where $m(h) = \mathbb{E}(z_{-\theta})$ and $d(h)^2 = \mathbb{V}(z_{-\theta})$. Then

$$\mathbb{P}\left(|\Delta z_{-\theta}| > \frac{ch}{2} - \frac{\varepsilon}{\Theta}\right) < \xi.$$

A key way in which NGA adapted to the QV context guarantees approximate strategy-proofness is that the choice of $\varepsilon(h, \Theta)$ in the lemma above typically makes it possible for someone who is most likely on the losing side because of everyone else's preferences to slightly reduce the probability of producing the public good, and makes the actions of someone who is most likely on the winning side able to push the probability of her desired outcome closer to 1. The cost of this is that it will generally be less than certain that the efficient outcome will take place. But as the number of agents becomes large, the expected distance from the efficient-outcome boundary needed to give incentives not too much in conflict with reporting one's true normalized gradient gets smaller and smaller. This expected distance from the boundary representing the efficient outcome can be calculated as follows for the case when it is (ex ante) socially optimal to produce the public good:

Theorem 3—Suppose $\lim_{h \rightarrow \infty} \mathbb{E}z_{\theta}(h) > 0$. Then the ex ante probability of producing the public good is bounded below by

$$\frac{1}{2} + \frac{z_{\theta}}{2\beta(\Theta)} - \xi(\Theta) \left(\frac{1}{\beta(\Theta)} - \frac{2}{ch\xi(\Theta)^{1/2}} \right) \mathbb{1}_{\left\{\xi(\Theta)^{-1/2} < \varepsilon(\Theta, h)\right\}},$$

where $\beta(\Theta) = \frac{1}{\Theta} + \frac{\Theta-1}{\Theta}|z_{\theta}| + \left(\frac{\Theta-1}{\xi(\Theta)\Theta^2}(z_{\theta})\right)^{1/2}$. Furthermore, if h is sufficiently large so that $\mathbb{E}z_{\theta}(h) > 0$ and $\xi(\Theta)$ satisfies $\xi(\Theta) \rightarrow 0$ and $\Theta\xi(\Theta) \rightarrow \infty$ as $\Theta \rightarrow \infty$, then this lower bound converges to 1 as $\Theta \rightarrow \infty$.

Remarks—The limit as $h \rightarrow \infty$ of the sign of $\mathbb{E}z_{\theta}(h)$ can be seen as giving the ex ante Samuelson criterion for large populations. An argument very similar to the above shows that if this limit is negative, the probability of producing the public good will go to 0 in large populations. Here we do not try to judge the efficiency of either possible decision when $\mathbb{E}z_{\theta}(h) \rightarrow 0$ as $h \rightarrow \infty$.

There is something odd about using the NGA mechanism to decide on whether to produce the public good when, in large populations, the distribution of preferences everyone is assumed to know would clearly indicate whether or not to produce the public good. But actually running the mechanism can be seen as a way of publicly verifying what everyone knows. More substantively, the NGA mechanism can get an efficiency result in the limit even if everyone is wrong about the distribution as long as the true distribution is more decisive in the sense of having a larger absolute value of $\mathbb{E}z_\theta(h)$. Everyone believing in the distribution used to choose ε or believing in one with a smaller $|\mathbb{E}z_\theta(h)|$ is needed for approximate strategy-proofness, but the efficiency result can be carried through if the true $|\mathbb{E}z_\theta(h)|$ is greater than what everyone believes. Unfortunately, if the true $|\mathbb{E}z_\theta(h)|$ is smaller than what everyone believes, there will be inefficiency: too much has been sacrificed for the sake of approximate strategy-proofness since people's beliefs of a higher-than true $|\mathbb{E}z_\theta(h)|$ require a small ε to get approximate strategy-proofness. Nevertheless, the exercise of attempting public verification would at least reveal problems with everyone's beliefs. Weyl (2016) suggests that QV may be more robust to departures from the assumption of common knowledge about the distribution of preferences in the population than NGA adapted to the QV context is, but a similar robustness exercise for NGA adapted to the QV context would be needed to fully make the comparison.

Note that either efficiency or full strategy-proofness with no approximation are easy to guarantee in isolation but require some delicacy to obtain in combination. A small enough step size to guarantee that the mechanism will stay within the restricted region regardless of what anyone chooses guarantees strategy-proofness, but can cause the decision about producing the public good to become probabilistic with substantial probabilities of either outcome—and thereby inefficient. On the other hand, if truth-telling is guaranteed despite strong incentives for strategic misreporting, then any step size that reaches or exceeds the boundaries for producing or not producing the public good will yield efficiency in the limit as the stretching parameter $h \rightarrow \infty$. To guarantee the combination of efficiency and even approximate strategy-proofness requires a limiting result with a large population size and all the apparatus of our theorems.

