Dynamically Consistent Voting Rules *

Madhav Chandrasekher

Current Draft: October 25, 2013

Abstract: This paper studies families of social choice functions (SCF's), i.e. a collection of social choice functions $\{\Phi_A\}$, where the family is indexed by the option set of choices. These (sets of) functions arise in sequential choice problems where at each stage a set of options is given to a population of voters and a choice rule must aggregate stated preferences to generate an aggregate choice. In such settings, the aggregate decision-making process should reflect some form of consistency across choice problems. We characterize the class of (sequences of) SCF's that satisfy two properties: (i) strategy-proofness and (ii) a notion of dynamic consistency inspired by Sen's α from choice theory. When the aggregate choice is anonymous, this class turns out to be exactly the set of q-rules, i.e. rules in which the selected alternative is the most preferred alternative of the voter at the q-th N-tile of the population (where N is the set of voters). This nests median voter schemes when no phantom voters are admitted in the decision rule. Without anonymity we obtain a class that we call "vote-by-committee" rules, the name due to some similarities with a class of SCF's axiomatized in Barberá et al. (1991).

Keywords: Aggregation of Preferences, Dynamic Consistency, Strategy-proofness.

1 Introduction

This paper analyzes strategy-proof voting rules under an additional consistency hypothesis that we call *dynamic consistency*. The motivation stems from situations in which voting procedures can involve multiple stages. Consider a simple example of a firm trying to fill a vacancy. The committee might vote an offer on a job candidate. If the candidate accepts, then the position is filled. If, on the other hand, he/sherejects then a subsequent offer is made to another candidate, who then makes an accept/reject decision. If the second candidate rejects, the committee convenes once more to vote on a subsequent offer, and so on the process continues until the position is filled. A sequential voting rule is a function that – at each stage – maps announced preferences over a set candidates to a selected candidate. What makes this procedure dynamic is that the pool of candidates in each round is (possibly) shrinking. Hence, announced preferences across rounds need not be consistent with one another, i.e. they can

*Acknowledgments: This is a draft of work still in progress. Comments are welcome. Mailing address: Department of Economics, W.P. Carey School of Business, P.O. Box 873806, Tempe, AZ 85287-3806; E-mail: madhav.chandrasekher@asu.edu

exhibit reversals of orderings between candidates. We will be interested in voting rules which induce truthful reporting of preferences at each stage. Under such rules, the preference announcement phase of the aggregation process is effectively static: it is as if we plug in voter preferences over the grand list of candidates at the outset and run the social choice rule as it were an automaton. If the round t candidate rejects an offer, the round $t+1$ social choice rule simply aggregates the restricted preferences over the remaining candidates in the pool. This process continues until either a selected candidate accepts an offer or there are no more candidates in the pool.

In addition to the property that the aggregation procedure elicits truthful reports of preferences, we will also require that the procedure be dynamically consistent. It will take some work to be more precise about this, but – in a nutshell – when we say a sequence of social choice rules (i.e. a social choice "procedure") is dynamically consistent we mean that the criteria used to make a choice in each stage are consistent with one another. There isn't a canonical measure of dynamic consistency in our setting. We propose two increasingly restrictive criteria and characterize choice procedures which (resp.) satisfy these criteria.[1] To describe these we introduce a little formalism. For each set of candidates A, let $\succeq_A$ denote a profile of preferences over candidates in the pool A. Let $\mathcal{M}$ denote the set of all subsets of the grand pool of candidates, X, and for each $A \in \mathcal{M}$ consider an aggregate choice function which takes as input reported preferences over options in A and outputs an outcome in A, i.e. putting $\mathcal{R}_A^N$ to be the set of preference profiles supported on A we have $\Phi_A : \mathcal{R}_A^N \to A$. In this setting a dynamic choice rule is a family of maps, $\{\Phi_A\}_{A\in\mathcal{M}}$. This is the primitive of our analysis. The dynamics of the aggregate choice process are induced by an (exogenous) background transition process which determines the period $t+1$ distribution of option pools given the period t option pool and aggregate choice. However, the details of the this process only affect how we (directly) implement a given family $\{\Phi_A\}$. It does not factor into the normative analysis of these families, which is the focus of this paper.

Since there is – in principle – a different voting rule applied to each pool of candidates we will want the family of voting rules to exhibit (dynamic) consistency across choice problems that are considered in sequence. More precisely, when considering two option pools $A \subseteq B$ where A is considered further down the process than B, the criteria used to select an aggregate choice from B are "similar" to those used to select an aggregate choice from A. Consider the fol-

[1]In the current draft, we only provide a characterization for choice rules which satisfy one of these criteria (viz. sequential consistency, see below). The characterization result for choice rules which satisfy the other criterion is work in progress.

lowing two notions of dynamic consistency: (let $\mathcal{R}_A$ denote the set of preferences supported on the set A)

1. Say that $\{\Phi_A\}$ is *rational* if

$$\forall \{\succeq_A\} \text{ s.t. } (\succeq_B)|_A = \succeq_A \ \exists \succeq \in X \times X \text{ s.t. } \Phi_A(\succeq_A) \in \text{arg max}_{x \in A} \succeq .$$

2. For each $\succeq_A \in \mathcal{R}_A$ consider $\succeq_B$ with $(\succeq_B)|_A = \succeq_A$. Say that $\{\Phi_A\}$ is *sequentially consistent* if $\Phi_A(\succeq_A) = \Phi_B(\succeq_B)$ whenever arg $\max_{x \in B} \succeq_i \in A, \forall i$.

The first restriction says that for any system $\{\succeq_A\}$ that is consistent, i.e. $(\succeq_B)|_A = \succeq_A$ (so that every voter's preferences are consistent), the social choice rule is induced by maximization of a single preference relation $\succeq \subseteq X \times X$. If the choice rule is strategy-proof then this restriction is tantamount to the condition that the map $\Phi_X : \mathcal{R}_X \to X$ is strategy-proof and rational, i.e. it maximizes a fixed preference on X. The second restriction relaxes rationality in a particular way, which is particularly easy to interpret when Φ_A are choice functions. In this case, it says that when the pool is B and the elicited preferences are such that all the most preferred candidates from B are still available in A, then the choice rule Φ_B selects the same candidate from A were the pool of candidates to be A. When it is a choice correspondence it says that when selected candidates (from the larger pool B) are still available in the smaller pool, then they are still selected – assuming population preferences over these candidates are the same. Sequential consistency clearly has a flavor of Arrow's well-known IIA (independence of irrelevant alternatives) axiom. The appearance is somewhat superficial however, as it is neither the case that sequential consistency implies IIA nor that IIA implies sequential consistency.

