Approximate implementation of Relative Utilitarianism via Groves-Clarke pivotal voting with virtual money

Marcus Pivato

30. October 2006

Approximate implementation of Relative Utilitarianism via Groves-Clarke pivotal voting with virtual money

Marcus Pivato
Department of Mathematics, Trent University*

October 30, 2006

Abstract

Relative Utilitarianism (RU) is a version of classical utilitarianism, where each person's utility function is rescaled to range from zero to one. As a voting system, RU is vulnerable to preference exaggeration by strategic voters. The Groves-Clarke *Pivotal Mechanism* elicits truthful revelation of preferences by requiring each voter to 'bid' a sum of real money to cast a pivotal vote. However, this neglects wealth effects and gives disproportionate power to rich voters. We propose a variant of the Pivotal Mechanism using fixed allotments of notional 'voting money'; this *Voting Money Pivotal Mechanism* (VMPM) is politically egalitarian and immune to wealth effects. In the large-population limit, the only admissible (i.e. weakly undominated) voting strategies in the VMPM are approximately truthful revelations of preferences; thus the VMPM yields an arbitrarily close approximation of RU.

Let $\mathcal{I}$ be a set of individuals and let $\mathcal{A}$ be a set of policy alternatives. Suppose that each $i \in \mathcal{I}$ has an ordinal preference relation over $\mathcal{A}$ and also over the set of all lotteries between elements in $\mathcal{A}$. If these lottery preferences satisfy the von Neumann-Morgenstern (vNM) axioms of minimal rationality, then we can define a cardinal utility function $u_i : \mathcal{A} \longrightarrow \mathbb{R}$ such that i's lottery preferences are consistent with maximization of the expected value of u_i.

A *utilitarian* social welfare function $U : \mathcal{A} \longrightarrow \mathbb{R}$ is one of the form

$$U(a) \quad := \quad \sum_{i \in \mathcal{I}} c_i u_i(a), \qquad \forall\ a \in \mathcal{A}, \tag{1}$$

where $c_i \in \mathbb{R}_{\not -}$ are nonnegative constants. Classic Utilitarianism prescribes the policy alternative $a^* \in \mathcal{A}$ which maximizes U. Utilitarianism has several philosophically appealing mathematical properties, such as those given by Harsanyi's (1953) Impartial Observer Theorem and (1955) Social Aggregation Theorem. It has also been characterized

*1600 West Bank Drive, Peterborough, Ontario, Canada, K9J 7B8; email: `marcuspivato@trentu.ca`.

by d'Aspremont and Gevers (1977), Maskin (1978), Myerson (1981) and Ng (1975, 1985, 2000) as the only social welfare function satisfying several combinations of axioms encoding 'fairness' and 'rationality'.

However, the von Neumann-Morgenstern cardinal utility functions u_i in eqn.(1) are only well-defined up to affine transformations: if $c \in \mathbb{R}_{\neq}$ and $d \in \mathbb{R}$ are constants, then the functions u_i and $\widetilde{u}_i := cu_i + d$ are equally consistent descriptions of the lottery preferences of individual i. Thus, Classic Utilitarianism suffers from two major problems:

1. There is no *a priori* reason to choose one set of constants $\{c_i\}_{i\in\mathcal{I}}$ versus another. Equivalently, there is no practical method for making accurate *interpersonal comparisons* of utility. (Indeed, it is not clear that such interpersonal comparisons could be well-defined, even in principle).

2. A classical utilitarian voting system is vulnerable to manipulation by 'strategic voters' who either exaggerate their preferences or misrepresent them in more subtle ways.

One solution to problem #1 is to insist that everyone rescale their personal utility function so that its range lies in a certain compact interval. Typically, all utilities are rescaled to range over a unit interval (e.g. from zero to one). In other words, for all $i \in \mathcal{I}$, we define $r_i := \max_{a\in\mathcal{A}} u_i(a) - \min_{a\in\mathcal{A}} u_i(a)$. We then substitute $c_i := 1/r_i$ in eqn.(1). This version of utilitarianism has been called *Relative Utilitarianism* (RU), and admits several appealing axiomatic characterizations; see Cao (1982), Dhillon (1998), Karni (1998), Dhillon and Mertens (1999) and Segal (2000).

However, RU is still susceptible to strategic misrepresentation of preferences. The scope for exaggeration of utilities is limited, but if the electorate is large, then each voter might try to maximize the influence of her vote by declaring a value of 'one' for all the alternatives she finds acceptable, and value 'zero' to all the alternatives she finds unacceptable (especially on a hard-fought issue). In this case RU devolves into the 'approval voting' of Brams and Fishburn (1983). Approval voting has many nice properties, but it does not satisfy the same axiomatic characterizations as RU. Furthermore, approval voting is an 'ordinal' voting system, so the impossibility theorem of Gibbard (1973) and Satterthwaite (1975) makes it susceptible to further forms of strategic voting.

If the space $\mathcal{A}$ of alternatives is a convex set of feasible allocations of economic resources, then Sobel (2001) has shown that the set of Nash equilibria of the resulting 'utility misrepresentation game' for RU contains the set of Walrasian equilibria of a pure exchange economy over these resources with equal initial endowments. However, the misrepresentation game also admits non-Walrasian Nash equilibria which are not even Pareto efficient.

In this paper, we introduce the *Voting Money Pivotal Mechanism* (VMPM), a version of the well-known Groves-Clarke pivotal mechanism which uses a virtual currency of 'voting money'. This voting money attains value because each voter must reuse her finite budget of voting money in a long sequence of consecutive referenda. Voters can still misrepresent their utilities in the VMPM, but a rational voter would only use a misrepresentation strategy which was *admissible* in the sense that it was not weakly dominated by some other strategy. We show that, when all voters adopt admissible voting strategies, the outcome of the VMPM is an approximation of Relative Utilitarianism; furthermore, this approximation becomes arbitrarily accurate as the number of voters increases to infinity.[1]

The paper is organized as follows. In §1 we review the standard Groves-Clarke pivotal mechanism and discuss its properties and drawbacks. In §2 we introduce the VMPM, and in §3 we prove our main result (Theorem 3). In §4, we sketch a simple protocol to protect the anonymity of voters. In §5, we propose a system to formulate the referendum ballots themselves using auctions priced in voting money; we then sketch a formal model of this 'Ballot Auction' system. In §6 we review previous work on 'point-based' voting systems, and contrast these systems with the VMPM.

1 The Groves-Clarke Pivotal Mechanism

The *Groves-Clarke Pivotal Mechanism* (GCPM) is a special case of the *demand-revealing mechanism* proposed by Groves (1973) and Clarke (1971), and later promoted by Tideman and Tullock (1976).[2] The GCPM is a hybrid between a referendum and an auction:

1. Each voter i assigns a monetary *valuation* $v_i(a)$ to each alternative $a \in \mathcal{A}$. We regard $v_i(a)$ as a proxy for the value of $u_i(a)$ in eqn.(1).

2. Society chooses the alternative $a \in \mathcal{A}$ which maximizes the aggregate valuation:

$$V(a) \quad := \quad \sum_{i \in \mathcal{I}} v_i(a). \tag{2}$$

3. Suppose that voter i is *pivotal*, meaning that alternative a wins only because of i's vote. In other words, $V(a) - V(b) < v_i(a) - v_i(b)$, so if i had voted differently (i.e. given a higher valuation to b and/or a lower one to a), then the alternative b would have won instead. Then voter i must pay a *Clarke tax* t_i defined:

$$t_i \quad := \quad \sum_{j \neq i} [v_j(b) - v_j(a)]. \tag{3}$$

Intuitively, $[v_j(b) - v_j(a)]$ is the 'net loss' in utility for voter j because society chose a instead of b; hence the Clarke tax t_i is the 'aggregate net loss' for everyone else besides i.

[1] This is similar in spirit to other results about approximate implementation in large-population limits, such as Roberts and Postlewaite (1976). However, our paper differs both in its goal and its methods.

[2] The GCPM is extensively analyzed in the collection by Tideman (1977) and the monograph by Green and Laffont (1979). See also §8.2 of Moulin (1988), §23.C of Mas-Colell et al. (1995), §5 of Tideman (1997), and §8.1 of Mueller (2003). Another special case of Groves' and Clarke's demand-revealing mechanism is the Vickrey (1961) auction; for this reason the demand-revealing mechanism is sometimes called the 'Vickrey-Groves-Clarke mechanism'.

Note that

$$\begin{aligned} t_i &= \sum_{j\neq i} v_j(b) - \sum_{j\neq i} v_j(a) \quad = \quad [V(b) - v_i(b)] - [V(a) - v_i(a)] \\ &= [v_i(a) - v_i(b)] - [V(a) - V(b)] \quad \leq \quad v_i(a) - v_i(b), \end{aligned}$$

(because $V(a) \geq V(b)$ by hypothesis). Thus, the Clarke tax never exceeds i's personal gain in obtaining a rather than b (assuming she expressed her preferences honestly); hence i should always be willing to pay the tax t_i in order to secure alternative a.

In most cases, the winning alternative will win by a margin of victory which far exceeds the valuation assigned by any single voter, so that step #3 will only rarely be implemented. However, in a very close electoral outcome, many voters may find themselves in the position of the 'swing' voter described in step #3 (i.e. each one could have single-handedly changed the outcome), and in these cases, all these voters must pay a Clarke tax.

Because of this possibility, each voter has a strong incentive to express her preferences honestly. If she understates her preference for a particular alternative, then she runs the risk that a less-preferred alternative may be chosen, even though she *could have* changed the outcome to her more preferred alternative had she voted honestly (and would have happily paid the resulting Clarke tax). Conversely, if she *overstates* her value for a particular alternative, then she risks paying more than it is worth for her to 'purchase' her preferred outcome. Thus, the GCPM acts as a kind of 'auction', where each valuation $v_i(a)$ functions not only as a 'vote', but also as a 'bid' for the option to change the referendum outcome. In most cases (e.g. landslide victories), this option will not be exercised, but in a close race, the option *will* be exercised, and the voter must pay her bid value. Just as in an ordinary auction, each voter neither wishes to 'underbid' (and risk unnecessary defeat) nor to 'overbid' (and risk paying too much). Her dominant strategy is always to bid honestly.

Formally, we can model the GCPM as a ***Bayesian game***, in which each player $i \in \mathcal{I}$ has a (secret) utility function $u_i : \mathcal{O} \longrightarrow \mathbb{R}$ (where $\mathcal{O}$ is some set of ***outcomes***), along with a ***strategy set*** $\mathcal{S}_i$, and the outcome of the game is determined by a function $o : \prod_{i\in\mathcal{I}} \mathcal{S}_i \longrightarrow \mathcal{O}$. Let $\mathcal{S}_{-i} := \prod_{j\in\mathcal{I}\setminus\{i\}} \mathcal{S}_j$, and regard o as a function $o : \mathcal{S}_i \times \mathcal{S}_{-i} \longrightarrow \mathcal{O}$. We say $s_i \in \mathcal{S}_i$ is a ***dominant strategy*** for player i if, for any $\mathbf{s}_{-i} \in \mathcal{S}_{-i}$,

$$u_i\left[o(s_i, \mathbf{s}_{-i})\right] \quad \geq \quad u_i\left[o(s_i', \mathbf{s}_{-i})\right], \quad \forall\ s_i' \in \mathcal{S}_i.$$

In other words, s_i is an optimal strategy for player i, given any possible choice of strategies for the other players.