Such a careful choice of ε based on a priori information about a distribution is not an implementation of NGA that we would actually recommend. We focus on it here because it is an adaptation of NGA that is reasonably close to QV as described by Lalley and Weyl (2016).

It might be possible to obtain a result even closer to Lalley and Weyl (2016) with a zero direct utility cost to misreporting and ε chosen carefully to give a particular equilibrium with strategic reporting that is different from the truth, but not by much. One substantial difficulty with such a calculation is that standard second-order conditions for an optimum can fail for agents on the losing side. Intuitively, if someone on the losing side partially gives up because the outcome seems like a foregone conclusion, the outcome is even more likely to be a foregone conclusion, which reinforces the temptation to give up. Thus there is the possibility of multiple optima for the actions of an agent on the losing side.

6.6. Asymptotic Expenditure on Voting in NGA Adapted to the QV Context

As the number of individuals goes to infinity, the choice of ε discussed above takes the mechanism closer and closer to the boundary of the restricted region. This allows a simple, informal calculation of what each individual will spend on voting in NGA asymptotically as both the number of individuals and h go to infinity. Asymptotically, as $h \rightarrow \infty$, $z_\theta = (1/ch)[u_\theta - (c/\theta)]$. The distance to the boundary from the starting place is equal to $ch/2$. One fact simplifying the asymptotic calculation is that as $\theta \rightarrow \infty$, the ratio between the sum over everyone and the sum over everyone but one person goes to 1. Then, as both $h \rightarrow \infty$ and $\theta \rightarrow \infty$, the step size ε goes asymptotically to $|(ch/2)/\{\sum_\theta (1/ch)[u_\theta - (c/\theta)]\}| = |(c^2 h^2/2)/\{\sum_\theta [u_\theta - (c/\theta)]\}|$ (assuming that the limit for the denominator is nonzero). Finally, asymptotically, the amount that an individual spends on voting goes to

$\varepsilon z_\theta^2/2 = [u_\theta - (c/\theta)]^2 / \{4|\sum u_{\theta'} - c|\}$. That is, asymptotically, the amount an individual spends on voting is one-quarter of that individual's valuation (net of the equal assessment) squared, divided by the absolute value of the social difference between benefit and cost. Equivalently, an individual's vote equals that individual's net valuation times one-quarter of the ratio between that individual's net valuation and the absolute value of the overall social difference between benefit and cost. This means that expenditure on voting dies away asymptotically as $1/\theta$ in a way inversely proportional to the number of people.

7. Concluding Remarks

Our analytic exercise in this paper sheds some new light on the view that QV and NGA are closely related mechanisms. Of course, much more can be done. Hylland et al. (2016) takes a further step by embedding both QV and NGA in a broader class of mechanisms in which the effective budget constraints are quadratic in some dimensions and linear in others. It would also be natural to investigate the connection of both mechanisms to the path-breaking work of Groves and Ledyard (1977), in which agents make a monetary payment proportional to the square of the difference between the amount they vote to add to a public good and others' votes.

Besides shedding light on the relationship between the two mechanisms, this paper has also provided some of the first formal analyses of QV as implemented with tokens rather than money, which is likely to be an attractive implementation of QV in some applications.

When applied to more than one collective decision, an important challenge for virtually all voting mechanisms is sensitivity of the outcome to agenda-setting. We are not aware of any research yet on the agenda-setting problem for QV. For NGA, something akin to the agenda-setting problem comes up in the choice of units for the different policy dimensions (see subsection 6.2 for an example of how changing the units will change the outcome of the mechanism). As Benjamin et al. (2013) and Benjamin et al. (2016) discuss, there are several ways in which the choice of units could be disciplined so as to remove some degrees of freedom. For example, a survey (of either experts or the general public) could be conducted to find out for what choice of units a one-unit change in each policy dimension would be judged intuitively as being of similar size. These intuitive judgments of distance in policy

space can also be cross-checked against the absolute size of the costs and benefits of a particular policy change using a monetary equivalent in a traditional cost-benefit approach.

While this paper yields a result that the combination of efficiency and strategy-proofness can be obtained in the limit as the population size becomes large, we believe that strategy-proofness and strategic manipulation more broadly may be less of a concern for some implementations of NGA than it is for other mechanisms. In particular, Benjamin et al. (2016) argue that eliciting implicit preferences through a combination of subjective well-being surveys and econometric analysis may make it harder to game that application of the NGA mechanism than NGA (or other mechanisms) implemented through direct voting would be. Of course, when it can be achieved at reasonable cost, strategy-proofness is obviously a useful property. In adapting NGA to the QV context, our notion of approximate incentive compatibility is having only a vanishingly small incentive to misreport, not a total absence of an incentive to misreport.