The main results of the paper characterize procedures $\{\Phi_A\}$ that are strategy-proof and dynamically consistent, viz. under sequential consistency. It will turn out that a rule which is strategy-proof and sequentially consistent must be a q-rule, viz. where the selected outcome lies at the q-th N-tile of the distribution of voter ideal points. These rules are also rational in the sense described above. Hence, while the opposite might seem true at first blush, rationality (coupled with strategy-proofness) is stronger than sequential consistency. We have, ex post, chosen the labels of the restrictions to reflect this fact.[2] From a positive viewpoint, the class of q-rules are important. For example, majority rule and super-majority rule are commonly used aggregation procedures

[2]This discussion states our characterization under the hypothesis of anonymity. When the family of SCF's is not anonymous, we get a generalization of q-rules that aggregates to a final choice in two steps: First, break up voters into committees and use a q rule to aggregate the recommended choice across committees (where different committees can have different q's). Second, aggregate recommendations across committees. Thus, removing anonymity introduces two forms of aggregation, (i) inter-committee aggregation and (ii) intra-committee aggregation.

which can both be represented as q-rules. What properties of these voting rules accounts for their use in aggregation problems? Moreover, are these the only rules which exhibit these properties? Our results suggest the following answer: majority, super-majority (and q-rules, more generally) are the *only* anonymous rules which are both non-manipulable and dynamically consistent (in the sense of sequential consistency).

Let us conclude by describing some limitations of our results. We do not explore whether we can extend our characterization to a larger domain of preferences. In other words, we do not investigate whether the single-peaked domain is a maximal domain of preferences for which our characterization holds. We also pass to the domain of strict preferences since this leads to clarity in the results and arguments. The case of ties shouldn't (we believe) present too many difficulties but we do not seek out such an extension here. In the next section, we (re)introduce the dynamic preference aggregation problem. We then state and prove our two characterization results, then compare our model to the only other (to our knowledge) characterization of strategy-proof committee-voting in the social choice literature. Section 4 concludes and the appendix collects proofs omitted from the main text.

1.1 Related Literature

This paper is most closely related to the positive political theory/social choice literature that seeks to characterize incentive compatible, i.e. strategy-proof, aggregate choice rules (voting rules). Several papers in this literature, see Barberá (2001) for a comprehensive survey, shows that when we interpret the voting rule as influenced by "phantom" voters, i.e. voters who are present in the population but who need not report to the mechanism, then the *only* incentive compatible voting rule is majoritarian (median) voting.[3] We add to this literature by seeking a normative foundation for rules that are incentive compatible and, additionally, internally consistent when the voting process involves several stages. When we add anonymity to our list of axioms we obtain q-rules, i.e. the aggregate choice is such that – under some ordering – at least a fraction $1-q$ of the population possesses an ideal point that is to the right of this choice. Chambers (2007) provides a characterization of q-rules when voter preferences are subject to uncertainty. Fix a distribution over profiles of voter preferences and an ex post aggregate decision rule f. Now consider the ex ante decision rule which selects a choice which lies at the qth-percentile of the distribution

[3]This result also requires other assumptions, e.g. oddness of the population of voters who report preferences and anonymity of the voting rule. We omit these extra hypotheses from this more coarse discussion.

(induced by the ex post choice function f) – this is the class of rules characterized in Chambers (2007). When we relax anonymity, we obtain a generalization of q-rules that we call "vote-by-committee" rules. These can be thought of as an analogue, for the single-peaked domain, of a class of rules introduced and axiomatized in Barberá et al. (1991). We will say more about committee voting rules and the connection with this paper in a subsequent section.

2 Model

The domain of the social choice problem is described as follows. Let X denote a finite subset of $\mathbf{R}$, say $X = \{x_1, \ldots, x_n\}$ with $x_i < x_{i+1}$. Fix a finite set of voters N and let $\mathcal{R}$ denote the space of profiles of *single-peaked* orders. Note that an order is a complete and transitive binary relation on X and a profile of orders is a vector, $\succeq := (\succeq_1, \ldots, \succeq_N)$, of orders – one for each of the voters in the population. A given order $\succeq_i$ is single-peaked on X if there is some $x^* = x_k \in X$ such that $x_1 \prec_i x_2 \prec_i x_k = x^*, x^* \succ_i x_{k+1} \succ_i x_{k+2} \succ_i \cdots \succ_i x_n$. That is, placing elements of X left-to-right along the number line (w.r.t. the natural ordering on $\mathbf{R}$) the preference $\succeq_i$ has a unique maximum at x^*, so that the utility function is hump-shaped on X with a peak at x^*. Henceforth, when we speak of preferences we will *always* have single-peaked preferences in mind – so that we suppress mention of this implicit requirement.

For each subset $A \subseteq X$ let $\mathcal{R}_A$ denote profiles of orders on the subset A. Here and elsewhere, $\hat{\succeq}_i$ will denote voter i's reported preference and $\succeq_i$ will denote his actual preference. Profiles of announcements will be similarly denoted $\hat{\succeq}$ and we take $\hat{\mathcal{R}}$ (resp. $\hat{\mathcal{R}}_A$) to be the space of reports. The primitive of the analysis is a sequence of social choice functions (SCF's), $\{\Phi_A\}$, where $\Phi_A : \mathcal{R}_A \rightrightarrows A$. The reason we refer to the collections $\{\Phi_A\}$ as a "sequence" is that the cross-sectional restrictions imposed on the family $\{\Phi_A\}$ only apply across pairs (Φ_A, Φ_B) where $A \subseteq B$. To fix ideas, a story to have in mind is one of options being considered in sequence. For instance, if an aggregate selection, x_B, from B is no longer feasible, then the subsequent decision problem is to aggregate preferences over $B \backslash x_B$ and make a choice from this set, and so on until we reach the leftover pool of options comprising menu A.

We impose three normative criteria on sequences $\{\Phi_A\}$ – (i) strategy-proofness, (ii) sequential consistency, and (iii) anonymity, and characterize voting rules which satisfy these three criteria. We then characterize rules which satisfy the first two criteria but are possibly non-anonymous. Moreover, we (intend to) examine two incarnations of dynamic consistency, resp. sequential consistency and a second that we dub "rationality". Rationality will require that aggregate choices are derived from maximization of some aggregate preference. The class of rules that are strategy-proof and satisfy sequential consistency will turn out to be rational, so that rationality is a more permissive criterion when restricting to the class of strategy-proof families $\{\Phi_A\}$. We do not yet have an answer for how much more permissive it is, viz. a characterization of families that are strategy-proof and rational.

Definition 1. A sequence $\{\Phi_A\}$ is *strategy-proof* if for each Φ_A we have $a \succeq_i b, \forall a \in \Phi_A(\succeq_i, \succeq_{-i}), \forall b \in \Phi_A(\hat{\succeq}_i, \succeq_{-i})$.