Let $\mathcal{V} := \mathbb{R}^{\mathcal{A}} = \{v : \mathcal{A} \longrightarrow \mathbb{R}\}$ be the set of all monetary ***valuations*** of the alternatives in $\mathcal{A}$. Consider the Bayesian game where $\mathcal{S}_i = \mathcal{V}$ for all i (each player's strategy is to declare some valuation in $\mathcal{V}$), and where the outcome of the game is a choice of policy in $\mathcal{A}$, and some Clarke tax for each player, as determined by the GCPM. In other words, $\mathcal{O} := \mathcal{A} \times \mathbb{R}^{\mathcal{I}}$, and for any vector of valuations $\mathbf{v} = (v_1, \ldots, v_I) \in \prod_{i \in \mathcal{I}} \mathcal{S}_i$, $o(\mathbf{v}) := (a; \mathbf{t})$, where $a \in \mathcal{A}$ is the alternative with the highest total valuation, and $\mathbf{t} := (t_1, \ldots, t_I) \in \mathbb{R}^{\mathcal{I}}$ is the vector of Clarke taxes computed using eqn.(3). Suppose that (after perhaps multiplying by a constant), each voter's utility function has the ***quasilinear*** form

$$u_i(a, -t_i) \quad = \quad w_i(a) - t_i, \qquad \forall a \in \mathcal{A} \text{ and } t_i \in \mathbb{R}, \tag{4}$$

where $w_i : \mathcal{A} \longrightarrow \mathbb{R}$ is her utility function over the policy alternatives and t_i is the Clarke tax she must pay. Then it makes sense to say that $w_i(a)$ is the monetary ***worth*** which voter i assigns to alternative $a \in \mathcal{A}$. Given assumption (4), the GCPM is a ***dominant strategy implementation*** of utilitarianism in the following sense:

Theorem 1 *Suppose all voters have quasilinear utility functions like eqn.*(4)*. Then for each $i \in \mathcal{I}$, a dominant strategy is to set $v_i := w_i$. In the resulting dominant strategy equilibrium, the GCPM chooses the same alternative as utilitarianism* (because then maximizing $V = \sum_{i \in \mathcal{I}} v_i$ is equivalent to maximizing $U = \sum_{i \in \mathcal{I}} w_i$).[3] □

The GCPM also satisfies other appealing axiomatic characterizations, due to Moulin (1986) and Sjostrom (1991). However, because it links voting to money, the GCPM has several major caveats:

Caveat #1. Theorem 1 only holds if voters have quasilinear utility functions like eqn.(4). This is false. Real people are *risk-averse*, which means their utility is highly *concave* as a function of money. At the very least, we should assume utility functions have the 'quasiconcave' form

$$u_i(a, t_i) \quad = \quad w_i(a) + c(E_i + t_i), \qquad \forall a \in \mathcal{A} \text{ and } t_i \in \mathbb{R}, \tag{5}$$

where c is some concave function (e.g. $c = \log$) and E_i is the initial ***endowment*** of player i (i.e. her current assets, plus the expected present value of all future earnings). But this leads to further problems:

(a) If c is strictly concave, then the GCPM clearly assigns much more 'voting power' to rich people than poor people. A rich person i might easily be willing to bid \$100,000 to change the outcome of the election from a to b, whereas a poor person j would only bid \$100 to change it from b to a, even though $w_i(a) = w_j(b)$ and $w_i(b) = w_j(a)$.

(b) If c is nonlinear, then Theorem 1 is false; indeed, a voter may not have *any* dominant strategy. For example, suppose $\mathcal{A} = \{a, b, c\}$, and

$$w_i(a) \quad = \quad 0 \quad < \quad w_i(b) \quad = \quad 2 \quad < \quad w_i(c) \quad = \quad 4.$$

[3] See Proposition 23.C.4 on p.877 of Mas-Colell et al. (1995) or Lemma 8.1 on p.204 of Moulin (1988).

Suppose c is a concave function such that $c(E_i) = 0$, $c(E_i - \$2) = -2$ and $c(E_i - \$3) = -4$. Thus, voter i would be willing to pay a \$2 Clarke tax to change outcome a to outcome b, and also \$2 to change outcome b to outcome c, but would only be willing to pay \$3 to change a to c. Suppose voter i declares valuations $v_i(a) = 0$ and $v_i(c) = 3$ (which is a truthful expression of her quasiconcave utility function with respect to a and c). What valuation should she declare for b? If she declares $v_i(b) < 2$, then she has 'undervalued' b versus a; if a ultimately wins by a margin of less than \$2 over b, then she will regret her choice. However, if she declares $v_i(b) > 1$, then she has 'overvalued' b versus c; if b ultimately wins by a margin of less than \$2 over c, then she will still regret her choice.

Suppose, then, that i declares $v_i(a) = 0$, $v_i(b) = 2$, and $v_i(c) = 4$; then she will be satisfied with any referendum outcome of a vs. b or b vs. c. But suppose c beats a by a margin between \$3 and \$4; then i will have to pay a Clarke tax greater than \$3, so once again she will regret her choice. In summary, there is *no* valuation of the alternatives $\{a, b, c\}$ which i will not regret under some circumstances. Her best strategy depends upon her expectations about how other people will vote. In other words, she has no dominant strategy.

In this situation, one or more Nash equilibria may still exist (some of which may even be truth-revealing). But the predictive relevance of a Nash equilibrium depends upon each voter making accurate predictions about the behaviour of every other voter, and in a 'voting game' involving millions of voters (e.g. a modern democracy) this is not very plausible.

(c) Like the quasilinear function (4), the quasiconcave function (5) 'solves' the problem of interpersonal utility comparison by implicitly assuming that all people have *identical* utility function c for money. This is false. Even if two people have the same initial endowment, their utility for money may differ. For example, a person with modest material needs (e.g. an ascetic monk) will assign less utility to each dollar than a hedonistic playboy. Hence we should assume each person's utility function has the form $u_i(a, t_i) = w_i(a) + c_i(E_i + t_i)$, where c_i is some concave function which may differ from person to person. This further confounds any interpretation of the aggregate monetary valuation of an alternative as its 'aggregate utility'.

Good (1977) has proposed a modified pivotal scheme which equalizes voting power between rich and poor or between ascetics and hedonists. Loosely speaking, we redefine $V(a) := \sum_{i \in \mathcal{I}} f_i[v_i(a)]$ in eqn.(2), where $f_i[t] := c_i[E_i] - c_i[E_i - t]$ measures the disutility of t lost dollars for voter i (for example, if the function c_i is linear with slope λ_i, then this simplifies to $V(a) := \sum_{i \in \mathcal{I}} \lambda_i v_i(a)$, where presumably the marginal utilities λ_i are smaller for rich people and larger for poor people). The problem, of course, is to estimate the functions f_i; clearly each person has considerable incentive to misrepresent her marginal utility. Good proposes we use some standard function like $f_i(t) = t/E_i$, but this seems somewhat procrustean. Also, the proof that Good's mechanism is a dominant-strategy truthful implementation of utilitarianism still implicitly assumes that voter's utility functions are quasilinear, so it is vulnerable to Caveat #1(b).

Tideman (1997) has proposed that Clarke taxes be paid in *time* (spent, say, in community service) rather than money. This gives the poor the same *a priori* political power as the rich, but it is still far from egalitarian. Different people value their time very differently. The retired and the unemployed have a lot of spare time (and hence, presumably, assign a low marginal utility to this time), whereas working parents and jet-setting professionals have almost no time to spare.

(d) Even the individualized quasiconcave utility functions in #1(c) assume that each person's preferences over the alternatives in $\mathcal{A}$ are totally *separable* from her wealth level E_i. This is false. For example, rich people and poor people have very different preferences concerning redistributive taxation schemes and publicly funded goods.

Caveat #2. Any revenue collected by the Clarke tax must be removed from the economy (e.g. destroyed or donated to a faraway country), because otherwise voters who expect *not* to pay a Clarke tax have an incentive to distort their valuations so as to inflate the amount of revenue which is collected; see (Riker, 1982, p.54) for example. Thus, the GCPM is never Pareto-efficient.

Caveat #3. As (Riker, 1982, p.56) notes, pivotal voting *cannot* be anonymous, because to implement the Clarke tax, we need a public record of each person's valuations of the alternatives. However, anonymity of voting is a crucial feature of modern democracy. Anonymity protects voters from discrimination and political extortion, and also prevents voters from selling their votes for material gain. The GCPM is clearly vulnerable to a scam where I pay a thousand people $5 each to declare a valuation of $100 for a particular outcome. If this outcome then wins by a 'landslide' margin of $100,000 (or indeed, by any margin greater than $100), then none of my accomplices needs to pay the Clarke tax (so they each profit $5), and the total cost for me is only $5,000 (which is much cheaper than personally paying a $100,000 Clarke tax to swing the outcome in my favour).

2 The Voting Money Pivotal Mechanism

To resolve these problems, we propose a modified Pivotal Mechanism, where people post anonymous valuations using a virtual, nontransferable currency we call 'voting money'. We call this the *Voting Money Pivotal Mechanism* (VMPM):

- Everyone in society starts with exactly the same initial endowment of voting money (e.g. everyone starts with 100 'voting cents', which make one 'voting dollar'). Your voting money is held in an an anonymous[4] escrow account, from which it is not transferable (so you can't sell, lend, or give your votes to anyone else). The only way you can 'spend' voting money is by paying a Clarke tax. The only way you obtain voting money is that, after you pay a Clarke tax, your voting money account is automatically but very gradually replenished (e.g. at a rate of one cent per week) until it returns to its initial amount (e.g. \$1.00).

- Voting money can be used to vote in regular (e.g. weekly) referenda to decide public policy questions.

- When you vote, the amount currently in your account must be sufficient to pay any Clarke tax you might incur —i.e. it must equal the largest difference between your valuations of any pair of alternatives. (In particular, this means we can assume that your valuations for each alternative are always between \$0 and \$1.00).

- After each referendum, the Electoral Commission (EC) aggregates the valuations of all voters, and computes the total valuation for each alternative. The EC then publicly announces these totals. As in the GCPM, the alternative with the highest total valuation is chosen.

- The EC computes Clarke taxes for each voter, as in eqn.(3). The EC then debits your voting money account of any Clarke tax you owe.

In §3, we will show that each voter assigns nonzero value to her voting money (because she must reuse it in successive referenda), so that Clarke taxes create a nontrivial incentive for truth revelation. Furthermore, we will show that, under reasonable assumptions, each voter's utility function is approximately quasilinear in voting money, which yields Theorem 3, an approximate version of Theorem 1. First, we will briefly discuss how the VMPM obviates some of the caveats we earlier identified for the Groves-Clarke pivotal mechanism.

To neutralize Caveat #2, note that all Clarke taxes are paid in a virtual currency with no real economic value (the only value of voting money is its influence on VMPM referenda). Thus, the 'destruction' of Clarke tax revenues does not imply any Pareto inefficiency in real economic resources.

To dispel Caveat #1(a), observe that there is no disparity in voting power between rich and poor, because everyone has the same initial 'wealth' in voting money and the same 'future earning potential'. Someone may temporarily become 'poor' if she pays a large Clarke tax, but this 'poverty' disappears fairly quickly because her account is automatically replenished at some constant rate. To obviate Caveat #1(d), note that a person's endowment of voting money has no influence on her policy preferences, because her voting money endowment has no relation to her role in society or her real economic status. Finally, to neutralize Caveat #1(c), note that voting money can only be spent on political action, not on physical goods. A single voting cent represents exactly the same

[4] See §4.

amount of political power for each person, and all voters receive the same endowment of voting money. Thus, if all voters as have the same intensity of preferences over policy alternatives (which is the implicit assumption of Relative Utilitarianism), then a single voting cent has exactly the same utility to for each voter.

We will address Caveat #3 in §4. Caveat #1(b) is addressed by Lemma 2 in the next section.

3 Approximate implementation of Relative Utilitarianism

It remains to show that each voter assigns nontrivial (but approximately quasilinear) utility to her voting money, so that something like Theorem 1 is true. To begin, note that the value of a quantity of voting money is the expected increase in utility it yields when spent in a strategically optimal manner to influence the outcome of future referenda. We will argue that, in the limit as the population size I tends to infinity, the expected utility gained in this fashion from a quantity x of voting money is a linear function of x.

After the current referendum $\mathcal{A}_0$, there will be an infinite sequence of further referenda with alternative sets $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \ldots$. Voter i is trying to decide how much money to risk on the current referendum $\mathcal{A}_0$, and how much to reserve for these future referenda. Recall that, whenever i's voting money account is depleted by a Clarke tax, it is slowly replenished up to some maximum —say E voting dollars. Suppose that, after each referendum, a depleted account is replenished with E/N voting dollars, for some $N \in \mathbb{N}$.[5] Thus, any Clarke tax which i pays in referendum $\mathcal{A}_0$ will be totally replenished after at most N referenda; thus, her voting strategy in $\mathcal{A}_0$ will affect her future voting power only in referenda $\mathcal{A}_1, \ldots, \mathcal{A}_N$.

Assume the alternative sets $\mathcal{A}_1, \ldots, \mathcal{A}_N$ are disjoint, and let $\mathcal{A} := \bigsqcup_{n=1}^{N} \mathcal{A}_n$. If $m > 0$, then any element of $[0, m]^{\mathcal{A}}$ represents some valuation of between 0 and m voting dollars for each alternative of the future referendum $\mathcal{A}_n$, for every $n \in [1...N]$. Thus, if society has I individuals, each with an initial endowment of E dollars, then the set of possible collective valuations is $[0, IE]^{\mathcal{A}}$. For simplicity, we could set $E := 1$. However, to represent the proportionate influence of each individual voter (which becomes infinitesimally small as I becomes large), it is convenient to treat each voter as having an endowment of $E := 1/I$, so that the set of collective valuations is normalized to $[0, 1]^{\mathcal{A}}$.