While no collective choice mechanism can be perfect, it may still be possible to dramatically improve upon collective decision-making as currently conducted in the real world. Indeed, we believe that the two closely related mechanisms of QV and NGA are promising as possible improvements in specific contexts. We are heartened to see the research interest that QV has justifiably garnered (as evidenced, for example, by this special issue). In our view, an important next step is more empirical tests of these mechanisms, like the experiments reported by Cárdenas et al. (2014), Goeree and Zhang (2016), and Quarfoot et al. (2016).

Acknowledgments

The authors would like to thank Glen Weyl for multiple rounds of detailed editorial feedback that significantly improved the paper; Itai Sher, Gabriel Carroll, and participants in the Becker-Friedman Institute's "Quadratic Voting and the Public Good" conference for helpful comments and suggestions; and Assaf Kott and especially Rebecca Royer for excellent research assistance.

References

- Arrow KJ. A difficulty in the concept of social welfare. *The Journal of Political Economy*. 1950; 58(4): 328–346.
- Benjamin, DJ., Carroll, G., Heffetz, O., Kimball, MS., Lougee, D. Aggregating Local Preferences to Guide Policy. Cornell University and University of Southern California mimeo; 2016.
- Benjamin DJ, Heffetz O, Kimball MS, Szembrot N. Aggregating Local Preferences to Guide Marginal Policy Adjustments. *American Economic Review*. 2013; 103(3):605–10. [PubMed: 24014888]
- Benjamin DJ, Heffetz O, Kimball MS, Szembrot N. Beyond Happiness and Satisfaction: Toward Well-Being Indices Based on Stated Preference. *American Economic Review*. 2014; 104(9):2698–2735. [PubMed: 25404760]
- Cárdenas, JC., Mantilla, C., Zárate, RD. Purchasing votes without cash: Implementing quadratic voting outside the lab. 2014. <https://www.aeaweb.org/aea/2015conference/program/retrieve.php?pdfid=719>
- Chung, H., Duggan, J. Directional Equilibria. University of Rochester mimeo; 2015.
- Finkelstein A, Taubman S, Wright B, Bernstein M, Gruber J, Newhouse JP, Allen H, Baicker K, Group OHS. The Oregon health insurance experiment: Evidence from the first year. *The Quarterly Journal of Economics*. 2012; 127(3):1057–1106. [PubMed: 23293397]
- Gibbard A. Manipulation of voting schemes: A general result. *Econometrica*. 1973; 41(4):587–601.
- Goeree, JK., Zhang, J. One man, one bid. Quadratic Voting and the Public Good Conference; 2016.

- Groves T, Ledyard J. Optimal allocation of public goods: A solution to the "free rider" problem. *Econometrica*. 1977; 45(4):783–808.
- Helliwell, J., Layard, R., Sachs, J., editors. *World Happiness Report*. The Earth Institute, Columbia University; 2015.
- Hylland A, Zeckhauser R. A Mechanism for Selecting Public Goods when Preferences must be Elicited. *Kennedy School of Government Discussion Paper D*. 1980
- Hylland, A., Zeckhauser, R., Weyl, EG. An efficient competitive economy with public goods. 2016. Contact Glen Weyl (glenweyl@microsoft.com)
- Kaplow, L., Kominers, SD. Who will vote quadratically? voter turnout and votes case under quadratic voting. *Quadratic Voting and the Public Good Conference*; 2016.
- Lalley, SP., Weyl, EG. Quadratic voting. SSRN; 2016. <http://ssrn.com/abstract=2003531>
- Laurence, B., Sher, I. Ethical considerations on quadratic voting. *Quadratic Voting and the Public Good Conference*; 2016.
- Ludwig J, Duncan GJ, Genetian LA, Katz LF, Kessler RC, Ling JR, Sanbonmatsu L. Neighborhood effects on the long-term well-being of low-income adults. *Science*. 2012; 337(6101):1505–1510. [PubMed: 22997331]
- OECD. Measuring well-being and progress: Well-being research. 2016.
- Quarfoot, D., von Kohorn, D., Slavin, K., Sutherland, R., Konar, E. Quadratic voting in the wild: Real people, real votes. *Quadratic Voting and the Public Good Conference*; 2016.
- Satterthwaite MA. 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.
- Weyl, EG. The efficiency of quadratic voting in finite populations. 2015. Available at SSRN: <http://ssrn.com/abstract=2571026>
- Weyl, EG. The robustness of quadratic voting. 2016. Available at SSRN: <http://ssrn.com/abstract=2571012>

A. Appendix: Proofs

A.1. Proofs from Subsection 6.4 (Efficiency)

Theorem 1

Denote by S_n an ellipsoid whose axis in the z dimension is squished by a factor of $\sqrt{n}$: $S_n = \{(z, \mathbf{w}) \in \mathbb{R}^{1+\theta} : nz^2 + w_1^2 + \dots + w_\theta^2 \leq 1\}$. Let $z_\theta(n)$ denote θ 's optimal contribution to the public good when the choice set is S_n . Then, $z_\theta(n) \rightarrow 0$ as $n \rightarrow \infty$ and $z_\theta(n)/z_{\theta'}(n) \rightarrow u_\theta/u_{\theta'}$ as $n \rightarrow \infty$.