Let us make two comments on this definition. First, a minor comment on correspondences vs. functions. While we will restrict to choice functions in our results, the definition of strategy-proofness for the family $\{\Phi_A\}$ does not necessitate that the individual maps Φ_A be functions. Note, however, that incentive compatibility is imposed across all selections from (resp.) $\Phi(\succeq_i, \succeq_{-i})$ and $\Phi(\hat{\succeq}_i, \succeq_{-i})$. This is evidently necessary if we want to allow the flexibility for Φ_A to be a strategy-proof choice correspondence as opposed to a function. A second comment regards the very definition of strategy-proofness in dynamic choice problems, where preferences are reported at each stage in which an aggregate choice is made. In this case, the definition that we have dubbed "strategy-proofness" is just an *ad hoc* criterion. It does not, by itself, imply that period-by-period truth telling is an SPNE in the extensive-form induced by the SCF's $\{\Phi_A\}$ and the ambient transition process which determines the current period's option set.

Formally define a transition process $p(\cdot, \cdot)$ as follows.[4] For each pair (A, a) where $A \subseteq X$ and $a \in A$ we have a distribution $p(\cdot|(A, a)) \in \Delta(\mathcal{M})$. This is the distribution on the next period's pool of options conditional on the current pool being A and the social choice in the period being $a \in A$. Fixing a strategy-proof family $\{\Phi_A\}$ and transition process $\{p(\cdot|(A, a))\}$, we complete the description of the extensive-form with a transfer function which determines payoffs at terminal nodes. To define payoffs in the game, we need to select cardinal representations of voter preferences (even though they only report ordinal preferences). Let $\succeq \in \mathcal{R}$ denote a preference profile and let $u_{\succeq}$ denote an associated cardinal representation.[5] Define a transfer function (for each reporting stage) $r_i^t(\hat{\succeq}_i, \hat{\succeq}_{-i})$, where the superscript denotes the stage t.

[4]Note: In our notation, we are taking the process $p(\cdot, \cdot)$ to be time-homogenous. This is just for brevity of notation. The time (non)homogeneity of the transition process does not factor whatsoever into the arguments, hence is wlog.

[5]Any $u_{\succeq}$ will work, the implementability of period-by-period truth-telling in SPNE is not dependent on a particular selection of $u_{\succeq}$.

Fixing a triple $(\{\Phi_A\}, p(\cdot,\cdot), r_i^t(\cdot,\cdot))$, the extensive-form game is defined in stages, where in each stage t there is a pool of options $A_t \subseteq X$. Voters then report preferences over these options to the (direct) mechanism, these preferences are aggregated and an aggregate choice is determined according to Φ_A, with transfers given out to each voter i according to r_i^t. With some (exogenous) probability this aggregate choice is implemented (e.g. the candidate accepts the job offer), in which case we reach a terminal node of the game tree. With complementary probability the aggregate choice is not implementable, in which case the game proceeds to another round of voting. Note that while the story we have in mind is that of a candidate $a \in A_t$ declining the offer and hence being removed from the pool, there is no such restriction on the transition probability $p(\cdot|(A,a))$. More precisely, the probability that the pool in round $t+1$ is B conditional on the round t pool being A and the offer having gone to $a \in A$ is given by $p(B|(A,a))$ (put $B = \emptyset$ if the game ends after offer a is made). We do not assume that $p(B|(A,a)) = 0$ if $a \in B$, although this is the leading example of the story we have in mind. This process of voting, selection, and subsequent implementation/rejection of the aggregate choice continues on for a maximum of T periods, at which point the process terminates. Denote the extensive-form game induced by the triple $(\{\Phi_A\}, p(\cdot,\cdot), r_i^t(\cdot,\cdot))$ as $\mathcal{G}(\{\Phi_A\}, p(\cdot,\cdot), r_i^t(\cdot,\cdot))$.

Lemma 1. *Fix a transition process $p(\cdot,\cdot)$ and a family of strategy-proof SCF's $\{\Phi_A\}$. Then, for each voter i we can find transfers $\{r_i^t(\cdot,\cdot)\}_{t=1}^T$ such that truthful reporting of preferences in each stage is an SPNE of the game $\mathcal{G}(\{\Phi_A\}, p(\cdot,\cdot), r_i^t(\cdot,\cdot))$.*

Via the mechanism constructed in the lemma, we can think of the period-by-period strategy-proof restriction on $\{\Phi_A\}$ as an equilibrium condition. Put another way, there are mechanisms such that static non-manipulability implies extensive-form non-manipulability. The main step in the proof of the lemma is to construct transfers that induce period-by-period truth-telling. Note that there is a trade-off involved in the decision to truthfully report in any given period. On the one-hand, truth telling is a static best-response on account of the strategy-proofness hypothesis. On the other hand, the current period's report changes the outcome via the SCF Φ_A and this, in turn, changes the next period's distribution over continuation utilities (i.e. $p(\cdot|(A,a))$, where $a = \Phi_A(\succeq))$. This means that misreporting preferences, while leading to a current utility drop, might induce a more favorable distribution of option pools in the next stage should the current offer be declined. Hence, static strategy-proofness does not, by itself, imply that truth-telling in every stage is an SPNE in the induced extensive-form game.

To show that it is an SPNE in the extensive form game we introduce transfers in the reporting mechanism. If the outcome functions, Φ_A, at each stage is strategy-proof (in the sense of Definition 1), then in the terminal stage of the game truth telling is a best-response regardless of what path has led to that node of the game tree. Now work backwards, assuming that from round $t+1$ onwards truth-telling is an SPNE of the continuation game. To make this precise, we assume players evaluate terminal payoffs by looking forward from any round t and compute the distribution over these terminal nodes taking continuation strategies and transition probabilities into account (and evaluate using expected utility). For this we fix *any* cardinal representation of a voter's true preference $\succeq$, call it $u(\cdot)$. His payoff from announcing $\hat{\succeq}$ in round t is the following $p(\cdot|(A,a))$-weighted sum: (let $p(\emptyset|(A,a))$ denote the acceptance probability)

$$p(\emptyset|(A,a))u(a) + \sum_{B} p(B|(A,a)) \cdot U_{t+1}(\succeq |B),$$

where $U_{t+1}(\succeq |B)$ is the continuation utility when voters truthfully report preferences in the continuation game where the next stage pool is B. By backwards induction, the continuation utility $U_{t+1}(\hat{\succeq}|B)$ is maximized when $\hat{\succeq} = \succeq$. Note that the report affects current period utility $u(a)$ directly and it affects continuation payoffs indirectly by shifting the distribution $p(\cdot|(A,a))$. We construct transfers so that net utility is maximized when the best-response in the current period, given that the best-response in the continuation game is to truth-tell, is to set $\hat{\succeq} = \succeq$. This is relatively straightforward (compared to other dynamic mechanism constructions) since we are gifted with a function, viz. Φ_A, that implements the stage outcome in dominant strategies were the game to be one-shot. Transfers just adjust (on path) continuation utility so that the stage N.E. remains a best-response once the voter is compensated in present payoffs for any temptation to deviate from truth-telling induced by continuation payoffs.[6]

Definition 2. A sequence of profiles $\{\succeq_A\}$, where each $\succeq_A \in \mathcal{R}_A$, is *constant* if whenever $A \subseteq B$ we have $(\succeq_B)|_A = \succeq_A$.