Thus, a vector in $[0, 1]^{\mathcal{A}}$ represents an aggregate valuation of each element of $\mathcal{A}$ by society, and hence, implicitly determines an outcome for each of the next N referenda. Let $\mathcal{B} := \prod_{n=1}^{N} \mathcal{A}_n$. Thus, an element $\mathbf{b} = (b_1, b_2, \ldots, b_N) \in \mathcal{B}$ represents a 'bundle' of policies determined by outcomes of the next N referenda. We assume that i has vNM utility function $\beta_i : \mathcal{B} \longrightarrow \mathbb{R}$. Define the *outcome* function $o : [0, 1]^{\mathcal{A}} \longrightarrow \mathcal{B}$ so that, for every

[5]For a discussion of how to set the parameter N, see Remark 9(R)[i] below.

bundle $\mathbf{b} \in \mathcal{B}$,

$$o^{-1}\{\mathbf{b}\} \quad := \quad \left\{\mathbf{x} \in [0,1]^{\mathcal{A}} \ ; \ \forall\, n \in [1...N], \ \forall\, a \in \mathcal{A}_n, \ x_{b_n} \geq x_a\right\}$$

is the set of all aggregate valuations such that alternative b_1 wins the first referendum, b_2 wins the second referendum, etc. Define $\widetilde{\beta}_i := \beta_i \circ o : [0,1]^{\mathcal{A}} \longrightarrow \mathbb{R}$; hence $\widetilde{\beta}_i(\mathbf{x})$ is the utility for voter i of the bundle of future referendum outcomes determined by collective valuation $\mathbf{x}$.

Voter i cannot completely predict the outcome of the next N referenda, but she can make an educated guess, based on her knowledge of the distribution of preferences in her society. Thus, let $f_i : \mathbb{R}^{\mathcal{A}} \longrightarrow \mathbb{R}_{\not-}$ be a probability density function (supported on $[0,1]^{\mathcal{A}}$), such that, for any measurable subset $\mathbf{Y} \subset [0,1]^{\mathcal{A}}$, $\int_{\mathbf{Y}} f(\mathbf{y})\ d\mathbf{y}$ is the subjective probability (as seen by voter i) that aggregate valuation of all other voters will lie in $\mathbf{Y}$. In particular, this means that her subjective expected utility (before voting) is given by

$$\int_{\mathbb{R}^{\mathcal{A}}} \widetilde{\beta}_i(\mathbf{x}) f_i(\mathbf{x})\ d\mathbf{x}. \tag{6}$$

We make the following assumptions:

(D) *f_i is twice differentiable, and $\|\partial_a \partial_b f_i\|_1 < \infty$ for all $a, b \in \mathcal{A}$.* (Here, $\|g\|_1 := \int_{\mathbb{R}^{\mathcal{A}}} |g(\mathbf{x})|\ d\mathbf{x}$).

(S) *Let $u_i : \mathcal{A}_0 \times \mathcal{B} \longrightarrow \mathbb{R}$ be individual i's vNM utility function over all bundles of policies in the referenda $\mathcal{A}_0, \mathcal{A}_1, \ldots, \mathcal{A}_N$. Then u_i has the 'separable' form:*

$$u_i(a_0, a_1, \ldots, a_n) \quad = \quad \alpha_i(a_0) + \beta_i(a_1, \ldots, a_n), \tag{7}$$

where $\alpha_i : \mathcal{A}_0 \longrightarrow \mathbb{R}$ is her vNM utility function for referendum $\mathcal{A}_0$ and $\beta_i : \mathcal{B} \longrightarrow \mathbb{R}$ is her vNM utility function for all the remaining referenda, as above.

(I) *For each voter i, the functions f_i, α_i and β_i do not depend on I (the number of other voters).*

Lemma 2 *Assume* **(D)**, **(S)**, *and* **(I)**. *For all $a_0 \in \mathcal{A}_0$ and $t \in [0,1]$, let $U_i(a, -t)$ be voter i's cardinal utility when society chooses policy a_0 and i pays a Clarke tax of t voting dollars. Then*

$$U_i(a_0, -t) \quad = \quad w_i(a_0) \ - \ m_i(t), \tag{8}$$

where $w_i : \mathcal{A}_0 \longrightarrow \mathbb{R}$ is a scalar multiple of α_i, and where $m_i : [-1,1] \longrightarrow \mathbb{R}$ is an increasing function such that $m_i(0) = 0$ and

$$m_i(t) \quad = \quad t \ \pm \ \mathcal{O}(h_i t^2 / I). \tag{9}$$

Here, I is the number of voters; $h_i \geq 0$ is constant not depending on I; and "$\mathcal{O}(h_i t^2/I)$" represents some function such that $|\mathcal{O}(h_i t^2/I)| \leq |h_i t^2/I|$ for all $t \in [-1,1]$ and $I \in \mathbb{N}$. $\square$

Heuristically speaking, if $\mathcal{I}$ is sufficiently large (e.g. a democracy of several million people), then the error term $\mathcal{O}(h_i t^2/I)$ will be extremely small, so Lemma 2 says each voter's utility function u_i will be very close to the quasilinear form (4). Hence, $w_i(a)$ roughly measures the 'worth' (in voting money) of alternative a to voter i. However, this doesn't mean that her dominant strategy will be truthful revelation (as in Theorem 1), because if u_i diverges even a little from being quasilinear, then voter i might *have no* dominant strategy [as in Caveat #1(b)]. Nevertheless, if I is large, then the only *admissible* strategies for i are the 'approximately honest' ones (see Corollary 6 below). To show this, we require one further assumption about the range of w_i and m_i:

(R) *Suppose i has the utility function $u_i(a, -t) = w_i(a) - m_i(t)$ as in eqn.(8). Then $|w_i(a) - w_i(b)| \in m_i[0, 1]$ for all $a, b \in \mathcal{A}$.*

If the hypotheses of Lemma 2 are satisfied, then $m_i(t) \approx t$ for all $t \in [0, 1]$, so that $m_i[0, 1] \approx [0, 1]$. In this case, assumption **(R)** is roughly equivalent to stipulating that $\max_{a\in\mathcal{A}} w_i(a) - \min_{a\in\mathcal{A}} w_i(a) \leq 1$.

A ***valuation strategy*** (or 'voting strategy') for voter i in the VMPM is a function $v_i : \mathcal{A}_0 \longrightarrow [0, 1]$ which declares a valuation (in voting money) for each alternative in $\mathcal{A}_0$. (Presumably this valuation should reflect i's utility function w_i, but it may not.) Thus, i's strategy set is $\mathcal{V}_i := [0, 1]^{\mathcal{A}}$. Let $\mathcal{V}_{-i} := \prod_{i\neq j\in\mathcal{I}} \mathcal{V}_j$. For any $\mathbf{v}_{-i} = (v_j)_{\substack{j\in\mathcal{I}\\ j\neq i}} \in \mathcal{V}_{-i}$ and $a \in \mathcal{A}$, we define $\mathbf{v}_{-i}(a) := \sum_{i\neq j\in\mathcal{I}} v_j(a)$. Let a be the element of $\mathcal{A}$ maximizing $\mathbf{v}_{-i}(a)$ [i.e. the outcome of a referendum involving everyone except i]. If $v_i \in \mathcal{V}_i$, then let b be the element of $\mathcal{A}$ maximizing $v_i(b) + \mathbf{v}_{-i}(b)$, and let $t := \mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(b)$ be the resulting Clarke tax (if $a = b$, then $t = 0$). Thus,

$$U_i(v_i, \mathbf{v}_{-i}) \quad := \quad w_i(b) - m_i(t)$$

is the utility of the valuation strategy v_i for voter i, given the profile $\mathbf{v}_{-i}$ of opposing strategies. If $\overline{v}_i, v_i \in \mathcal{V}_i$, then $\overline{v}_i$ ***weakly dominates*** v_i if, for any $\mathbf{v}_{-i} \in \mathcal{V}_{-i}$, $U_i(\overline{v}_i, \mathbf{v}_{-i}) \geq U_i(v_i, \mathbf{v}_{-i})$, and this inequality is strict for some $\mathbf{v}_{-i} \in \mathcal{V}_{-i}$. That is: regardless of the strategies of the other voters, the outcome of $\overline{v}_i$ (i.e. the referendum winner and resulting Clarke tax) is always at least as good for i as the outcome of v_i, and is strictly better than v_i for some possible strategies by other voters.

The valuation strategy $\overline{v}_i$ is ***admissible*** if it is *not* weakly dominated by any other valuation strategy. This is a weak condition, and voter i may have many admissible strategies. However, it is clearly irrational for her to use any *inadmissible* strategy. For example, suppose $v_i \in \mathcal{V}_i$ was inadmissible, and weakly dominated by $\overline{v}_i$. Let ρ be any probability measure on $\mathcal{V}_{-i}$. Then

$$\int_{\mathcal{V}_{-i}} U_i(v_i, \mathbf{v}_{-i})\ d\rho[\mathbf{v}_{-i}] \quad \leq \quad \int_{\mathcal{V}_{-i}} U_i(\overline{v}_i, \mathbf{v}_{-i})\ d\rho[\mathbf{v}_{-i}].$$

If ρ has full support on $\mathcal{V}_{-i}$, then this inequality is strict. Hence i's expected utility is never maximized by an inadmissible valuation strategy.

If $\epsilon > 0$ and $x, y \in \mathbb{R}$, then "$x \underset{\epsilon}{\sim} y$" means $|x - y| < \epsilon$. Our main result is this:

Theorem 3 *Assume* **(D)***,* **(S)***,* **(I)** *and* **(R)***, and fix $\epsilon > 0$. For all $i \in \mathcal{I}$, let voter i have utility function $w_i : \mathcal{A}_0 \longrightarrow \mathbb{R}$ as in Lemma 2, and let $\overline{v}_i : \mathcal{A}_0 \longrightarrow [0, 1]$ be any valuation strategy that is admissible for i. Let*

$$W \;:=\; \frac{1}{I} \sum_{i \in \mathcal{I}} w_i : \mathcal{A}_0 \longrightarrow \mathbb{R} \qquad \textit{and} \qquad V \;:=\; \frac{1}{I} \sum_{i \in \mathcal{I}} \overline{v}_i : \mathcal{A}_0 \longrightarrow [0, 1]$$

be, respectively, the average utility function and the average valuation of all voters. If $\mathcal{I}$ is sufficiently large, then:

- **(a)** *There is a constant $K \in \mathbb{R}$ such that $W(a) \underset{\epsilon}{\sim} V(a) + K$ for all $a \in \mathcal{A}_0$.*
- **(b)** *Thus, if $a^* \in \mathcal{A}_0$ maximizes W and $b^* \in \mathcal{A}_0$ maximizes V, then $V(a^*) \underset{\epsilon}{\sim} V(b^*)$.*

In other words, in the large population limit, the VMPM induces all voters to declare valuations such that the outcome yields an average social utility within ϵ of the optimal outcome specified by Relative Utilitarianism. (Presumably, much of the time, the referendum outcome is actually identical to that of RU. However, the outcome itself is less important than the aggregate utility level which we achieve.)

Proof of Lemma 2: Suppose voter i had an unlimited budget, and could declare any valuations she wanted for the alternatives in the future referenda $\mathcal{A}_1, \ldots, \mathcal{A}_N$. If voter i declared a valuation vector $\mathbf{v} \in \mathbb{R}^{\mathcal{A}}$, then she would translate the density function f_i in expression (6) by the vector $\mathbf{v}$; hence her expected utility (after voting) would be

$$\int_{\mathbb{R}^{\mathcal{A}}} \widetilde{\beta}_i(\mathbf{x}) f_i(\mathbf{x} - \mathbf{v}) \; d\mathbf{x}.$$

This means that the expected gain in her utility obtained by declaring $\mathbf{v}$ would be

$$\begin{aligned} \gamma_i(\mathbf{v}) \;&:=\; \int_{\mathbb{R}^{\mathcal{A}}} \widetilde{\beta}_i(\mathbf{x}) f_i(\mathbf{x} - \mathbf{v}) \; d\mathbf{x} - \int_{\mathbb{R}^{\mathcal{A}}} \widetilde{\beta}_i(\mathbf{x}) f_i(\mathbf{x}) \; d\mathbf{x} \\ &=\; \int_{\mathbb{R}^{\mathcal{A}}} \widetilde{\beta}_i(\mathbf{x}) \left[f_i(\mathbf{x} - \mathbf{v}) - f_i(\mathbf{x}) \right] \; d\mathbf{x}. \end{aligned}$$

This defines a nonlinear, smooth function $\gamma_i : \mathbb{R}^{\mathcal{A}} \longrightarrow \mathbb{R}$, which we will approximate with a first-order Taylor expansion. Observe that $\gamma_i(0) = 0$. Define

$$\mathbf{g} \;:=\; \nabla \gamma_i(0) \;=\; -\int_{\mathbb{R}^{\mathcal{A}}} \widetilde{\beta}_i(\mathbf{x}) \nabla f_i(\mathbf{x}) \;\in\; \mathbb{R}^{\mathcal{A}}, \qquad \text{and} \qquad \widehat{h}_i \;:=\; \max_{a,b \in \mathcal{A}} \|\partial_a \partial_b \gamma_i\|_\infty.$$

(Here, $\|\gamma\|_\infty := \sup_{\mathbf{x}\in\mathbb{R}^{\mathcal{A}}} |\gamma(\mathbf{x})|$). Assumption **(D)** says that $\mathbf{g}$ is well-defined. Also,

$$\widehat{h}_i \quad \underset{(*)}{\leq} \quad \left\|\widetilde{\beta}_i\right\|_\infty \cdot \max_{a,b\in\mathcal{A}} \|\partial_a\partial_b f\|_1 \quad \underset{(\dagger)}{<} \quad \infty.$$

where $(*)$ is the Hölder inequality and $(\dagger)$ is by assumption **(D)**.