Proof—From the perspective of type θ , the relevant constraint is in fact $nz^2 + w_\theta^2 = 1$. From (2), type θ faces the following optimization problem:

$$\max_{z_\theta, s_\theta} u_\theta \Psi \left(\frac{2}{c} \frac{\varepsilon}{\Theta} (z_{-\theta} + z_\theta) \right) + \frac{\varepsilon}{\Theta} s_\theta$$

subject to $nz_\theta^2 + s_\theta^2 = 1$, where $z_\theta = \sum_{\hat{\theta} \neq \theta} z_{\hat{\theta}}$. Using $s_\theta = \sqrt{1 - nz_\theta^2}$, the relevant first-order condition is

$$u_\theta \frac{2}{c} \psi \left(\frac{\varepsilon}{\Theta} (z_{-\theta} + z_\theta) \right) - \frac{nz_\theta}{\sqrt{1-nz_\theta^2}} = 0, \quad (5)$$

where $\psi \equiv \Psi'$.

Recall that $\varepsilon > 0$ is assumed to be sufficiently small so that $\psi(2z/c) \approx (0)$, where

$\Delta z = \frac{\varepsilon}{\Theta} \sum_\theta z_\theta$. That is, we can treat $\bar{\psi} \equiv \frac{2}{c} \psi(0)$ as a constant. This allows us to solve (5) as a quadratic equation:

$$z_\theta(n) = u_\theta \frac{\bar{\psi}}{\sqrt{n^2 + nu_\theta^2 \bar{\psi}^2}}. \quad (6)$$

As n becomes large, the n^2 term in the denominator of (6) dominates. Thus for n sufficiently large, $z_\theta(n) \rightarrow 0$ is approximately proportional to u_θ . To be precise,

$$\frac{z_\theta}{z_\theta} = \frac{u_\theta \bar{\psi} / \sqrt{n^2 + nu_\theta^2 \bar{\psi}^2}}{u_\theta \bar{\psi} / \sqrt{n^2 + nu_\theta^2 \bar{\psi}^2}} \rightarrow \frac{u_\theta}{u_\theta}$$

as $n \rightarrow \infty$.

Corollary 1

The quadratic approximation in (3) is asymptotically valid as $n \rightarrow \infty$.

Proof—Recall that the exact relationship in (3) is given by the Taylor expansion

$$1 - \sqrt{1 - z_\theta^2} = \frac{1}{2} z_\theta^2 + R_4(z_\theta),$$

where R_4 is the remainder term. In Lagrange form, there exists $\bar{z} \in (-|z_\theta|, |z_\theta|)$ such that

$$R_4(z_\theta) = \frac{\bar{z}^6 (4\bar{z}^2 + 3)}{8(1 - \bar{z}^2)^{9/2}} < \frac{z_\theta^6 (4z_\theta^2 + 3)}{8(1 - z_\theta^2)^{9/2}}.$$

Note that the upper bound of $R_4(z_\theta)$ is increasing in z_θ with $R_4(0) = 0$. Thus since $z_\theta \rightarrow 0$ as $n \rightarrow \infty$ by Theorem 1, $R_4(z_\theta) \rightarrow 0$ as $n \rightarrow \infty$.

A.2. Preliminary Results

Lemma 2 (From Benjamin et al., 2016.)

Fix $\mathbf{x}_{-\theta} = \Sigma \hat{\theta} - \theta \mathbf{x}_{\hat{\theta}}$ and let $\mathbf{x}_{\theta}, \tilde{\mathbf{x}}_{\theta} \in S$. Define $\Delta \mathbf{x} = \frac{\varepsilon}{\theta} (\mathbf{x}_{-\theta} + \mathbf{x}_{\theta})$ and $\Delta \tilde{\mathbf{x}} = \frac{\varepsilon}{\theta} (\mathbf{x}_{-\theta} + \tilde{\mathbf{x}}_{\theta})$, and define $\mathbf{x}_R$ and $\tilde{\mathbf{x}}_R$ to be the associated restricted NGA outcomes. Then $(\mathbf{x} - \tilde{\mathbf{x}}) \cdot (\mathbf{x}_R - \tilde{\mathbf{x}}_R) \geq 0$ and $\|\mathbf{x} - \tilde{\mathbf{x}}\| \geq \|\mathbf{x}_R - \tilde{\mathbf{x}}_R\|$.