That is, a constant family of profiles is just induced (via restricting the orders of each voter) by a single profile of orders, $\succeq \in \mathcal{R}_X$, on the full pool of candidates X. Rationality requires that, across a constant family of announced profiles $\{\succeq_A \}$, the SCF is induced by maximization of a single order $\succeq_\Phi \subseteq X \times X$, i.e. there is a representative agent for the SCF.

[6]Two putative extensions of the lemma would be to (i) make the transfers budget-balanced and/or (ii) seek a mechanism that strongly implements the SCF family $\{\Phi_A\}$. The former might be possible, but a well-known impossibility result in static mechanism design says that ex post equilibrium and budget balance are incompatible. Hence, we might be able to adjust transfers to obtain time 0 (ex ante) budget balance, even though transfers don't cancel out along every branch of the game tree. Regarding the implementation question (i.e. strong implementability of the SCF family $\{\Phi_A\}$), this is quite possibly doable. It would also, to the best of our knowledge be an interesting contribution to the implementation literature since, in general, there is not much work on (exact) dynamic implementation.

Definition 3. A sequence of SCF's $\{\Phi_A\}$ is *rational* if for any constant family $\{\succeq_A\}$ there is a single order $\succeq^*$ such that $\Phi_A(\succeq_A) = \arg\max_{x \in A} (\succeq^*)|_A$.

Finally, we consider an alternative definition of dynamic consistency. In conjunction with strategy-proofness this turns out to be a formal weakening of rationality. However, while this might be intuitively clear, our argument is indirect since it passes through the representation theorem. For a profile of preferences $\succeq_B$ over candidates in pool B, let $(a_1(\succeq_B), \ldots, a_N(\succeq_B))$ denote the profile of peaks of (resp.) each individual preference $\succeq_i$ comprising the profile.

Definition 4. A sequence of SCF's $\{\Phi_A\}$ is *sequentially consistent* if for any pair (A, B) where $A \subseteq B$ we have $(a_1(\succeq_B), \ldots, a_N(\succeq_B)) \subseteq A$, then $\Phi(\succeq_B) = \Phi(\succeq_A)$.

To motivate this condition, recall Sen's α from choice theory: (here $C(\cdot)$ denotes a choice correspondence)

$$C(B) \cap A \subseteq C(A) \text{ if } A \subseteq B.$$

The idea behind this condition is that if choices are derived from some form of preference maximization, then – according to this preference – elements of $C(B)$ beat, via head-to-head comparison, elements of $B \backslash C(B)$, and hence must also beat elements of A. Thus, if the criteria, viz. from an aggregate choice perspective we represent criteria with preferences, being used to aggregate a choice when the option pool is A are similar to the ones being used to make a choice when the set is B, then elements of $C(B)$ should still be selected from A when $A \subseteq B$.

The sequential consistency hypothesis says that when all the ideal points of voters are present in A, then the choices agree. This hypothesis is clearly related to Sen's α, but there isn't a set-theoretic relationship in either direction. There is a sense in which it is a strengthening of Sen's α. To see this replace $C(B)$ with $\Phi_B(\succeq)$ and note that, since $\{\Phi_A\}$ is a family of SCF's, applying Sen's α directly would yield,

$$\Phi_B(\succeq) \cap A \subseteq \Phi_A(\succeq)$$

provided that the intersection on the LHS is non-empty. Note that, since Φ_A is an SCF there is no distinction here between set containment and equality. Hence, Sen's α applied to the aggregate choice correspondence $\{\Phi_A\}$ would say that $\Phi_A(\succeq) = \Phi_B(\succeq)$ so long as $\Phi_B(\succeq) \in A$. In contrast, sequential consistency says that if the ancillary condition that all the ideal points from B are present in A, then it must be the case that $\Phi_B(\succeq) = \Phi_A(\succeq)$ (in particular, $\Phi_B(\succeq) \in A$). In this sense, sequential consistency strengthens Sen's α.

There is also a sense in which it weakens Sen's α. If $a_i(\succeq) = \Phi_B(\succeq)$ for some voter i, then the hypothesis that $(a_1(\succeq_B), \ldots, a_N(\succeq_B)) \subseteq A$ implies that $\Phi_B(\succeq) \cap A \subseteq \Phi_A(\succeq)$. Hence, when the choice from B is voter i's ideal point, we require that the aggregate choice from A agrees with the aggregate choice from B only when *both* the choice from B is present in A and all ideal points (relative to B) are present in the contracted set A. Hence, if we delete some of these points and still keep $\Phi_B(\succeq_B)$ in A, it may nevertheless be the case that the aggregate choice selects something other than $\Phi_B(\succeq_B)$. Compliance with Sen's α would, on the other hand, require that $\Phi_A(\succeq_A) = \Phi_B(\succeq_B)$. There is another possible way to interpret a family of SCF's as "dynamically consistent" that is a weakening of both sequential consistency and of Sen's α.

Definition 5. A sequence of SCF's $\{\Phi_A\}$ is *weakly sequentially consistent* if for any pair (A, B) where $A \subseteq B$ we have $(a_1(\succeq_B), \ldots, a_N(\succeq_B)) \subseteq A$ and $\Phi_B(\succeq_B) \in A$, then $\Phi(\succeq_B) = \Phi(\succeq_A)$.

In work in progress, we intend to characterize SCF families $\{\Phi_A\}$ which satisfy strategy-proofness, weak sequential consistency, and anonymity. To sum on the preceding discussion, sequential consistency is inspired by but formally distinct from Sen's α from choice theory. Our main task is to characterize SCF's which are strategy-proof and alternately satisfy rationality vs. sequential rationality. To this end, we introduce some classes of SCF's $\{\Phi_A\}$.

Definition 6. An SCF $\Phi : \mathcal{R}_X \rightrightarrows X$ is a *vote by committee* rule if there is (i) a collection of committees $\Pi := \{\pi_i\}_{i=1}^k$, where $\pi_i \subseteq N$ and $\cup_i \pi_i = N$, and (ii) a collection of "quantiles" $(q_1, \ldots, q_k)$ such that

$$\Phi(\succeq) = \text{arg max}_{\pi \in \Pi}\, q - \text{min}_{j \in \pi}\, a_j(\succeq), \forall \succeq$$

where $a_j(\succeq)$ denotes the peak of agent j.