Now recall that i only has a small endowment, and that $1/I$ is the weight of i's endowment, divided by the population size. Then her set of feasible future valuation vectors is actually $\mathbf{V} := [0, 1/I]^{\mathcal{A}}$ (assuming she never has to pay a Clarke tax). If $\mathbf{v} \in \mathbf{V}$ and $t \in \mathbb{R}$, then the multivariate Taylor inequality and assumption **(D)** imply:

$$\gamma_i(t\mathbf{v}) \quad = \quad \gamma_i(0) + \langle t\mathbf{v}, \nabla\gamma_i(0)\rangle \pm \mathcal{O}(\widehat{h}_i t^2|\mathbf{v}|^2) \quad = \quad t\langle \mathbf{v}, \mathbf{g}\rangle \pm \mathcal{O}(\widehat{h}_i t^2/I^2). \qquad (10)$$

Thus, for any $t \in [0, 1]$, the optimal deployment of t voting dollars for i is the valuation vector $t\overline{\mathbf{v}}$, where $\overline{\mathbf{v}} \in \mathbf{V}$ is the vector in $\mathbf{V}$ which maximizes $t\langle\overline{\mathbf{v}}, \mathbf{g}\rangle$. Let $||\mathbf{g}||_\infty := \max_{a\in\mathcal{A}} |g_a|$; then $\overline{\mathbf{v}} := \mathbf{g}/(I||\mathbf{g}||_\infty)$ (independent of t). Let $||\mathbf{g}||_2^2 := \sum_{a\in\mathcal{A}} |g_a|^2$ and let $k_i := ||\mathbf{g}||_2^2/||\mathbf{g}||_\infty$; then $\langle\overline{\mathbf{v}}, \mathbf{g}\rangle = k_i/I$, so $t\langle\overline{\mathbf{v}}, \mathbf{g}\rangle = tk_i/I$.

Thus, $\widehat{m}_i(t) := \gamma_i(t\overline{\mathbf{v}})$ is the utility for i of t voting dollars: the increase in expected utility when t dollars are deployed optimally in future referenda (voting money has no other value, because it is not transferable and cannot be spent on physical goods). Substituting into eqn.(10), we get

$$\widehat{m}_i(t) \quad := \quad \gamma_i(t\overline{\mathbf{v}}) \quad = \quad k_i t/I \pm \mathcal{O}(\widehat{h}_i t^2/I^2).$$

This, combined with hypothesis **(S)**, yields a combined utility function $\widehat{U}_i : \mathcal{A}_0 \times \mathbb{R} \longrightarrow \mathbb{R}$ of the form

$$\widehat{U}_i(a_0, -t) \quad = \quad \alpha_i(a_0) + \widehat{m}_i(-t) \quad = \quad \alpha_i(a_0) - k_i t/I \pm \mathcal{O}(\widehat{h}_i t^2/I^2).$$

Finally, since vNM cardinal utilities are equivalent up to affine transformations, we can divide $\widehat{U}_i$ and $\widehat{m}_i$ by the scalar (k_i/I) to get an equivalent utility function $U_i(a_0, -t) = w_i(a_0) - m_i(t)$, where $w_i(a_0) := I\alpha_i(a_0)/k_i$ and $m_i(t) := -t \pm \mathcal{O}(h_i t^2/I)$, and $h_i := \widehat{h}_i/k_i$. □

If the quasilinear approximation in Lemma 2 was exact (i.e. if $m_i(t) = t$ for all $t \in [0, 1]$), then i's 'truthful' valuation strategy would be $v_i = w_i + k$ (where k is any constant), so that

$$v_i(a) - v_i(b) \quad = \quad w_i(a) - w_i(b) \qquad \text{for all } a, b \in \mathcal{A}.$$

More generally, if $u_i(a, -t) = w_i(a) - m_i(t)$ as in Lemma 2, then a 'truthful' valuation strategy would be any v_i such that

$$v_i(a) - v_i(b) \quad = \quad m^{-1}[w_i(a) - w_i(b)], \qquad \text{for all } a, b \in \mathcal{A}. \qquad (11)$$

[Recall that m_i is an increasing function, hence invertible, so assumption **(R)** implies that the right-hand side of eqn.(11) is well-defined.] However, if m_i is a nonlinear function, and $\mathcal{A}$ has more than two elements, then it is generally impossible for v_i to simultaneously satisfy eqn.(11) for all possible pairs $a, b \in \mathcal{A}$ —see Caveat #1(b). Fortunately, Lemma 2 says that m_i is still 'approximately' linear, so we could still try to 'approximately' satisfy eqn.(11). If $v_i \in \mathcal{V}_i$, then for any $a, b \in \mathcal{A}$, we define the *discrepancy*

$$\delta v_i(a,b) \quad := \quad v_i(a) - v_i(b) - m^{-1}\left[w_i(a) - w_i(b)\right].$$

Heuristically, i's 'honest' valuation strategy is one which minimizes these discrepancies. Let $v_i, \overline{v}_i \in \mathcal{V}_i$. We write "$\delta\overline{v}_i \lll \delta v_i$" if, for all $a, b \in \mathcal{A}$,

$$\operatorname{sign}(\delta\overline{v}_i(a,b)) \;=\; \operatorname{sign}(\delta v_i(a,b)) \quad \text{and} \quad |\delta\overline{v}_i(a,b)| \;\leq\; |\delta v_i(a,b)|. \tag{12}$$

We write "$\delta\overline{v}_i \ll \delta v_i$" if this inequality is strict for some $a, b \in \mathcal{A}$.

Lemma 4 *Let $v_i, \overline{v}_i \in \mathcal{V}_i$. If $\delta\overline{v}_i \ll \delta v_i$, then $\overline{v}_i$ weakly dominates v_i.*

Proof: Let $\mathbf{v}_{-i} \in \mathcal{V}_{-i}$. For any $a \in \mathcal{A}$, recall that $\mathbf{v}_{-i}(a) := \sum_{i \neq j \in \mathcal{I}} v_j(a)$. Let a be the element of $\mathcal{A}$ maximizing $\mathbf{v}_{-i}(a)$, and let $c \in \mathcal{A}$, $c \neq a$. Thus, $\mathbf{v}_{-i}(a) \geq \mathbf{v}_{-i}(c)$. Let b be the element of $\mathcal{A}$ maximizing $v_i(b) + \mathbf{v}_{-i}(b)$ and let $\overline{b}$ be the element of $\mathcal{A}$ maximizing $\overline{v}_i(\overline{b}) + \mathbf{v}_{-i}(\overline{b})$. There are three cases.

Case 1: $(\mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c) > 1)$

In this case, i's vote can make no difference to the outcome, so $\overline{b} = a = b$ and hence $U_i(\overline{v}_i, \mathbf{v}_{-i}) = w_i(a) = U_i(v_i, \mathbf{v}_{-i})$.

Case 2: $(m_i^{-1}\left[w_i(c) - w_i(a)\right] < \mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c) \leq 1)$

In this case, either $w_i(c) < w_i(a)$ [i.e. i actually prefers a over c] or $w_i(a) > w_i(c)$, but i's utility improvement from a to c is not worth the Clarke tax of $[\mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c)]$ that she must pay to achieve it. Hence i's goal is *not* to overbid. Equation (12) implies that the quantity $[\overline{v}_i(c) - \overline{v}_i(a)]$ is always closer to $m_i^{-1}\left[w_i(c) - w_i(a)\right]$ than the quantity $[v_i(c) - v_i(a)]$. Thus, one of four subcases occurs:

$$\begin{array}{lccccccc}
\text{(2a)} & v_i(c) - v_i(a) & \leq & \overline{v}_i(c) - \overline{v}_i(a) & \leq & m_i^{-1}\left[w_i(c) - w_i(a)\right] & \leq & \mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c). \\
\text{(2b)} & m_i^{-1}\left[w_i(c) - w_i(a)\right] & \leq & \overline{v}_i(c) - \overline{v}_i(a) & \leq & v_i(c) - v_i(a) & \leq & \mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c). \\
\text{(2c)} & m_i^{-1}\left[w_i(c) - w_i(a)\right] & \leq & \overline{v}_i(c) - \overline{v}_i(a) & \leq & \mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c) & \leq & v_i(c) - v_i(a). \\
\text{(2d)} & m_i^{-1}\left[w_i(c) - w_i(a)\right] & \leq & \mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c) & \leq & \overline{v}_i(c) - \overline{v}_i(a) & \leq & v_i(c) - v_i(a).
\end{array}$$

These four subcases yield the following outcomes:

(2a) and (2b): $\overline{b} = a = b$, and there are no Clarke taxes, so $U_i(\overline{v}_i, \mathbf{v}_{-i}) = w_i(a) = U_i(v_i, \mathbf{v}_{-i})$.

(2c): $\overline{b} = a$ but $b = c$. Thus, $\overline{v}_i$ incurs no tax, but v_i incurs a tax of $\mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c)$, so

$$U_i(\overline{v}_i, \mathbf{v}_{-i}) \quad = \quad w_i(a) \quad \underset{(*)}{>} \quad w_i(c) - m_i\left[\mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c)\right] \quad = \quad U_i(v_i, \mathbf{v}_{-i}).$$

Here, $(*)$ is because $w_i(c) - w_i(a) < m_i\left[\mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c)\right]$, because m_i is increasing and because $m_i^{-1}\left[w_i(c) - w_i(a)\right] < \mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c)$ by the hypothesis of Case 2.

(2d): $\overline{b} = c = b$, so both $\overline{v}_i$ and v_i incur the same Clarke tax, so that $U_i(\overline{v}_i, \mathbf{v}_{-i}) = w_i(c) - m_i\left[\mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c)\right] = U_i(v_i, \mathbf{v}_{-i})$.

In every subcase, $U_i(\overline{v}_i, \mathbf{v}_{-i}) \geq U_i(v_i, \mathbf{v}_{-i})$, and in subcase (2c) this inequality is strict.

Case 3: $\left(m_i^{-1}\left[w_i(c) - w_i(a)\right] \geq \mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c) \geq 0\right)$

In this case, $w_i(c) \geq w_i(a)$, and the utility improvement for i from a to c is worth the Clarke tax of $\left[\mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c)\right]$ that she must pay to achieve it. Hence i's goal is not to *under*bid. Equation (12) implies that one of four subcases occurs:

$$\begin{array}{lccccccc}
\text{(3a)} & v_i(c) - v_i(a) & \geq & \overline{v}_i(c) - \overline{v}_i(a) & \geq & m_i^{-1}\left[w_i(c) - w_i(a)\right] & \geq & \mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c). \\
\text{(3b)} & m_i^{-1}\left[w_i(c) - w_i(a)\right] & \geq & \overline{v}_i(c) - \overline{v}_i(a) & \geq & v_i(c) - v_i(a) & \geq & \mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c). \\
\text{(3c)} & m_i^{-1}\left[w_i(c) - w_i(a)\right] & \geq & \overline{v}_i(c) - \overline{v}_i(a) & \geq & \mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c) & \geq & v_i(c) - v_i(a). \\
\text{(3d)} & m_i^{-1}\left[w_i(c) - w_i(a)\right] & \geq & \mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c) & \geq & \overline{v}_i(c) - \overline{v}_i(a) & \geq & v_i(c) - v_i(a).
\end{array}$$

These four subcases yield the following outcomes:

(3a) and (3b): $\overline{b} = c = b$, so both $\overline{v}_i$ and v_i incur the same Clarke tax, so $U_i(\overline{v}_i, \mathbf{v}_{-i}) = w_i(c) - m_i\left[\mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c)\right] = U_i(v_i, \mathbf{v}_{-i})$.

(3c): $\overline{b} = c$ but $b = a$. Thus, v_i incurs no tax, but $\overline{v}_i$ incurs a tax of $\mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c)$, so

$$U_i(\overline{v}_i, \mathbf{v}_{-i}) \quad = \quad w_i(c) - m_i\left[\mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c)\right] \quad \underset{(*)}{\geq} \quad w_i(a) \quad = \quad U_i(v_i, \mathbf{v}_{-i}).$$

here $(*)$ is because $w_i(c) - w_i(a) \geq m_i\left[\mathbf{v}_{-i}(a) - \mathbf{v}_{-i}(c)\right]$, because m_i is increasing and by the hypothesis of Case 3.

(3d): $\overline{b} = a = b$, and there are no Clarke taxes, so $U_i(\overline{v}_i, \mathbf{v}_{-i}) = w_i(a) = U_i(v_i, \mathbf{v}_{-i})$.