Proof—By assumption, the restricted policy space R is convex. Furthermore, by construction, $\mathbf{x} - \mathbf{x}_R$ is normal to R at $\mathbf{x}_R$. Thus, $\mathbf{x} - \mathbf{x}_R$ defines a supporting hyperplane with R contained in the lower half-space. Similarly, $\tilde{\mathbf{x}} - \tilde{\mathbf{x}}_R$ defines a supporting hyperplane with R contained in the lower half-space. Since $\tilde{\mathbf{x}}_R \in R$, it is contained in the lower half-space defined by $\mathbf{x} - \mathbf{x}_R$. Thus, $(\mathbf{x} - \mathbf{x}_R) \cdot (\tilde{\mathbf{x}}_R - \mathbf{x}_R) \leq 0$. Similarly, $(\tilde{\mathbf{x}} - \tilde{\mathbf{x}}_R) \cdot (\mathbf{x}_R - \tilde{\mathbf{x}}_R) \leq 0$.

We can rewrite these equations as

$$\begin{aligned} (\Delta \mathbf{x}_R - \Delta \mathbf{x}) \cdot (\Delta \tilde{\mathbf{x}}_R - \Delta \mathbf{x}_R) &\geq 0, \\ (\Delta \tilde{\mathbf{x}} - \Delta \tilde{\mathbf{x}}_R) \cdot (\Delta \tilde{\mathbf{x}}_R - \Delta \mathbf{x}_R) &\geq 0. \end{aligned}$$

Furthermore,

$$(\Delta \tilde{\mathbf{x}}_R - \Delta \mathbf{x}_R) \cdot (\Delta \tilde{\mathbf{x}}_R - \Delta \mathbf{x}_R) \geq 0.$$

Summing,

$$(\Delta \tilde{\mathbf{x}} - \Delta \mathbf{x}) \cdot (\Delta \tilde{\mathbf{x}}_R - \Delta \mathbf{x}_R) \geq 0.$$

Summing the first two equations above,

$$\underbrace{(\Delta \mathbf{x}_R - \Delta \mathbf{x} + \Delta \tilde{\mathbf{x}} - \Delta \tilde{\mathbf{x}}_R)}_{\equiv \mathbf{A}} \cdot \underbrace{(\Delta \tilde{\mathbf{x}}_R - \Delta \mathbf{x}_R)}_{\equiv \mathbf{B}} \geq 0.$$

Note that $\mathbf{A} + \mathbf{B} = \tilde{\mathbf{x}} - \mathbf{x}$. Therefore,

$$\begin{aligned} \|\Delta \tilde{\mathbf{x}} - \Delta \mathbf{x}\|^2 &= (\mathbf{A} + \mathbf{B}) \cdot (\mathbf{A} + \mathbf{B}) = \|\mathbf{A}\|^2 + \|\mathbf{B}\|^2 + 2\mathbf{A} \cdot \mathbf{B} \\ &\geq \|\mathbf{B}\|^2 = \|\Delta \tilde{\mathbf{x}}_R - \Delta \mathbf{x}_R\|^2. \end{aligned}$$

Lemma 3

Let X be a bounded random variable, and let $\bar{x}$ be an upper bound of the support of X . Then for any $k > 0$,

$$X - \min\{X, k\} \leq \frac{\bar{x} - k}{k^2} \mathbf{1}\{k < \bar{x}\}.$$

Proof—We can write $\mathbb{E}X = \mathbb{E}(X; X \leq k) + \mathbb{E}(X; X > k)$ and

$$\min\{X, k\} = (X; X \leq k) + (k; X > k) = (X; X \leq k) + k\mathbb{P}(X > k).$$

Subtracting, the two equations,

$$X - \min\{X, k\} = (X; X > k) - k\mathbb{P}(X > k).$$

Since X is bounded above by $\bar{x}$,

$$(X; X > k) = \int_k^{\bar{x}} x d\mathbb{P}(x) \leq \int_k^{\bar{x}} \bar{x} d\mathbb{P}(x) = \bar{x}\mathbb{P}(X > k).$$

Thus,

$$X - \min\{X, k\} \leq \bar{x}\mathbb{P}(X > k) - k\mathbb{P}(X > k) = (\bar{x} - k)\mathbb{P}(X > k) \leq \frac{\bar{x} - k}{k^2} \mathbf{1}\{k < \bar{x}\}$$

where the final inequality follows from Chebyshev's inequality for $k < \bar{x}$, and clearly $\mathbb{P}(X > k) = 0$ if $k > \bar{x}$.