The committees are allowed to be non-disjoint, so that one voter can – in principle – serve on more than one committee. If we require committees to be partitions then there is an additional requirement on the family of SCF's $\{\Phi_A\}$. The notation $q - \text{min}\{a_1, \ldots, a_l\}$ denotes the peak of the voter in cell π whose bliss point is at the q-th l-tile (where l is the size of the committee), measured w.r.t. the other voters in the same committee. The interpretation of the voting rule is as follows. Note that the outcome space is uni-dimensional, i.e. given a menu of actions A (where we label $\{a_1 < a_2 < \cdots < a_k\} := A$), the aggregate choice rule selects an action "level" – e. g. to fix ideas we will think of the action as a level of military intervention. The aggregate choice, when determined by vote-by-committee, selects a level of intervention using the following "hawks among doves" procedure: First, break up the voters into groups (committees).

Second, within each committee consider the recommended action by the "q-dove", i.e. the voter whose most preferred intervention level is at most q levels higher than the committee member who least prefers intervention. Third, consider only the ideal points of the set of q-doves, taken across all committees. Implement the highest recommended intervention level from this set, i.e. choose the level of intervention preferred by the most hawkish of the q-doves.

There is axiomatic work on voting by committees in the social choice literature, viz. Barberá et al. (1991). Our "vote-by-committee" rules can be thought of as an analogue, for the single-peaked domain, of the Barberá et al. (1991) committee voting rules. However, an important formal difference between the model in Barberá et al. (1991) and ours is in the domain of preference aggregation. In Barberá et al. (1991) the domain of aggregation is the set of profiles of *menu preferences*, i.e. preferences over sets of candidates and the outcome space of the SCF is the set of menus, e.g. subsets of candidates. With this as primitive, an SCF is a vote-by-committee rule in the sense of Barberá et al. (1991) if, there is – for each candidate – a collection of minimal decisive coalitions such that, fixing the population profile (of menu preferences), each winning candidate is an element of the top ranked menu for all voters in some decisive coalition. By contrast, the primitive in this paper is a family of SCF's $\{\Phi_A\}$. There is a single family of committees $\{\pi_i\}_{i=1}^k$ attached to the aggregate choice function. Hence, both how we model aggregate choices and the form in which these choices are generated by some form of aggregation across committees is quite different. There are, nevertheless, some interesting similarities between the concepts in Barberá et al. (1991) and this paper. We will revisit their model in a subsequent section, after we present our characterization results.

3 Main Results

We first characterize rules that satisfy the main restrictions plus anonymity, then extend the result by relaxing anonymity.[7]

Theorem 1. *An SCF family $\{\Phi_A\}$ is strategy-proof, sequentially consistent, and anonymous if and only if it is a q-rule for some quantile $q \in [0, 1]$.*

A q-rule is a special case of a vote-by-committee rule when there is a single committee comprised of all voters. Taking N to be the size of the voting popula-

[7]On account of its ubiquity, we haven't set aside a formal definition of anonymity: For any domain of preference profiles, $\mathcal{R}^N$, let Π denote a permutation of the identities of the N voters and let $\succeq_\Pi := (\succeq_{\Pi(1)}, \succeq_{\Pi(2)}, \ldots, \succeq_{\Pi(N)})$. Anonymity is the requirement that $\Phi(\succeq) = \Phi(\succeq_\Pi)$, i.e. the aggregate choice depends on the set of preferences comprising the population, but not on the identity of the voter who possesses any given preference in this set.

tion the set of q-rules breaks up into equivalence classes as follows. Set $q^i_\Phi := \frac{i}{N}$ and partition $[0,1] = \sqcup_i [q^i_\Phi, q^i_\Phi + \frac{1}{N})$. For any preference profile $\succeq$ and collection of peaks $\{a_1(\succeq), a_2(\succeq), \ldots, a_N(\succeq)\}$ we let $a_{q^i_\Phi}(\succeq)$ denote the q^i_Φ-th N-tile w.r.t. the set $\{a_1(\succeq), \ldots, a_N(\succeq)\}$. Any $q \in [q^i_\Phi, q^i_\Phi + \frac{1}{N})$ picks out the same choice, $a_{q^i_\Phi}(\succeq)$. Note that when $q = 1/2$ we recover the median rule (when N is odd). The difference between this and the generalized median is that there are no phantom voters invoked in the aggregation rule. This is an obvious implication of the sequential rationality hypothesis. The (perhaps) surprising part of the result is that this property in conjunction with strategy-proofness provides a distinct "parameterization" of the median voter class, as one among a class of q-rules. Since all q-rules are rational we obtain that strategy-proofness, sequential rationality, and anonymity are together more restrictive than strategy-proofness, rationality, and anonymity. Moreover, note that by Moulin (1980) we know that strategy-proofness and anonymity imply that, for each menu A, we have (for odd integers n)

$$(*)\ \Phi_A(\succeq_A) = \text{med}\{\alpha^1_A, \ldots, \alpha^k_A, a_1(\succeq_A), \ldots, a_n(\succeq_A)\}$$

where the α^i_A denote the ideal points of "phantom" voters, i.e. voters whose preferences influence the outcome, even though their reports are not explicitly an argument in the mechanism.

Note that we have a (potentially) distinct set of phantom voters at each choice set and, absent any restriction tying together the preferences of phantom voters across choice sets, the set of phantom voters in – say – menu B need not bear any relation to the set of phantom voters in menu A (where $A \subseteq B$). Moreover, the ideal points of phantoms can coincide with those of the voters in the population. Since the class $(*)$ nests q-rules, in order to express a $q \neq 1/2$ rule as a generalized median the phantoms in different menus are chosen to coincide with voter ideal points, with multiplicities selected such that the qth N-tile among voter ideals is exactly the median with an appropriate choice of phantoms. Finally, note that not all generalized median voter rules satisfy rationality since we can choose phantoms arbitrarily as we pass from menu to menu. It would be interesting to see (we have not done so here) what class of rules one obtains (for odd integers n) when we impose strategy-proofness, anonymity, and rationality (as opposed to sequential rationality) on the family of SCF's $\{\Phi_A\}$.

Proof of Theorem 1. Necessity of the axioms is immediate, hence we just present the sufficiency argument.[8] We make some initial reductions. Note that, via standard arguments, strategyproofness implies that the SCF only depends on the reported peaks, and not on the underlying orders. Second, sequential consistency

[8]Our argument uses a technique, viz. constructing sequences of "perturbed" population profiles, which we learnt from Schmeidler and Sonnenschein (1978).

implies that *for any menu A* we must have $\Phi_A(\succeq_A) \in \{a_1(\succeq_A), \ldots, a_N(\succeq_A)\}$. That is, the SCF always chooses one of the voter peaks. Next, we call a voter i pivotal at $\succeq$ if (i) $\Phi(\succeq) = a_i(\succeq)$, (ii) if we have some $\succeq'$ such that $\succeq_{-i} = \succeq'_{-i}$ and $a_i(\succeq') \neq a_i(\succeq)$, and (iii) $\Phi(\succeq) \neq \Phi(\succeq')$. That is, at $\succeq$ the selected candidate is voter i's top choice. Moreover, if population preferences are such that only voter i's preference changes, then the SCF changes as well. Proceed in the following steps.