Again, in every subcase, $U_i(\overline{v}_i, \mathbf{v}_{-i}) \geq U_i(v_i, \mathbf{v}_{-i})$. □

The relation $\lll$ is a partial ordering on $\mathcal{V}_i$. A valuation strategy is called *discrepancy-minimizing* if it is minimal with respect to $\lll$.

Lemma 5 *For any $v_i \in \mathcal{V}_i$, there is some discrepancy-minimizing $\overline{v}_i \in \mathcal{V}_i$ such that $\overline{v}_i \lll v_i$.*

Proof: The ordering $\lll$ is continuous with respect to the topology of $\mathcal{V}_i = [0,1]^{\mathcal{A}}$. Hence for any $v_i \in \mathcal{V}_i$, the set $\mathbf{C} = \{v'_i \,;\, v'_i \lll v_i\}$ is closed. Furthermore, $\mathcal{V}_i$ is compact, so $\mathbf{C}$ is compact, so $\mathbf{C}$ has a $\lll$-minimal element. $\square$

Corollary 6 *Any admissible valuation strategy is discrepancy-minimizing.*

Proof: If $v_i \in \mathcal{V}_i$ is not discrepancy-minimizing. Then Lemma 5 yields some $\overline{v}_i \in \mathcal{V}_i$ with $\overline{v}_i \ll v_i$. Then Lemma 4 says that $\overline{v}_i$ weakly dominates v_i, so v_i is not admissible. $\square$

It follows that a rational voter will only declare discrepancy-minimizing valuations. Note that i may prefer different discrepancy-minimizing valuations, depending upon her subjective probability estimates for the valuations of other players. However, as I becomes large, every discrepancy-minimizing valuation looks very close to w_i plus a constant, and in this sense, the VMPM induces 'honest' voting by all voters:

Proposition 7 *Assume* **(D)**, **(S)** *and* **(I)**, *so that i has vNM utility function $U_i(a,-t) = w_i(a) - m_i(t)$ as in Lemma 2. Also assume* **(R)**, *and fix $\epsilon > 0$. If I is large enough, then for any discrepancy-minimizing $\overline{v}_i \in \mathcal{V}_i$, there is some constant $k_i \in \mathbb{R}$ so that $\overline{v}_i(a) \underset{\epsilon}{\sim} w_i(a) + k_i$ for all $a \in \mathcal{A}$.*

Proof: Let $A := \#(\mathcal{A})$, let $\eta := \frac{\epsilon}{2(3A+1)}$, and let $\theta := \eta/3$. If I is large enough, then Lemma 2 implies that $m_i(t) \underset{\theta}{\widetilde{\quad}} t$, for all $t \in [0,1]$. It follows immediately that

$$m_i^{-1}(r) \quad \underset{\theta}{\widetilde{\quad}} \quad r, \qquad \text{for all } r \in m_i[0,1]. \tag{13}$$

From this we deduce the following 'approximate cocycle' properties for discrepancies:

Claim 1: *Let $v \in \mathcal{V}_i$. Then:*

(a) *For all $a,b,c \in \mathcal{A}$, $\delta v(a,c) \underset{\eta}{\widetilde{\quad}} \delta v(a,b) + \delta v(b,c)$.*

(b) *For all $a,b \in \mathcal{A}$, $\delta v(a,b) \underset{\eta}{\widetilde{\quad}} -\delta v(b,a)$.*

Proof: (a) $\delta v(a,c) \ := \ v(a) - v(c) - m_i^{-1}[w_i(a) - w_i(c)]$

$$\begin{aligned}
&\underset{\theta}{\widetilde{\quad}} \quad v(a) - v(c) - [w_i(a) - w_i(c)] \\
&= \quad v(a) - v(b) + v(b) - v(c) - [w_i(a) - w_i(b) + w_i(b) - w_i(c)] \\
&= \quad v(a) - v(b) - [w_i(a) - w_i(b)] + v(b) - v(c) - [w_i(b) - w_i(c)] \\
&\underset{2\theta}{\widetilde{\quad}} \quad v(a) - v(b) - m_i^{-1}[w_i(a) - w_i(b)] + v(b) - v(c) - m_i^{-1}[w_i(b) - w_i(c)] \\
&=: \quad \delta v(a,b) + \delta v(b,c).
\end{aligned}$$

Here, both "$\sim$" are by eqn.(13). Thus, $\delta v(a,c) \underset{3\theta}{\widetilde{\quad}} \delta v(a,b) + \delta v(b,c)$. But $3\theta = \eta$.
(b) $\delta v(a,b) + \delta v(b,a) \underset{\eta}{\widetilde{\quad}} \delta v(a,a) \ = \ 0$ by (a). $\diamond$ `Claim 1`

If $v \in \mathcal{V}_i$, then a v-*partition* of $\mathcal{A}$ is a pair of nonempty disjoint subsets $\mathcal{B}, \mathcal{C} \subset \mathcal{A}$ such that $\mathcal{A} = \mathcal{B} \sqcup \mathcal{C}$, and such that, for all $b \in \mathcal{B}$ and $c \in \mathcal{C}$, $\delta v(b,c) > 0$.

Claim 2: *If $\mathcal{A}$ can be v-partitioned, then v is not discrepancy-minimizing.*

Proof: Suppose $\mathcal{A} = \mathcal{B} \sqcup \mathcal{C}$ is a v-partition. Then $\mu := \min\limits_{b\in\mathcal{B},\ c\in\mathcal{C}} \delta v(b,c) > 0$ (because $\mathcal{B}$ and $\mathcal{C}$ are finite). Define $v' \in \mathcal{V}_i$ as follows:

$$\begin{aligned} \forall\ b \in \mathcal{B}, \quad v'(b) &:= v(b) - \mu/2. \\ \forall\ c \in \mathcal{C}, \quad v'(c) &:= v(c). \end{aligned}$$

Thus, $\delta v'(b_1, b_2) = \delta v(b_1, b_2)$ for all $b_1, b_2 \in \mathcal{B}$, and $\delta v'(c_1, c_2) = \delta v(c_1, c_2)$ for all $c_1, c_2 \in \mathcal{C}$. But for any $b \in \mathcal{B}$ and $c \in \mathcal{C}$, we have $0 < \delta v'(b,c) = \delta v(b,c) - \mu/2 < \delta v(b,c)$. Thus, $\delta v' \ll \delta v$, so v is not discrepancy-minimizing. $\diamond$ `Claim 2`

Let $\overline{\mathcal{V}}_i(\epsilon) := \{v_i \in \mathcal{V}_i\ ;\ |\delta v_i(a,b)| < \epsilon/2,\ \forall\ a, b \in \mathcal{A}\}$.

Claim 3: *Every discrepancy-minimizing valuation is in $\overline{\mathcal{V}}_i(\epsilon)$.*

Proof: By contradiction, suppose $v \notin \overline{\mathcal{V}}_i(\epsilon)$. Then $v(b,c) \geq \epsilon/2$ for some $b, c \in \mathcal{A}$. Recall that $\eta = \epsilon/2(3A+1)$.

Claim 3.1: *There is some $k \in [1...A]$ such that $\mathcal{A} = \mathcal{B} \sqcup \mathcal{C}$, where*

$$\begin{aligned} \mathcal{C} &:= \{c' \in \mathcal{A}\ ;\ \delta v(c', c) < 3k\eta\}, \\ \textit{and}\ \ \mathcal{B} &:= \{b' \in \mathcal{A}\ ;\ \delta v(b', c) \geq 3(k+1)\eta\}. \end{aligned}$$

Proof: For all $k \in [1...A]$, let $\mathcal{A}_k := \{a \in \mathcal{A}\ ;\ 3(k-1)\eta \leq \delta v(a,c) < 3k\eta\}$. Note that b is not in any of these sets (because $v(b,c) \geq \epsilon/2 = (3A+1)\eta$). Thus, $\bigsqcup_{k=1}^{A} \mathcal{A}_k \subset \mathcal{A} \setminus \{b\}$ is a union of A disjoint sets, but it contains at most $A-1$ elements. Thus, the Pigeonhole Principle says $\mathcal{A}_k = \emptyset$ for some $k \in [1...A]$. Thus, if $\mathcal{B}$ and $\mathcal{C}$ are defined as above, then $\mathcal{A} = \mathcal{B} \sqcup \mathcal{C}$. $\triangledown$ `Claim 3.1`

Claim 3.2: *$\mathcal{B} \sqcup \mathcal{C}$ is a v-partition of $\mathcal{A}$.*

Proof: First note that both $\mathcal{B}$ and $\mathcal{C}$ are nonempty, because $b \in \mathcal{B}$ and $c \in \mathcal{C}$. Now, let $b' \in \mathcal{B}$ and $c' \in \mathcal{C}$. Then

$$\begin{aligned} \delta v(b', c') &\underset{\eta}{\sim} \delta v(b', c) + \delta v(c, c') \underset{\eta}{\sim} \delta v(b', c) - \delta v(c', c) \\ &> 3(k+1)\eta - 3k\eta = 3\eta, \end{aligned}$$

where the first "$\underset{\eta}{\sim}$" is by Claim 1(a), the second "$\underset{\eta}{\sim}$" is by Claim 1(b), and "$>$" is by definition of $\mathcal{B}$ and $\mathcal{C}$. Thus, $\delta v(b', c') > 3\eta - 2\eta = \eta > 0$, as desired. $\triangledown$ `Claim 3.2`

Claims 2 and 3.2 together imply that v is *not* discrepancy-minimizing. By contradiction, if v is discrepancy-minimizing, then $v \in \overline{\mathcal{V}}_i(\epsilon)$. $\diamond$ `Claim 3`

Now suppose $\overline{v}_i$ is discrepancy-minimizing. Fix $b \in \mathcal{A}$ and let $k_i := \overline{v}_i(b) - w_i(b)$. Then for any other $a \in \mathcal{A}$,

$$\begin{array}{rcl} \overline{v}_i(a) - (w_i(a) + k_i) & = & \overline{v}_i(a) - \overline{v}_i(b) - [w_i(a) - w_i(b)] \\ & \underset{\epsilon/2}{\widetilde{\quad}} & m_i^{-1}\,[w_i(a) - w_i(b)] - [w_i(a) - w_i(b)] \\ & \underset{\theta}{\widetilde{\quad}} & [w_i(a) - w_i(b)] - [w_i(a) - w_i(b)] \quad = \quad 0. \end{array}$$

Here, "$\underset{\epsilon/2}{\widetilde{\quad}}$" is because $\overline{v}_i \in \overline{\mathcal{V}}_i(\epsilon)$ by Claim 3; and "$\underset{\theta}{\widetilde{\quad}}$" is by eqn.(13).

Thus, $v_i(a) \underset{\epsilon}{\sim} w_i(a) + k_i$ for all $a \in \mathcal{A}$ (because $\epsilon/2 + \theta < \epsilon$). □

Proof of Theorem 3: Corollary 6 implies that any rational voter $i \in \mathcal{I}$ will declare a discrepancy-minimizing valuation $v_i \in \mathcal{V}_i$. Proposition 7 says that, for each $i \in \mathcal{I}$, there is some constant k_i such that $v_i \underset{\epsilon}{\sim} w_i + k_i$. Let $K := \frac{1}{I} \sum_{i \in \mathcal{I}} k_i$. The result follows. □

Remark 8: (a) The 'quasilinear approximation' in Lemma 2 depends on the idea that each voter can exert only an 'infinitesimal' influence over the referendum outcome. Thus, Theorem 3 does *not* apply if either:

1. The population of voters is relatively small, or

2. Some voter has a disproportionately large amount of voting money.

This has two implications:

1. Theorem 3 is only applicable to large-population referenda, and not to small committees.

2. It is vitally important to strictly limit the amount of voting money any single voter can accumulate. This is why we stipulated that voting money must be nontransferable, and it should only be obtainable through some government-provided 'income', and only up to some modest maximum endowment per voter.

(b) Note that we do *not* assume that the probability density function f_i in equation (6) is a correct or consistent description of the behaviour of other voters (we are not looking for a Bayesian Nash equilibrium). The point of Lemma 2 is that as long as voter i predicts the future voting behaviour of other voters with *some* probability density function satisfying **(D)**, **(S)**, and **(I)**, she will assign quasilinear utility to her voting money. Corollary 6 then says that, in fact, it doesn't really matter *what* prediction i makes about the other voters; her admissible valuation strategies will always be discrepancy-minimizing, and hence (by Proposition 7) very close to truthful preference revelation. (In particular, of course, this will be true in any Bayesian Nash equilibrium.)

(c) Strictly speaking, Theorem 3 only yields an 'ϵ-approximation' of Relative Utilitarianism. There are two other ways in which VMPM may deviate from RU:

- Not all voters will have the same endowment of voting money, because some may be temporarily impoverished by Clarke taxes levied in previous referenda. In this case, the VMPM yields a *weighted* utilitarianism, where each voter's 'weight' is her current endowment.

  Such impoverishment will be an infrequent, because Clarke taxes will rarely be levied except in very close races. However, even when such impoverishment does occur, it is arguably 'fair', because a voter would only be impoverished if she had recently been pivotal in one or more closely fought referenda.