A.3. Proofs for Subsection 6.5 (Strategy-proofness)

Theorem 2

Let $\mathbf{x}_\theta(h)$ denote θ 's truthful response given stretching parameter h , let $\mathbf{x}_\theta^*(h, \Theta)$ denote θ 's optimal response given stretching parameter h , when there are Θ agents in the economy, and suppose the probability each agent perceives that the outcome of the mechanism from others actions, $\mathbf{z}_0 + \mathbf{z}_{-\theta}$ will be at least ε/Θ from the boundary is at least $1 - \xi$. Then

$$u_\theta(\mathbf{x}_{-\theta}, \mathbf{x}_\theta(h, \Theta)) - u_\theta(\mathbf{x}_{-\theta}, \mathbf{x}_\theta^*(h)) \leq \xi \frac{\varepsilon}{\Theta} \|\nabla \mathbf{u}_\theta(\mathbf{x}_{-\theta}, \mathbf{x}_\theta(h))\|.$$

Proof—Note that $u(x_{-\theta}, x_\theta^*)$ is bounded above by the utility θ can obtain by reporting ex post. So suppose that agent θ is able to report their contribution after observing $\mathbf{x}_{-\theta}$. Conditional on not being able to reach the boundary of the constrained space, Benjamin et al. (2016) show that truth-telling is optimal. So consider the case that $\mathbf{x}_{-\theta}$ is such that θ can reach the boundary of the restricted space.

Lemma 2 with $\mathbf{x}_\theta = \mathbf{0}$ shows that $\nabla \mathbf{u}_\theta \cdot (\mathbf{x}_R - \mathbf{x}_\theta) = 0$ (since $\mathbf{x}_\theta$ is equal to the normalized gradient). Therefore, $u_\theta(\mathbf{x}_{-\theta}, \mathbf{0}) = u_\theta(\mathbf{x}_{-\theta}, \mathbf{x}_\theta)$; that is, reporting nothing is dominated by truth-telling.

Note that since utility is locally linear, for any $\hat{\mathbf{x}}_\theta \in S$,

$$\begin{aligned} u(\mathbf{x}_{-\theta}, \hat{\mathbf{x}}_\theta) - u(\mathbf{x}_{-\theta}, \mathbf{0}) &= \frac{\varepsilon}{\Theta} \|\hat{\mathbf{x}}_\theta\| \nabla_{\hat{\mathbf{x}}_\theta} \mathbf{u}_\theta \leq \frac{\varepsilon}{\Theta} \nabla_{\nabla \mathbf{u}_\theta} \mathbf{u}_\theta \\ &= \frac{\varepsilon}{\Theta} \nabla \mathbf{u}_\theta \cdot \frac{\nabla \mathbf{u}_\theta}{\|\nabla \mathbf{u}_\theta\|} = \frac{\varepsilon}{\Theta} \|\nabla \mathbf{u}_\theta\|. \end{aligned}$$

The inequality follows from the fact that the direction derivative $\nabla_{\mathbf{x}} \mathbf{u}_\theta$ is maximized at $\mathbf{x} = \nabla \mathbf{u}_\theta$ and since $\hat{\mathbf{x}}_\theta \in S$ implies $\|\hat{\mathbf{x}}_\theta\| \leq 1$. In particular, this holds for the optimal report $\mathbf{x}_\theta^*$.

Thus,

$$u_\theta(\mathbf{x}_{-\theta}, \mathbf{x}_\theta^*) - u_\theta(\mathbf{x}_{-\theta}, \mathbf{x}_\theta) = (u_\theta(\mathbf{x}_{-\theta}, \mathbf{x}_\theta^*) - u_\theta(\mathbf{x}_{-\theta}, \mathbf{0})) - (u_\theta(\mathbf{x}_{-\theta}, \mathbf{x}_\theta) - u_\theta(\mathbf{x}_{-\theta}, \mathbf{0})) \leq \frac{\varepsilon}{\Theta} \|\nabla \mathbf{u}_\theta\|.$$

Since truth-telling is the optimal when the agent cannot interact with the boundary, this provides a count on utility when reports are made ex post. The result follows from the fact that the probability of interacting with the boundary is ξ , and that reporting ex post provides an upper bound on attainable utility.

Lemma 1

Fix h, Θ , θ and $k = \xi(\Theta)^{-1/2} > 0$. Suppose agents (except perhaps θ) are truthfully reporting. Define $\varepsilon(h, \Theta)$ by $|m(h)| + kd(h) = \frac{ch}{2} - \frac{\varepsilon}{\Theta}$ where $m(h) = \mathbb{E}(z_{-\theta})$ and $d(h)^2 = \mathbb{V}(z_{-\theta})$. Then

$$\mathbb{P}\left(|\Delta z_{-\theta}| > \frac{ch}{2} - \frac{\varepsilon}{\Theta}\right) < \xi.$$

Proof

$$\begin{aligned} \mathbb{P}\left(|z_{-\theta}| > \frac{ch}{2} - \frac{\varepsilon}{\Theta}\right) &= \mathbb{P}(|z_{-\theta}| > |m(h)| + kd(h)) \\ &= \mathbb{P}(|z_{-\theta} - m(h)| > kd(h)) \end{aligned}$$

From the reverse triangle inequality, $|z_{-\theta} - m(h)| \geq |z_{-\theta}| - |m(h)|$. Therefore,

$$\mathbb{P}(|z_{-\theta}| - |m(h)| > kd(h)) \leq \mathbb{P}(|z_{-\theta} - m(h)| > kd(h)) < \frac{1}{k^2}$$

where the last inequality follows from Chebyshev's inequality.