Step 1: Reduction to two peaks.
Take any profile $\succeq$ and let $\Phi(\succeq) \in \{a_1(\succeq), \ldots, a_N(\succeq)\}$. Consider two groups, $\Sigma_1(\succeq) := \{i \in N : a_i(\succeq) \leq \Phi(\succeq)\}, \Sigma_2(\succeq) := \{i \in N : a_i(\succeq) > \Phi(\succeq)\}$. Now consider $\succeq'$ where we just change the preference of some agent i and $\succeq_{-i} = \succeq'_{-i}$. Observe that, due to strategy-proofness, we must have

$$(*)\ \Phi(\succeq') \geq \Phi(\succeq) \text{ if } a_i(\succeq') > a_i(\succeq), i \in \Sigma_1(\succeq)$$

and similarly

$$(**)\ \Phi(\succeq) \leq \Phi(\succeq') \text{ if } a_i(\succeq) > a_i(\succeq'), i \in \Sigma_2(\succeq).$$

Fix a profile $\succeq$ and let i^* be such that $\Phi(\succeq) = a_{i^*}(\succeq)$. Now amend this profile as follows. Let $a(2) := \min_{i \in \Sigma_1(\succeq)} a_i(\succeq)$ and for each $i \in \Sigma_2(\succeq)$ let $\succeq_i$ denote any single-peaked order on X with maximum at $a(2)$. Similarly, for any $i \in \Sigma_1(\succeq)$ let $\succeq_i$ denote a single-peaked order on X with maximum at $a_{i^*}(\succeq)$. Let $\succeq'$ be the associated population profile. We claim that $\Phi(\succeq') = \Phi(\succeq)$. Show this by slowly perturbing $\succeq$ into $\succeq'$. Consider the intermediate profile $\succeq''$ where for all $i \in \Sigma_1(\succeq)$ we have $\succeq_i = \succeq''_i$ and $\succeq_i$ has peak $a(2)$ for all $i \in \Sigma_2(\succeq)$. Strategyproofness implies that $\Phi(\succeq) = \Phi(\succeq'')$. Hence, we reduce to consider $\succeq'$ and $\succeq''$.

We claim that $\Phi(\succeq') = \Phi(\succeq'')$. Find any sequence of profiles $\succeq(0) := \succeq''$ $, \ldots, \succeq(n) := \succeq'$, where along each step of the sequence we change just one voter's preference to an order with peak at $a_{i^*}(\succeq)$. We claim that $\Phi(\succeq'') = \Phi(\succeq(1))$. By $(*)$ we know that $\Phi(\succeq(1)) \geq a_{i^*}(\succeq)$. If $\Phi(\succeq(1)) > a_{i^*}(\succeq)$, then take the agent i such that $\succeq_i(1) \neq \succeq_i(0)$ and note that if $\Phi(\succeq) > a_{i^*}(\succeq)$ he will misreport his preference as $\succeq_i(0)$, which contradicts strategy-proofness. Hence, $\Phi(\succeq(0)) = \Phi(\succeq(1))$, and by induction $\Phi(\succeq) = \Phi(\succeq(0)) = \cdots = \Phi(\succeq(n)) = \Phi(\succeq')$.

Step 2: Characterization on two-peaked profiles.
We verify that the rule Φ restricts to a q-rule on the set of population profiles that have two peaks, i.e. the set $\{a_1(\succeq), \ldots, a_N(\succeq)\}$ has cardinality 2. Let $\{a(\succeq), b(\succeq)\}$ denote the two peaks with $a(\succeq) < b(\succeq)$. Now put $a := \min\{a_i : a_i \in X\}, b := \max\{a_i : a_i \in X\}$. Consider a new profile $\succeq'$, constructed as follows. For every voter i such that $a_i(\succeq) = a(\succeq)$ we replace $\succeq_i$ with the (unique) single-peaked order $\succeq'_i$ with peak a. Similarly, for every voter i with $a_i(\succeq) = b(\succeq)$ replace $\succeq_i$ with

the unique single-peaked order with peak at b. Call the resulting profile $\succeq'$. We claim that $\Phi(\succeq') = b$ (resp. a) if and only if $\Phi(\succeq) = b(\succeq)$. This follows from a near mirror argument as in the previous step. Perturb $\succeq$ to $\succeq'$ sequentially, putting $\succeq(0) := \succeq, \ldots, \succeq(n) := \succeq'$. Here we choose a specific sequence of profiles $\succeq(i)$ as follows. Let $k(\succeq) := |\{i \in N : a_i(\succeq) = b(\succeq)\}|$ and label voters so that those with peak at $b(\succeq)$ are numbered from 1 to $k(\succeq)$. For $1 \leq i \leq k(\succeq)$ let $\succeq(i)$ perturb $\succeq(i-1)$ by keeping all preferences $\succeq_i$ for $i > k(\succeq)$ unchanged. At step $i \leq k(\succeq)$ perturb $\succeq_i$ to $\succeq_i'$, where $\succeq_i'$ is the unique single-peaked order with peak at b. This defines the sequence of amended profiles, $\{\succeq(i)\}_{i=0}^{k(\succeq)}$. Now for $k(\succeq) + 1 \leq i \leq N$ define $\succeq(i)$ as follows. For each $i \leq k(\succeq)$ keep $\succeq_i$ the same as in $\succeq(i-1)$. For each $i > k(\succeq)$ sequentially change $\succeq_i$ to the single peaked order with peak at a. This then describes the sequence $\{\succeq(i)\}_{i=0}^{N}$. Observe that , by $(**)$, for $i = k(\succeq)$ we must have $\Phi(\succeq(i)) = b$ (resp. $a(\succeq)$) if and only if $\Phi(\succeq) = b(\succeq)$ (resp. $a(\succeq)$) – for brevity, call this a *co-monotonicity condition.* To show this, first consider the case where there is some $i < k(\succeq)$ for which $\Phi(\succeq(i)) = b(\succeq)$. Consider the first such i. If $i = 0$, then $(**)$ implies that $\Phi(\succeq(k(\succeq)) = b$. It follows that the co-monotonicity condition holds in this case. If $1 \leq i < k(\succeq)$ is the first such that $\Phi(\succeq(i)) = b(\succeq)$, then when the population profile is $\succeq(i-1)$ voter i will report an order with peak b, so that the reported profile is $\succeq(i)$. This yields a profitable deviation as reporting honestly yields $\Phi(\succeq(i-1)) = a(\succeq)$ and misreporting yields $\Phi(\succeq(i)) = b(\succeq)$, which is voter i's peak – hence, we have a contradiction to strategy-proofness. If the first i with $\Phi(\succeq(i)) \neq a(\succeq)$ is such that $\Phi(\succeq(i)) = b$ – viz. there is a "jump" in the SCF – then again consider incentives of agent i to report truthfully. Here is where "off-peak" preferences matter insofar as strategy-proofness is concerned. We know that for any two profiles $\succeq(i-1), \succeq'(i-1)$ for which the peaks of the orders are the same, we have (via strategy-proofness) $\Phi(\succeq(i-1)) = \Phi(\succeq'(i-1))$.