- If voter i is somewhat apathetic about referendum $\mathcal{A}_0$, then it is likely that $\max_{a\in\mathcal{A}_0} w_i(a) - \min_{a\in\mathcal{A}_0} w_i(a) < 1$, whereas technically, RU requires all voters' utility functions to range over the entire unit interval. In this sense, the VMPM technically fails to implement RU when confined to a single referendum. However, if we instead consider voter i's 'extended' utility function $u_i : \prod_{n=0}^{N} \mathcal{A}_n \longrightarrow \mathbb{R}$ over all the set of all referenda simultaneously, then the VMPM effectively normalizes u_i to range over the unit interval (i.e. the size of voter i's endowment). Hence, in this long term sense, the VMPM does implement RU.

Remark 9: In Caveat #1, we complained that Theorem 1 requires the highly dubious assumption that all voters have quasilinear utility functions like (4) with respect to real money. We have replaced this dubious assumption with assumptions **(D)**, **(S)**, **(I)** and **(R)**. Are these assumptions not equally dubious?

(D) This is an assumption about voter i's subjective probability estimates concerning future referenda —in other words, it is an assumption about her *psychology.* To the extent that it is psychologically plausible to model a voter as rationally maximizing her expected utility with respect to some well-defined subjective probability density function, it seems just as plausible to assume that this subjective probability density functions has the nice properties specified by **(D)**.

(S) This assumption simply translates the 'separability' property of the monetary utility function (4) into an analogous separability property for joint utility functions over multiple political issues. Clearly such a separability assumption is false: a decision on one political issue may affect your stance on other issues. However:

- Although your joint utility function over multiple political issues is not completely separable, it is arguably much *more* separable than your joint utility function over political issues and personal wealth level.

- Governments must consider issues one at a time. Thus, eventually, every social choice mechanism must act as if people's preferences on different issues are separable.

If issue nonseparability is a serious concern, then one solution is to define the referenda over 'bundles' of policies which simultaneously encompass several nonseparable issues.

(I) This assumption just says that voter i understands the Law of Large Numbers. Based on her knowledge of her own culture, she estimates a certain subjective probability for statements like, "More than 60% of voters in my society would strongly prefer policy a to policy b". If $\mathcal{I}$ is a sufficiently large random sample of these voters, then voter i would estimate the *same* subjective probability for the statement, "More than 60% of voters in $\mathcal{I}$ would strongly prefer policy a to policy b". Her subjective probability estimate will be the same, whether $\mathcal{I}$ contains 100 000 voters or 100 000 000 voters. In other words, f_i does not depend on I.

(R) This assumption means that voter i can accurately express her preferences between the alternatives in $\mathcal{A}$ by using monetary valuations within her endowment limit $[0, 1]$. This is true as long as the constant k_i in the proof of Lemma 2 is large enough —in other words, as long as i assigns roughly the same importance to her participation in foreseeable future referenda as she does to the present referendum. If i feels much more intensely about the present referendum than she does about foreseeable future referenda, or if she believes that these future referenda will all be landslides where her own vote is irrelevant, then assumption **(R)** will fail. In this case, she will chose an 'extreme' valuation strategy on $\mathcal{A}_0$, which bids her whole endowment for some alternatives (and nothing for others), but this valuation will still be insufficient to adequately express the intensity of her preferences over $\mathcal{A}_0$.

To obviate this problem, we must make value of voting money large enough for each person that she can always fully convey the intensity of her preferences using valuations within her endowment. There are two ways to do this:

[i] By slowing the replenishment rate of depleted endowments, we can extend the 'referendum horizon' N of each voter far into the future. Thus, each voter can always foresee the possibility of future referenda in which her participation is at least as important to her as the present referendum. A Clarke tax incurred on the present referendum would inhibit her influence on these future referenda; this will make the constant k_i in Lemma 2 large enough to verify assumption **(R)**.

[ii] By making voting money fungible for some other kind of political influence, we can increase its value enough to verify assumption **(R)**.

Method [i] suggests a natural feedback mechanism: if we observe too many voters declaring 'extreme' valuations, then this means that assumption **(R)** is failing because these voters are discounting future referenda too much. This means their voting money endowments are being replenished too quickly, so we can slow the replenishment rate until the frequency of extreme valuations drops to an acceptable level.

Method [ii] suggests combining the VMPM with some other form of political participation. We will do this in §5.

4 Anonymity in the Pivotal Mechanism

Caveat #3 of the Groves-Clarke Pivotal Mechanism was its lack of anonymity. In this section we propose a protocol which uses public key cryptography and an intermediating layer of 'brokers' to to preserve the anonymity of pivotal mechanism voting. (This protocol applies equally to the GCPM and the VMPM, but obviously we have in mind the latter).

In a *public key cryptosystem*, a message is encrypted using an *encryption key* k_e, but can only be decrypted by using a different (but matching) *decryption key* $k_d \neq k_e$. Thus, even if k_e is public knowledge, k_d can be a secret which is only known by one person, say i. Thus, anyone can encrypt a message (using k_e) and send it to i (perhaps over an insecure channel), but only i can decrypt the message (using k_d). Encrypting a document using k_e is like putting it into a 'sealed envelope' which only i can open. The security of the cryptosystem is based on the fact that, even with knowledge of k_e, it is extremely difficult to reconstruct k_d —just as difficult as trying to decrypt an encrypted message by 'brute force'. A general introduction to public key cryptography is Schneier (1996); more technical introductions are Koblitz (1994) and Stinson (2006).

The anonymity protocol works as follows:

1. Every voter stores her voting money in an escrow account managed by a 'vote broker'. The broker's job is to protect the anonymity of her clients. We assume that the existence of many private, competing brokers, each of whom depends on her reputation for trustworthiness to attract clients.

2. Before each referendum, the Electoral Commission (EC) generates an encryption key k_e and corresponding decryption key k_d, and publishes k_e.

3. To vote, you record your valuations for each alternative in $\mathcal{A}$, and encrypt this document using k_e. You then pass this encrypted vote to your broker.

4. On your behalf, your broker sends your encrypted votes to the EC, along with the current balance of your voting money account (but does not reveal your identity).

5. The EC decrypts your vote using k_d. If the difference between your maximum valuation and minimum valuation is greater than the current balance of your voting money account, then the vote is rejected as invalid.

6. The Electoral Commission (EC) aggregates all the valid valuations it received from all brokers, and computes the total valuation for each alternative. The EC then publicly announces these totals. The alternative with the highest total valuation is chosen.

7. The EC computes Clarke taxes for each (anonymous) voter, as in eqn.(3), and communicates this information to the relevant brokers. On your behalf, your broker pays

the EC any Clarke taxes you owe, and then debits your escrow account accordingly.[6]

Because your vote is encrypted, no one can read your votes except the EC. Thus, your broker cannot reveal your votes to a third party. Thus, just as with a conventional secret ballot, it is not possible to 'sell' your vote (because the 'buyer' can't verify how you voted).

The EC generates a new key pair (k_e, k_d) for each referendum to maximize cryptographic security. (It is easy to generate many new key pairs, using a simple computation).

Because your broker communicates your vote anonymously to the EC, it is not possible for the EC to identify a particular vote with a particular voter. Thus, it is not possible for the government to identify and persecute people with particular political views.

If a single 'brokerage layer' is not deemed sufficient to protect voter anonymity, then we can interpose multiple layers. For example, each private broker can communicate her anonymized clients' valuations to a private 'metabroker', who then communicates these anonymized valuation bundles to a private 'metametabroker', who finally communicates them to the Electoral Commission. With sufficiently many anonymizing brokerage layers, it would be almost impossible to trace a particular valuation to a particular voter, even if some brokers were corruptible.

The gradual replenishment of voting money accounts is also anonymized. The EC reimburses each broker in small installments of voting money for each Clarke tax paid by that broker. The broker then passes these reimbursements on to the appropriate client.

Presumably, the per-capita operating costs of a broker will be small. Nevertheless, each voter could be given an (anonymous) government voucher with which to pay her brokerage costs, so that voting imposes no financial burden.

5 Ballot Auctions

The VMPM is an attractive mechanism for democratically choosing one policy from a ballot of alternatives. But how is this ballot created in the first place? Clearly the outcome of a referendum is strongly influenced by the wording of the question and the list of responses presented to the voters. An inferior or biased ballot will yield an inferior or biased outcome, no matter how 'democratic' the referendum process is itself.

One possibility is for the ballots to be designed by a 'Legislative Committee'. But how should we select the members of this Committee? The electoral systems presently used in most democracies are less than satisfactory, and frequently degenerate into puerile popularity contests and partisan publicity stunts. We could use the VMPM itself to elect the members of the Legislative Committee, but how would we nominate the list of candidates for this VMPM-based election? In modern democracies, most candidates are nominated by organized political parties, so that meaningful participation in government is restricted

[6] Note that, from the EC's public announcement, you can easily compute your own Clarke taxes, so your broker cannot cheat you.

to an exclusive club of loyal party apparatchiks.[7]

Instead of relying on a Legislative Committee to design the referenda, we could allow *any* citizen to propose a new referendum question, to add a new response to the ballot of an existing referendum, or to propose an amendment to an existing response (we will refer to any of these three possibilities as a ***ballot initiative***). However, this would lead to a bewildering profusion of thousands of referenda, each containing thousands of alternatives. Many referendum questions would be ambiguous or severely biased; many responses on each ballot would be incomplete, incoherent, poorly designed, and/or poorly articulated, and many more would be frivolous variations upon one another. In the VMPM, 'vote-splitting' is no longer a concern, but information overload is a universal problem. Furthermore, any interest group which lost a referendum could simply keep reintroducing new referenda on the same topic until it got its way.

We could require anyone proposing a ballot initiative to pay a fee (which is made as large as necessary to keep manageable the number of referenda and the number of responses on each referendum). However, this method obviously favours the ballot initiatives of the wealthy. Alternately, we could require each ballot initiative to be supported by petition with some minimum number of signatures. But signatures are cheap and easy to obtain (and difficult to verify).

A better method is to combine these two methods: we could implicitly require a 'petition' in support of each initiative, manifested by a 'fee' paid in *voting money.* For example, suppose the fee to add a new response to the ballot of an existing referendum was one hundred voting dollars (and each person's endowment was one dollar). To pay this fee, you would need to exhaust the voting money endowments at least one hundred people. Thus, to successfully add a new response to a ballot, you must either obtain extremely strong support from a coalition of one hundred people (expressed in an unfalsifiable manner by expending their entire political capital), or you must obtain a weaker expression of support from all members of a larger coalition (e.g. a donation of 10 'voting cents' from each of a thousand supporters). The fee to introduce an entirely new referendum would be considerably larger (say, one thousand voting dollars), so it would require an even larger show of public support.

Since these fees are paid in voting money, this system does not favour wealthy vested interests. Instead, it favours coalitions which can mobilize sufficient popular political support around their ballot initiative. The fees should be made large enough to keep the number of ballot initiatives manageable, without becoming so large that they altogether suppress popular political participation. Indeed, instead of charging a fixed fee, the best approach is to hold a public auction for the right to introduce a ballot initiative.

For example, suppose there are to be 52 referendums per year (i.e. one per week). Then every week, the Electoral Commission (EC) will hold an auction, and the winner of

[7] In principle, any citizen can nominate herself as an 'independent' candidate, by paying a candidacy fee and/or obtaining sufficiently many signatures on petition. In practice, without the backing of a political party and its well-financed campaign machine, she is almost certain to lose the election.

this auction will dictate the question for the referendum to be held exactly 26 weeks later. Furthermore, suppose that we decide that there should be an average of *ten* responses on the ballot of each referendum (perhaps more on some and less on others). Then in each weekly auction, the EC will sell ten 'response licenses', each of which gives its owner the right to add one response to the ballot of any referendum to be held between 10 and 20 weeks hence (this provides at least six weeks to prepare responses after a given referendum question is announced, and also provides at least ten weeks for voters to consider all the responses to each referendum). Any response licenses unsold in a particular week will remain available for sale in following weeks (hence, in some weekly auctions, there may be more than ten licenses available for sale). Presumably, concerned citizens will pool their voting money and form *coalitions* which bid in these auctions, obtain licenses, and then introduce various ballot initiatives.

A standard printed page contains about 4000 printed characters (counting all letters, punctuation, and whitespace). Suppose we decide that on average, referendum questions (including explanatory preamble, etc.) should be no more than 1000 printed characters in length (i.e. a quarter page), and that each response (including clarifying clauses and supplementary information) should, on average, be no more than 2000 characters in length (i.e. half a page). Thus, the ballot for an average referendum would be less than 21000 characters in total (i.e. five or six pages). In this case, we can initially give the winner of each referendum auction the right to dictate a 500 character question, and the winner of each response license the right to dictate a 1000 character response. During each weekly auction, the EC will sell an additional 500 'question characters' (which can be used to increase the length of a referendum question), and an additional 10000 'response characters' (which can be used to increase the length of a response). For example, a coalition proposing an 800 character referendum question must also purchase 300 question characters, while a coalition introducing a 1400 character response must likewise purchase 400 response characters. Any characters unsold during a particular week will remain available for sale in following weeks.