Theorem 3

Suppose $\lim_{h \rightarrow \infty} \mathbb{E} z_{\theta}(h) > 0$. Then the ex ante probability of producing the public good is bounded below by

$$\frac{1}{2} + \frac{z_{\theta}}{2\beta(\theta)} - \xi(\theta) \left(\frac{1}{\beta(\theta)} - \frac{2}{ch\xi(\theta)^{1/2}} \right) 1_{\{\xi(\theta)^{-1/2} < \varepsilon(\theta, h)\}},$$

where $\beta(\theta) = \frac{1}{\theta} + \frac{\theta-1}{\theta} |z_{\theta}| + \left(\frac{\theta-1}{\xi(\theta)\theta^2} (z_{\theta}) \right)^{1/2}$. Furthermore, if h is sufficiently large so that $\mathbb{E} z_{\theta}(h) > 0$ and $\xi(\theta)$ satisfies $\xi(\theta) \rightarrow 0$ and $\theta\xi(\theta) \rightarrow \infty$ as $\theta \rightarrow \infty$, then this lower bound converges to 1 as $\theta \rightarrow \infty$.

Proof—Denote by A the event that the public good is produced. Then,

$$\mathbb{P}(A | \Delta z_R) = \frac{1}{2} + \frac{1}{ch} \Delta z_R.$$

The ex ante probability of producing the public good is then

$$\mathbb{P}(A | \Delta z_R) = \frac{1}{2} + \frac{1}{ch} \Delta z_R.$$

Note that $z_R = \max\{-ch/2, \min\{z, ch/2\}\} = \min\{z, ch/2\}$, implying $\mathbb{E} z_R = \mathbb{E} \min\{z, ch/2\}$. Furthermore, from the definition of z , there is a natural upper bound $z \leq \varepsilon$. Thus, Lemma 3 applied to $\mathbb{E} \min\{z, ch/2\}$ implies

$$\Delta z_R \geq \min\{\Delta z, ch/2\} \geq \Delta z - \frac{\varepsilon - k}{k^2} 1_{\{k < \varepsilon\}}.$$

Combining,

$$\mathbb{P}(A | \Delta z_R) \geq \frac{1}{2} + \frac{1}{ch} \Delta z - \frac{1}{ch} \frac{\varepsilon - k}{k^2} 1_{\{k < \varepsilon\}}.$$

Lemma 1 implies $\varepsilon = \frac{ch}{2\beta}$ where $\beta(\theta) = \frac{1}{\theta} + \frac{\theta-1}{\theta} |z_{\theta}| + \left(\frac{\theta-1}{\xi(\theta)\theta^2} (z_{\theta}) \right)^{1/2}$. Furthermore, $\mathbb{E} z = \mathbb{E} z_{\theta}$. Thus, taking $k = \xi(\theta)^{-1/2}$,

$$\mathbb{P}(A | \Delta z_R) \geq \frac{1}{2} + \frac{1}{2} \beta^{-1} z_{\theta} - \frac{\beta^{-1} - 2(\xi^{1/2} ch)^{-1}}{\xi^{-1}} 1_{\{\xi^{-1/2} < \varepsilon\}},$$

giving the desired bound. Note that $\Theta\xi \rightarrow \infty$ implies $\frac{\Theta-1}{\xi\Theta^2}(z_\theta) \rightarrow 0$ as $\Theta \rightarrow \infty$, which implies $\beta(\Theta) \rightarrow \mathbb{E}z_\theta$. Additionally, $\xi(\Theta) \rightarrow 0$ as $\Theta \rightarrow \infty$ implies

$\frac{\beta^{-1}-2(\xi^{1/2}ch)^{-1}}{\xi^{-1}}1_{\{\xi^{-1/2}<\varepsilon\}} \rightarrow 0$. Thus, for h sufficiently large so that $\mathbb{E}z_\theta(h) > 0$,

$$\mathbb{P}(A|\Delta z_R) \geq \frac{1}{2} + \frac{1}{2}\beta^{-1}z_\theta - \frac{\beta^{-1}-2(\xi^{1/2}ch)^{-1}}{\xi^{-1}}1_{\{\xi^{-1/2}<\varepsilon\}} \rightarrow 1.$$

Fig. 1. The NGA Mechanism

Figure 1 shows a graphical representation of how the NGA mechanism works with $P=2$ policy parameters and $\Theta=2$ individuals. The circle shown is of radius $\varepsilon > 0$. The dashed vectors represent the contributions of the two individuals, with each of their votes multiplied by $1/2$, their fraction in the population. The mechanism's output, x , is then the sum of these vectors.