Fix an order with the associated peaks for all agents $j \neq i$. Consider two different orders for agent i, both with peak at $b(\succeq)$. For $\succeq_i$ we select an order with $a(\succeq) \succ_i b$, and for $\succeq_i'$ we select an order with $a(\succeq) \prec_i b$.[9] Now consider an amended sequence of preference profiles where we keep all agents $j \neq i$ fixed at $\succeq_j$ and change agent i's starting preference from $\succeq_i$ to $\succeq_i'$. Let $\succeq$ (resp. $\succeq'$) denote the corresponding population profiles and let $\{\succeq(i)\}, \{\succeq'(i)\}$ denote the corresponding sequences. Note that strategy-proofness implies that

[9]Note that this can always be done. Graphically, draw a utility curve representing a generic order by linearly interpolating the utility values of the numbers $u(x_1), \ldots, u(x_k)$ with $x^* := b(\succeq)$. This yields a tent-shaped shaped function. We select an order which underlies $u(\cdot)$ where the slope, i.e. left-derivative, of the tent is everywhere steeper and at a higher level to the left of $b(\succeq)$ than it is to the right. The underlying order is $\succeq_i$. Similarly, consider a tent peaked at $b(\succeq)$ for which the right-derivative is everywhere steeper and at a higher level than points to the left of $b(\succeq)$. The order underlying this function is our chosen $\succeq_i'$.

$\Phi(\succeq(i)) = \Phi(\succeq'(i)), \forall i$. Hence, if i is the first such that (i) $\Phi(\succeq(i)) \neq a(\succeq)$ and (ii) $\Phi(\succeq(i)) = b$, then this same voter i is that first such that (i) $\Phi(\succeq'(i)) \neq a(\succeq)$ and (ii) $\Phi(\succeq(i)) = b$. This implies that voter i has an incentive to misreport when the population profile is $\succeq'(i-1)$. By reporting $\succeq'_i$ the outcome is $\Phi(\succeq'(i-1)) = a(\succeq)$. By reporting the order with peak at b he gets $\Phi(\succeq'(i)) = b$, where – by construction of $\succeq'_i$ – we have $b \succ'_i a(\succeq)$. Hence, voter i will not truthfully report his preference when the true profile is $\succeq'(i-1)$, contradicting strategy-proofness.

The preceding argument shows that if $\Phi(\succeq) = a(\succeq)$, then for all $i \leq k(\succeq)$ we have $\Phi(\succeq(i)) = a(\succeq)$. Now consider the values of $\{\Phi(\succeq(i))\}_{i=k(\succeq)+1}^{N}$. If $\Phi(\succeq(k(\succeq))) = a(\succeq)$, then it follows from $(*)$ that $\Phi(\succeq(N)) = a$. Hence, we reduce to $\Phi(\succeq(k(\succeq))) = b$. The symmetric argument as in the preceding paragraph now shows that we must have, in this case, $\Phi(\succeq(i)) = b, \forall i \geq k(\succeq)+1$. This concludes the co-monotonicity argument. To recapitulate what has been shown: Given any two-peaked profile $\succeq$ with peaks $\{a(\succeq, b(\succeq)\}$ (with $a(\succeq) < b(\succeq)$) consider the associated two-peaked profile where we replace all $\succeq_i$ with $a_i(\succeq) = a(\succeq)$ with the order with peak at a, and all $\succeq_i$ with $a_i(\succeq) = b(\succeq)$ with the order with peak at b. Denote this amended profile as $\succeq_{\text{ext}}$ (where "ext" is for *extreme*). The co-monotonicity argument (denoted CMN) states the following:

$$\text{(CMN)}\ \Phi(\succeq) = a(\succeq)\ (\text{resp.}\, b(\succeq)) \Leftrightarrow \Phi(\succeq_{\text{ext}}) = a\ (\text{resp.}\, b).$$

Step 3: Characterizing the rule on two-peaked extreme profiles.
This is the step of the argument where we apply anonymity. Importantly, the co-monotonicity condition (CMN) does not invoke anonymity. Given any two-peaked $\succeq$ pass to the extreme companion $\succeq_{\text{ext}}$ (omit the subscript for the remander of the argument – with the caution that when we write $\succeq$ we are only considering extreme two-peaked profiles). Consider two sets, (abusing notation) $\Sigma_1 := \{\succeq : \Phi(\succeq) = a\}, \Sigma_2 := \{\succeq : \Phi(\succeq) = b\}$. Consider a profile $\succeq \in \Sigma_1$ that is minimal w.r.t. $|\{i \in N : a_i(\succeq) = a\}|$. We claim that Φ is a q-rule for

$$q_\Phi = \frac{|\{i \in N : a_i(\succeq) = a\}|}{N}.$$

Note that there is some indeterminacy in the choice of q as any $q \in (q_\Phi - \frac{1}{N}, q_\Phi + \frac{1}{N})$ does the job. Now we check that for any profile $\succeq$ such that $|\{i \in N : a_i(\succeq) = a\}| \geq q_\Phi + \frac{1}{N}$ we have $\Phi(\succeq) = a$ as well. Towards contradiction, say there is such a profile where $\Phi(\succeq') = b$. Denote the profile where $\Phi(\succeq) = a$ and $\succeq \in \Sigma_1$ is minimal w.r.t. $|\{i \in N : a_i(\succeq) = a\}|$ as $\succeq(q)$. Note that, by anonymity, the value of $\Phi(\cdot)$ is the same on any two profiles $\succeq, \succeq'$ in which the cardinality $|\{i \in N : a_i(\succeq) = a\}|$ is the same. Hence, when comparing $\succeq(q)$ with the alleged counterexample $\succeq$ we can always – via anonymity – arrange the peaks such that

$$\{i \in N : a_i(\succeq(q)) = a\} \subseteq \{i \in N : a_i(\succeq) = a\}.$$

Now construct a perturbation argument as before, let $\succeq(0) := \succeq(q), \ldots, \succeq(K) := \succeq$, where $K = |\{i \in N : a_i(\succeq) = a\} \backslash \{i \in N : a_i(\succeq(q)) = a\}|$. Sequentially flip preferences of voters in the set $\{i \in N : a_i(\succeq) = a\} \backslash \{i \in N : a_i(\succeq(q)) = a\}$ from $\succeq_i (q)$ (which has a peak at b) to $\succeq_i'$ (which has a peak at a). Consider the first i such that $\Phi(\succeq(i)) = b$. Note that voter i then has an incentive to report $\succeq_i(q)$ rather than $\succeq_i'$ (when the true profile is $\succeq(i)$) – contradicting strategy-proofness. We conclude by verifying that whenever $\succeq$ is such that $|\{i \in N : a_i(\succeq) = a\}| < q_\Phi - \frac{1}{N}$, we have $\Phi(\succeq) = b$. By anonymity, we can again assume that individual preferences are such that

$$\{i \in N : a_i(\succeq) = a\} \subseteq \{i \in N : a_i(\succeq(q)) = a\}$$

Hence, if $|\{i \in N : a_i(\succeq) = a\}| < q_\Phi - \frac{1}{N}$ we must have strict containment. Now apply the hypothesis that q_Φ is minimal s.t. $\succeq \in \Sigma_1$. It follows that $\Phi(\succeq) = b$, so that Φ is a q_Φ-rule on the set of profiles with two extreme peaks.