This 'Ballot Auction' system has three major advantages over referendum design by a Legislative Committee:

1. The Ballot Auction totally democratizes the process of referendum design, while still keeping the result manageable.

2. Like the VMPM itself, the Ballot Auction is politically egalitarian, and does not favour wealthy vested interests or party apparatchiks.

3. By making voting money fungible for the right to propose new legislation, we have increased the utility of voting money (perhaps greatly). This may obviate the problem raised in Remark 9(R) of §3.

Remark 10: (a) Note that the aforementioned length limitations only apply to the *text of the ballot* (i.e. the question and responses). First, the coalition behind a ballot initiative can publish as much supplementary information as it desires, clarifying, justifying, and promoting its initiative. It can create a web page, mail out leaflets, or buy advertising space. Second, two thousand characters will clearly not be sufficient to completely and precisely encode a complex piece of legislation. The two thousand character response which appears on the ballot should be seen as a legally binding declaration of intent to draft a longer and more detailed legal text. Following the referendum, this legal text should be drafted by a team of lawyers (hired by the government, but presumably approved by the coalition which originally proposed the winning response). When it is finished, this legal text should then approved by an independent tribunal of judges, who should determine whether it represents a fair and accurate legal expression of the original wording of the winning response.

(b) After the 'licenses' and 'characters' are sold in the initial auctions, we cannot allow them to be resold. A resale market would create the opportunity for speculators to accumulate disproportionately large quantities of voting money, thereby undermining the intended political egalitarianism of the VMPM.

(c) To avoid the possibility of coordination failure (see below), we stipulate that each coalition must return to its contributing members any voting money it fails to spend on their behalf. The purpose of a coalition is to purchase licenses by bidding in an auction. If the coalition fails in this purpose, then it spends no money, and it must return to each member her entire contribution.

A formal model of these ballot auctions is beyond the scope of this paper, but we will offer an outline. Such a formal model must address at least four questions:

(a) How do coalitions form?

(b) How is the text of a ballot initiative determined by a coalition?

(c) How do the auctions work?

(d) How does the use of voting money in ballot auctions affect the conclusion of Lemma 2, that each voter's utility function is 'approximately quasilinear' in voting money?

The main issue in **(a)** is whether nascent coalitions will encounter coordination failures or free riding. Remark 10(c) stipulated that each coalition must return any unused voting money to its contributors. In particular, no contributor risks losing any money if the coalition fails to purchase the licenses it requires due to insufficient funds. Thus, each potential contributor's decision about whether and how much to contribute is independent of her beliefs about the actions of other contributors. This eliminates the threat of a coordination failure.

Nevertheless, free riding may be a problem, because a successful ballot initiative is a public good: a supporter of this initiative will benefit equally from the coalition's success whether she personally contributes or not. However, this differs from a standard public good problem in three ways:

1. The success of a ballot initiative is all-or-nothing. Unlike, say, the success of a public park, it is not a continuously increasing function of the level of funding.

2. A contributor only pays if the initiative is successful; she pays nothing if it fails.

3. As discussed under **(b)** below, a contributor can exert some control over the text of the initiative, by making her contribution conditional on certain amendments. Thus, by contributing, she obtains some degree of 'ownership' (unlike a truly public good).

If coalitions suffer from free riding, then they will have smaller budgets, and will necessarily make smaller bids in auctions. However, to succeed in its ballot initiative, a coalition merely has to out-bid other, competing coalitions, so if all coalitions are equally afflicted by free riding, then no coalition is especially disadvantaged, and its success or failure in the auctions will be more or less the same as if there was no free riding at all.

The problem is that some coalitions may suffer from more free riding than others. For example, suppose ballot initiative A is favoured by a small, homogeneous, politically extreme group, while initiative B is favoured by a much larger, heterogeneous, and politically moderate group. In this case, Olson (1965) has observed that the coalition promoting B will likely suffer more free riding than the coalition promoting A. If Olson's analysis is applicable, then extremist minorities may exert disproportionate influence over the text of referendum ballots. This influence will perhaps be countered by the fact that the whole electorate (dominated by a moderate majority) will decide the outcomes of these referenda.

In the simplest model of **(b)**, some person or group first finalizes the text for a ballot initiative, and then tries to gather together a coalition to financially support this initiative. However, in a more realistic model, each coalition member may make her contribution conditional on certain amendments to the text of the initiative. Thus, the ultimate text of the initiative arises out of some process of multilateral bargaining within the coalition, where the bargaining strength of each coalition member is perhaps proportional the Shapley Value (1953) of her financial contribution towards the coalition's bid. A precise model of this bargaining process must explicitly represent how the utility function of each coalition member depends upon the text of the ballot initiative.

To answer **(c)**, one could draw on the theory of auctions; see McAfee and McMillan (1987) or Milgrom (1989). Presumably, the licenses should go to the coalitions which have the most intense political preferences; thus, the auction mechanism should elicit truthful revelation of preference intensity from each bidder. One obvious choice would be Vickrey's (1961) second-price, sealed bid auction.

To answer **(d)**, one must replace the utility function (7) with a utility function $u_i : \mathcal{A}_0 \times \widetilde{\mathcal{B}} \longrightarrow [0,1]$, where

$$\widetilde{\mathcal{B}} \quad := \quad \bigcup \left\{ \prod_{n=1}^{N} \mathcal{A}_n \; ; \; \begin{array}{c} \mathcal{A}_1, \ldots, \mathcal{A}_N \text{ is any sequence of possible future} \\ \text{referenda, with any questions and responses} \end{array} \right\}.$$

The success or failure of various ballot initiatives determines which possible sequence of future referenda is actually realized. Voter i can then vote on this particular sequence of referenda, using the VMPM. Her utility function m_i over voting money is thus induced by her utility function over $\widetilde{\mathcal{B}}$ in two ways:

(i) As in Lemma 2, her money enables her to participate in the VMPM, which contributes infinitesimally to the probabilities that her preferred alternatives will win the referenda.

(ii) Voter i's contributions help to determine the success or failure of ballot initiatives, which in turn determines which future referenda will actually occur.

Lemma 2 used Taylor's theorem to show that the utility function induced by **(i)** was 'approximately quasilinear'. A similar method could probably be applied to **(ii)**, if we assume that each contributor's effect on the success of a ballot initiative is small enough. The difficulty here is that $\widetilde{\mathcal{B}}$ is an infinite set, so that the space $[0,1]^{\widetilde{\mathcal{B}}}$ of voting money allocations is infinite-dimensional. This makes it difficult to compute an expected utility with an integral like eqn.(6), and means that Taylor's theorem doesn't apply without further technical restrictions. Furthermore, if $\widetilde{\mathcal{B}}$ is infinite, then the utility function $\beta_i : \widetilde{\mathcal{B}} \longrightarrow \mathbb{R}$ could be unbounded in hypothesis **(S)** of §3; we must exclude this possibility or else $\mathbf{g}$ (and thus, the function m_i) could become infinite.

6 Concluding remarks

Other point voting systems. Systems where citizens vote by allocating a budget of 'voting money' are at least a century old; the earliest known description is Charles Dodgson's (1873) 'Method of Marks'. Musgrave (1959) briefly sketched a system of 'point voting' [p.130-131], while Coleman (1970) suggested that a currency of 'fungible votes' could supersede vote-trading just as money superseded barter [§III(b), p.1084]. 'Point voting' was also suggested by Mueller (1971, 1973), and is implicit in 'probabilistic' voting schemes like Intriligator (1973) and Nitzan (1975), as well as in the 'Walrasian equilibrium' model of vote-trading proposed by Mueller (1967, 1973) and studied in Philpotts (1971, 1972) and Mueller et al. (1972).

However, without some mechanism to encourage honesty, each voter will misrepresent her preferences; this was recognized by Dodgson (1873) and Mueller (1973, 1977), and further studied by Laine (1977), Nitzan et al. (1980) and Nitzan (1985). For example, in allocating her voting money over the alternatives of a single ballot, each voter might simply pile all her money onto her most-preferred alternative amongst the subset of alternatives she considers most likely to win (in particular, she may not allocate *any* money towards her favourite alternative, if she considers it doomed to lose). Thus, her allocation will not accurately represent her utility function.

Allen (1977, 1982) proposed a 'modified method of marks' (MMM), where, instead of allocating a fixed 'budget' of voting money, voters can give each alternative any numerical score within a certain range (like the VMPM, but without Clarke taxes). Allen claimed his MMM was less susceptible to strategic misrepresentation than Dodgson's Method of Marks, but this was refuted by Hardin (1982). (Indeed, we earlier argued that the MMM would in fact devolve into 'approval voting').

Theorem 3 shows that the truth-revealing property of the Groves-Clarke mechanism eliminates this problem of strategic misrepresentation of preferences, as long as voters must reuse the same budget of voting money (minus any Clarke taxes) to vote on a sequence of consecutive referenda. Hylland and Zeckhauser (1979) have proposed another 'point-based' voting system which truthfully reveals each voter's preferences for public goods. In the Hylland-Zeckhauser system, each voter has a budget of 'points' which she can allocate towards voting for various public expenditures. The amount of government money spent on each public expenditure is then proportional to the sum of the *square roots* of the point scores it receives from all voters[8]. Like the Groves-Clarke mechanism, the Hylland-Zeckhauser mechanism makes it optimal for voters to truthfully reveal their preferences, and implements a utilitarian outcome. Like the VMPM (but unlike Groves-Clarke), the Hylland-Zeckhauser mechanism relies on voting money rather than real money, so it does not favour wealthy voters. However, the Hylland-Zeckhauser mechanism is only designed for allocating a finite budget of resources amongst various preapproved public expenditures; it is not appropriate for making discrete, all-or-nothing choices between policies or between government candidates. Also, the Hylland-Zeckhauser mechanism relies on an iterative process (where voters repeatedly receive feedback and modify their votes), and it is not guaranteed that this iterative process will converge.

Vote-trading. It is well-known that vote-trading ('logrolling') can yield Pareto-improving outcomes. Each voter buys votes on issues of greater importance to her by selling her vote on issues of lesser importance; hence vote-trading allows 'ordinal' voting systems like plurality vote to indirectly detect 'cardinal utility' information. This was observed by Buchanan and Tullock (1990 [1962]), Coleman (1966), Mueller (1973) and Schwartz (1975), and confirmed in computer simulations by Philpotts (1971, 1972) and Mueller et al. (1972). However, Tullock (1959) and Brams and Riker (1973) have constructed examples where vote trading perversely leads to Pareto-*inferior* outcomes[9], as well as 'cyclical trading' scenarios wherein three or more factions endlessly trade votes on three or more referenda without ever converging to equilibrium[10].

It is tempting to think that the voting money of the VMPM could be embedded in some kind of 'currency exchange', thereby implementing the 'Walrasian' vote-trading mechanism of Mueller (1967); this could enable Pareto-improving vote-trades while avoiding perverse

[8]See §8.3, p.170 of Mueller (2003) for details. See also §4 of Tideman (1997).

[9]See Chapt. 10 of Buchanan and Tullock (1990 [1962]), §6E, p.161 of Riker (1982), §4.6, p.137 of Brams (2003 [1975]), or p.13 of Jones (1988).

[10]See §4.7, p.140 of Brams (2003 [1975]) or §5.9, p.107 of Mueller (2003).

or cyclical outcomes. However, as observed by Philpotts (1971) and Mueller (1973), such a currency exchange is again vulnerable to strategic misrepresentation of preferences. In allocating her voting money between different referenda, each voter will simply concentrate all her money on the referendum where it yields the highest expected utility gain, either because this is the issue where she holds the strongest preferences, or because this is the most closely contested ballot, where she has the highest chance of casting a pivotal vote. Either way, her distribution of voting money will not accurately reflect the intensities of her preferences on the different issues.

Nevertheless, the VMPM provides a mechanism whereby each voter can divert her 'political capital' from issues of lesser importance to her, so as to concentrate it on issues of greater importance, thereby achieving an outcome similar to a Pareto-improving vote trade. For example, suppose there are two referenda, $\mathcal{A}_0$ and $\mathcal{A}_1$, and that, in equation (7), voter i's utility function $u_i : \mathcal{A}_0 \times \mathcal{A}_1 \longrightarrow \mathbb{R}$ can be written as $u_i(a_0, a_1) = w_0(a_0) + w_1(a_1)$, where $w_0 : \mathcal{A}_0 \longrightarrow [0, 1]$ and $w_1 : \mathcal{A}_1 \longrightarrow [0, 1]$ are her utility functions over the two referenda considered separately. Suppose that her preferences over referendum $\mathcal{A}_1$ are 'ten times as intense' as her preferences over $\mathcal{A}_0$. For example, if $\max[w_1(\mathcal{A}_1)] - \min[w_1(\mathcal{A}_1)] = 0.8$, then $\max[w_0(\mathcal{A}_0)] - \min[w_0(\mathcal{A}_0)] = 0.08$. Suppose that, when voting in the VMPM, i declares valuations $v_0 : \mathcal{A}_0 \longrightarrow [0, 1]$ and $v_1 : \mathcal{A}_1 \longrightarrow [0, 1]$. Corollary 6 predicts that v_0 and v_1 will be discrepancy-minimizing, which means that $\max[v_1(\mathcal{A}_1)] - \min[v_1(\mathcal{A}_1)] \approx 0.8$ and $\max[v_0(\mathcal{A}_0)] - \min[v_0(\mathcal{A}_0)] \approx 0.08$. Thus, voter i will bid ten times more money on $\mathcal{A}_1$ than on $\mathcal{A}_0$. In other words, she willingly cedes her political influence (in a sense, 'selling her vote') over $\mathcal{A}_0$, so as to preserve her voting money for more important referenda like $\mathcal{A}_1$.