Fig. 2. The Restricted NGA Mechanism

Figure 2 graphically depicts the restricted NGA mechanism. The restricted policy space, R , is shown as the area below the line. The unrestricted NGA mechanism, as shown, would result in a new (in this case impossible) policy outside of the restricted policy space. This point is then projected back down onto the boundary of the restricted set (the orange line). The final restricted NGA output is the vector from x_0 to the point within R that is nearest to $x_0 + \Delta x$, the latent policy the unrestricted mechanism would deliver.

Fig. 3. Restricted NGA With Multiple Referenda

Figure 3 depicts the resulting latent policy adjustment vector $\mathbf{x}$ for values x_i . As shown, if the latent policy change $x_i \notin [0, 1]$ for every i , then $\mathbf{x}$ will be outside of the shaded cross (depicted as the solid arrow), and the closest point onto which the mechanism will project is a vertex. As a result, the outcome is deterministic. If $x_i \notin [0, 1]$ for some but not all i (as could result from an intermediate step size), then $\mathbf{x}$ will stay within the shaded cross region and will be projected onto the closest point within the restricted region (the middle square), resulting in a probabilistic outcome for at least one referendum. As long as the sum of votes is nonzero in every dimension, there exists a sufficiently large step size ϵ to ensure that the latent policy-change vector $\mathbf{x}$ will be outside of the shaded region.

Fig. 4. Introducing the Public Good

Figure 4 depicts person 1's wealth, w_1 , versus person 2's wealth, w_2 . Each person wants to maximize their wealth by moving in a direction parallel to his or her wealth vector. The public good dimension is shown vertically, where moving upward indicates an increase in the probability of building the public good. The dotted line indicates the line at which the public good will be built. The thin solid line is the line at which there is a 50% chance of building the public good. The height of the figure is $h = \frac{\sqrt{2}}{2}$. Height is a key parameter of the mechanism, which we will change later on.

Fig. 5. Gradients for Person 1 and Person 2

Figure 5 depicts the gradients before normalization, when they are linear functions of the valuation. Notice that the outside line traces out all possible gradients for different valuations, and the gradients are equidistant from one another as they hit the outside line. The normalized gradient, shown by the shorter arrows which end when each vector hits the edge of the semicircle, is non-linear in the valuation, but close to linear as long as the individual is close to indifferent about whether or not the public good is built (i.e., when the gradient is close to horizontal). This nonlinearity of the normalized gradient is an important departure from QV, where the actions are linear in the valuation. But in the limiting case in subsection 5.4, the normalized gradients become linear in the valuations because they become very close to horizontal.

Fig. 6. Cases: Will the Public Good be Produced?

Figure 6 shows four unique scenarios in which person 1 and person 2 have different valuations of the public good. In each case, the sum of vectors is depicted with varying step sizes, ϵ , corresponding to different arrow heads. In case 1, the sum of person 1's and person 2's valuations is large enough that the public good will be produced. In case 2, the sum of the vectors hits the no-public-good line with a large enough step size. In case 3, despite person 2's strong dislike of the public good, the public good will be produced. Case 4 shows a knife-edge scenario in which the public good has a 50% chance of being produced because person 1's and person 2's gradients are equal and opposite.

Fig. 7. Transforming the Policy Space

Figure 7 illustrates graphically the effects of adjusting height on the NGA mechanism adapted to the QV context. As the height, h , increases, the indifference curves become more nearly vertical. Because the gradients are orthogonal to the indifference curves, the gradients become more horizontal, implying little voting in the public good dimensions. Conversely, squishing the policy space vertically makes the indifference curves more nearly horizontal and thus makes the gradients more nearly vertical, implying that NGA votes focus on the public good dimension.

Fig. 8. Policy Space with Θ People

Figure 8 serves as an intuitive explanation for what the vertical distance represents. For each individual, the vertical distance shows the effect (proportional to $\frac{\varepsilon}{\Theta}$) of changing her report on the probability of building the public good. From the gradient formula

$\nabla H_{\theta} = \frac{\varepsilon}{\Theta} (1, \frac{1}{ch}(u_{\theta} - \frac{c}{\Theta}))$, we see that as h goes to 0, the mechanism becomes majority voting. On the other end of the limit, as h approaches ∞ , the mechanism becomes QV.