Preferences versus judgements. Jones (1988) compares point voting with several other ways that 'preference intensity' can influence public decisions, such as logrolling, pressure group activity, and electioneering. However, he questions the admissibility of subjective preference intensity into public decisions which also depend upon objective 'judgements' about matters of fact. Jones briefly entertains the possibility of a point voting system where a voter's point allocation expresses not only the intensity of her 'preferences', but also her confidence in her 'judgements'. However, such a system is clearly vulnerable to strategic exaggeration (of both intensity and confidence).

Nevertheless, the VMPM does indirectly measure each voter's 'confidence'. To see this, suppose that the utility function $w_i : \mathcal{A}_0 \longrightarrow [0, 1]$ in eqn.(8) arises as follows. Let $\mathbf{X}$ be set of possible 'states of the world', and suppose that i has a joint utility function $U_i : \mathcal{A}_0 \times \mathbf{X} \longrightarrow [0, 1]$. Let μ be a probability measure on $\mathbf{X}$. We assume:

(a) For all $x \in \mathbf{X}$, $\max_{a \in \mathcal{A}_0} U_i(a, x) \;-\; \min_{a \in \mathcal{A}_0} U_i(a, x) \;\approx 1.$

(b) There is some constant $m \in [0, 1]$ such that, for all $a \in \mathcal{A}_0$,

$$\int_{\mathbf{X}} U_i(a, k)\; d\mu[x] \quad = \quad m.$$

Hypothesis (a) says that, in each specific state of the world, some alternatives are clearly much better for i than others. However, hypothesis (b) says that all alternatives are equally good for i, when we average their payoffs over all possible world states. (In particular, i deems different alternatives to be optimal in different world states.)

The precise state of the world, $x \in \mathbf{X}$, is unknown to i —her incomplete knowledge is represented by a subset $\mathbf{K}_i \subseteq \mathbf{X}$ containing x. Thus, the utility which i assigns to an alternative $a \in \mathcal{A}_0$ is its expected utility:

$$w_i(a) \quad = \quad \int_{\mathbf{K}_i} U_i(a,k)\ d\mu[k].$$

Now, if the set $\mathbf{K}_i$ is quite large in $\mathbf{X}$ (meaning that i has little confidence about her knowledge of x), then hypothesis (b) implies that $w_i(a) \approx m$ for all $a \in \mathcal{A}$. Thus, $\max[w_i(\mathcal{A}_0)] - \min[w_i(\mathcal{A}_0)] \approx m - m = 0$. Thus, if i declares a discrepancy-minimizing valuation $v_i : \mathcal{A}_0 \longrightarrow [0,1]$ (as predicted by Corollary 6), then $\max[v_i(\mathcal{A}_0)] - \min[v_i(\mathcal{A}_0)] \approx 0$, so her vote will have little impact, due to her lack of confidence. On the other hand, if $\mathbf{K}_i$ is very small (meaning that i is highly confident the she has precise knowledge of x), then $w_i(a) \approx U_i(a,x)$ for all $a \in \mathcal{A}$, and then hypothesis (a) implies that $\max[w_i(\mathcal{A}_0)] - \min[w_i(\mathcal{A}_0)] \approx 1$. Thus, $\max[v_i(\mathcal{A}_0)] - \min[v_i(\mathcal{A}_0)] \approx 1$, so her vote will then have greater impact, reflecting her greater confidence.

References

Abeles, F. F. (Ed.), 2002. The Political Pamphlets and Letters of Charles Lutwidge Dodgson and Related Pieces. Lewis Carroll Society of North America.

Allen, G. O., October 1977. Beyond the Voter's Paradox. Ethics 88, 50–61.

Allen, G. O., January 1982. Formal decision theory and majority rule. Ethics 92, 199–206.

Black, D. S., 1958. Theory of committees and elections. Cambridge UP, Cambridge, UK.

Brams, S. J., 2003 [1975]. Game Theory and Politics. Dover, New York.

Brams, S. J., Fishburn, P. C., 1983. Approval voting. Birkhäuser Boston, Mass.

Brams, S. J., Riker, W. H., 1973. The paradox of vote trading. American Political Science Review 66.

Buchanan, J. M., Tullock, G., 1990 [1962]. The Calculus of Consent. Liberty Fund, Indianapolis.

Cao, X., 1982. Preference functions and bargaining solutions. In: Proceeedings of the 21st IEEE Conference on Decision and Control. Vol. 1. pp. 164–171.

Clarke, E., Fall 1971. Multipart pricing of public goods. Public Choice 11, 17–33.

Coleman, J. S., December 1966. The possibility of a social welfare function. American Economic Review 56, 1105–1122.

Coleman, J. S., 1970. Political money. American Political Science Review 64, 1074–1087.

d'Aspremont, C., Gevers, L., 1977. Equity and the informational basis of collective choice. Review of Economic Studies 44, 199–209.

Dhillon, A., 1998. Extended Pareto rules and relative utilitarianism. Soc. Choice Welf. 15 (4), 521–542.

Dhillon, A., Mertens, J.-F., 1999. Relative utilitarianism. Econometrica 67 (3), 471–498.

Dodgson, C. L., 1873. A discussion of the various methods of procedure in conducting elections. Reprinted in Black (1958) or Abeles (2002).

Gibbard, A., 1973. Manipulation of voting schemes. Econometrica 41, 587–601.

Good, I., 1977. Justice in voting by demand-revalation. Public Choice 11 (2), 65–70.

Green, J. R., Laffont, J.-J., 1979. Incentives in Public Decision Making. Vol. 1 of Studies in Public Economics. North-Holland, Amsterdam.

Groves, T., 1973. Incentives in teams. Econometrica 41, 617–631.

Hardin, R., January 1982. Comment on formal decision theory and majority rule. Ethics 92, 207–210.

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

Harsanyi, J., 1955. Cardinal welfare, individualstic ethics and interpersonal comparisons of utiltiy. Journal of Political Economy 63, 309–321.

Hylland, A., Zeckhauser, R., August 1979. A mechanism for selecting public goods when preferences must be elicited. Tech. Rep. KSG Discussion Paper 70D, Harvard University.

Intriligator, M. D., October 1973. A probabilistic model of social choice. Review of Economic Studies 40 (4), 553–560.

Jones, P., March 1988. Intense preferences: Strong beliefs and democratic decision-making. Political Studies XXXVI (1), 7–29.

Karni, E., 1998. Impartiality: definition and representation. Econometrica 66 (6), 1405–1415.

Koblitz, N., 1994. A course in number theory and cryptography, 2nd Edition. Vol. 114 of Graduate Texts in Mathematics. Springer-Verlag, New York.

Laine, C. R., August 1977. Strategy in point voting: A note. Quarterly Journal of Economics 91 (3), 505–507.

Mas-Colell, A., Whinston, M. D., Green, J. R., 1995. Microeconomic Theory. Oxford UP, Oxford, UK.

Maskin, E., 1978. A theorem on utilitarianism. Rev. Econom. Stud. 45 (1), 93–96.

McAfee, R., McMillan, J., 1987. Auctions and bidding. Journal of Economic Literature 25, 699–738.

Milgrom, P. R., 1989. Auction theory. In: Bewley, T. (Ed.), Advances in economic theory: Fifth World Congress. Vol. 12 of Econom. Soc. Monogr. Cambridge Univ. Press, Cambridge, pp. 1–32.

Moulin, H., October 1986. Characterizations of the pivotal mechanism. Journal of Public Economics 31 (1), 53–78.

Moulin, H., 1988. Axioms of cooperative decision making. Cambridge University Press, Cambridge.

Mueller, D. C., December 1967. The possibility of a social welfare function: comment. American Economic Review 57, 1304–1311.

Mueller, D. C., Fall 1971. Fiscal federalism in a constitutional democracy. Public Policy 19 (4), 567–593.

Mueller, D. C., February 1973. Constitutional democracy and social welfare. Quarterly Journal of Economics 87 (1), 60–80.

Mueller, D. C., August 1977. Strategy in point voting: Comment. Quarterly Journal of Economics 91 (3), 509.

Mueller, D. C., 2003. Public Choice III. Cambridge UP, Cambridge.

Mueller, D. C., Philpotts, G. C., Vanek, J., Fall 1972. The social gains from exchanging votes: A simulation approach. Public Choice 13, 55–79.

Musgrave, R. A., 1959. The theory of public finance. McGraw-Hill, New York.

Myerson, R. B., 1981. Utilitarianism, egalitarianism, and the timing effect in social choice problems. Econometrica 49 (4), 883–897.

Ng, Y.-K., October 1975. Bentham or Bergson? Finite sensibility, utility functions, and social welfare functions. Review of Economic Studies 42, 545–569.

Ng, Y.-K., 1985. The utilitarian criterion, finite sensibility, and the weak majority preference principle. A response. Soc. Choice Welf. 2 (1), 37–38.

Ng, Y.-K., 2000. From separability to unweighted sum: a case for utilitarianism. Theory and Decision 49 (4), 299–312.

Nitzan, S., Winter 1975. Social preference ordering in a probabilistic voting model. Public Choice 24 (24), 93–100.

Nitzan, S., 1985. The vulnerability of point-voting schemes to preference variation and strategic manipulation. Public Choice 47 (2), 349–370.

Nitzan, S., Paroush, J., Lampert, S. I., 1980. Preference expression and misrepresentation in points voting schemes. Public Choice 35 (4), 421–436.

Olson, M., 1965. The logic of collective action. Harvard UP, Cambridge, MA.

Philpotts, G. C., 1971. Vote-trading in a competitive model. Ph.D. thesis, Cornell University.

Philpotts, G. C., August 1972. Vote trading, welfare, and uncertainty. Canadian Journal of Economics 5 (3), 358–372.

Riker, W. H., 1982. Liberalism against populism. Waveland Press, Prospect Heights, Illinois.

Roberts, D. J., Postlewaite, A., January 1976. The incentives for price-taking in large exchange economies. Econometrica 44 (1), 115–127.

Satterthwaite, M., 1975. Strategy-proofness and Arrow's condition: Existence and correspondence theorems for voting procedures and social welfare functions. Journal of Economic Theory 10, 187–217.

Schneier, B., 1996. Applied Cryptography, 2nd Edition. John Wiley & Sons, New York.

Schwartz, T., Winter 1975. Vote trading and Pareto efficency. Public Choice 24, 101–109.

Segal, U., 2000. Let's agree that all dictatorships are equally bad. Journal of Political Economy 108 (3), 569–589.

Shapley, L. S., 1953. A value for n-person games. In: Kuhn, H., Tucker, W. (Eds.), Contributions to the theory of games II. Princeton UP, Princeton, NJ, pp. 307–317, (Reprinted on pp. 31-40 of *The Shapley Value*, Cambridge UP, 1988).

Sjostrom, T., July 1991. A new characterization of the Groves-Clarke mechanism. Economics Letters 36 (3), 263–267.

Sobel, J., 2001. Manipulation of preferences and relative utilitarianism. Games Econom. Behav. 37 (1), 196–215.

Stinson, D. R., 2006. Cryptography, 3rd Edition. Discrete Mathematics and its Applications (Boca Raton). Chapman & Hall/CRC, Boca Raton, FL, theory and practice.

Tideman, T. N. (Ed.), Special Supplement to Spring 1977. Public Choice. Vol. 29(2). Center for the Study of Public Choice, Virginia Polytechnic Institute, Blacksburg, Virginia, (Special issue on the Groves-Clarke demand-revealing process).

Tideman, T. N., 1997. Voting and the revelation of preferences for public activities. In: Mueller, D. C. (Ed.), Perspectives on Public Choice: A Handbook. Cambridge UP, New York, Ch. 11, pp. 226–244.

Tideman, T. N., Tullock, G., December 1976. A new and superior method for making social choices. Journal of Political Economy 84 (6), 1145–1159.

Tullock, G., December 1959. Some problems of majority voting. Journal of Political Economy 67, 571–579.

Vickrey, W., 1961. Counterspeculation, auctions and competitive sealed tenders. Journal of Finance 16, 8–37.