Step 4: Extension to all profiles.
We bootstrap to any profile to claim that Φ is a q_Φ-rule on the full space of profiles. Take any $\succeq$ and consider $\Phi(\succeq) = a_{i^*}(\succeq)$. Ordering peaks of the individual preferences $\succeq_i$ from bottom-to-top (counting multiplicity) we put $\{a_1(\succeq), \ldots, a_i(\succeq), \ldots, a_N(\succeq)\}$. Let a_i denote the ith N-tile w.r.t. this ordering, i.e. q_1 is the $\frac{1}{N}$th-tile, q_2 is the $\frac{2}{N}$th N-tile, and so on. Abusing notation, let a_{k_q} denote the peak located at the k_qth N-tile where $q := \frac{k_q}{N}$. We claim that $\Phi(\succeq) = a_{k_{q_\Phi}}$.[10] Consider three cases, (i) $\Phi(\succeq) > a_{k_{q_\Phi}}$, (ii) $\Phi(\succeq) < a_{k_{q_\Phi}}$, and (iii) $\Phi(\succeq) = a_{k_{q_\Phi}}$. We check that the first two cases yield a contradiction to the hypothesis that $\Phi(\cdot)$ restricts to a q_Φ rule on two-peaked profiles. In case (ii), we pass to a two-peaked profile $\succeq'$ with the property that $\Phi(\succeq) = \Phi(\succeq')$ and $\Phi(\succeq') = a_q = a(\succeq')$, where $a_{q+\frac{1}{N}} = b(\succeq') > a(\succeq')$. Passing to the extreme companion of this profile (which we, by abuse, also denote $\succeq'$) we obtain that $\Phi(\succeq') = a = a_q$ where $q < q_\Phi$. This contradicts the hypothesis that Φ restricts to a q_Φ-rule on two-peaked extreme profiles. Now consider case (i), where we allege $\Phi(\succeq) > a_{q_\Phi}$. Put $a(\succeq) :=$ max$\{a_i : a_i(\succeq) < \Phi(\succeq)\}, b(\succeq) = \Phi(\succeq)$ and consider the two-peaked profile where all voters with $a_i(\succeq) \geq \Phi(\succeq)$ have $\succeq_i$ replaced with an order peaked at $b(\succeq)$. Similarly, for each i with $a_i < \Phi(\succeq)$ replace each $\succeq_i$ with an order peaked at $a(\succeq)$. Call the resulting amended population profile $\succeq'$. The same argument as given in step 1 shows that $\Phi(\succeq) = \Phi(\succeq')$. Applying the CMN condition we obtain that – taking $\succeq''$ to be the extreme companion of $\succeq'$ – $\Phi(\succeq'') = b$, where $|\{i \in N : a_i(\succeq'') = a\}| > q_\Phi{\cdot}N$. This again contradicts the hypothesis that Φ restricts to a q_Φ-rule on two-peaked extreme profiles.

[10]Note that there is some indeterminacy (due to labeling of the peaks) when there are multiple voters whose peak coincides with $a_{k_{q_\Phi}}$. However, in this case the labeling of these voters does not affect the value of the rule Φ, so we ignore this issue in the argument.

Step 5: Conclusion.
The preceding argument shows that for any profile $\succeq$ of single-peaked preferences supported on X, we have a q_Φ such that the function $\Phi(\cdot)$ is a q_Φ-rule. The claim of the Theorem is that when we restrict to single-peaked profile supported on $A \subseteq X$ we get a q-rule with the same q_Φ. Fix a menu $A \subseteq X$ and given a profile $\succeq_A$ (supported on A) find a profile $\succeq_X$ supported on X which extends this profile in a manner such that the set of peaks $\{a_1(\succeq_X), \dots, a_n(\succeq_X)\}$ is the same as the set $\{a_1(\succeq_A), \dots, a_n(\succeq_A)\}$. By the preceding argument, Φ_X selects the q_Φ-th element (counting from the lowest-to-highest peak) from the set $\{a_1(\succeq_X), \dots, a_n(\succeq_X)\}$, which is the same as the q_Φ-th element from $\{a_1(\succeq_A), \dots, a_n(\succeq_A)\}$. $\square$

We now show that when we drop anonymity, but maintain strategy-proofness and sequential consistency, then the resulting class of SCF families $\{\Phi_A\}$ is precisely the set of vote by committee rules.

Theorem 2. *An SCF family $\{\Phi_A\}$ is strategy-proof and sequentially consistent if and only if it is a vote by committee rule.*

Proof. Necessity of sequential consistency is still obvious, but on account of the committee structure it might not be as immediate that strategy-proofness is also necessary, hence we provide proof for this here. Fix a vote-by-committee rule $(\mathcal{C}_i, \{q_i\})$ and a menu A to which this rule is being applied. Let $\Sigma(i)$ denote all the committees $\mathcal{C}_j$ to which voter i belongs. Fix the aggregate choice $\Phi(\succeq)$ and voter i's ideal $a_i(\succeq)$ (note: we suppress dependence on the ambient menu of options since that does not affect the argument). If $a_i(\succeq) < \Phi(\succeq)$, then voter i clearly has no incentive to misreport. If $a_i(\succeq) > \Phi(\succeq)$, then if voter i announces $\hat{\succeq}_i$ with a higher peak than $a_i(\succeq)$, the aggregate choice $\Phi(\succeq)$ is unchanged. Hence, the only way to shift the aggregate choice is to announce $\hat{\succeq}_i$ with a lower peak than $a_i(\succeq)$. Moreover, to shift $\Phi(\succeq)$ voter i must be on some committee $\mathcal{C}_j$ for which $q_j - \min_{k \in \mathcal{C}_j} a_k(\succeq) = \Phi(\succeq)$.[11] If he announces $\hat{\succeq}_i$ with a peak less than $\Phi(\succeq)$, then this shifts the value $q_j - \min_{k \in \mathcal{C}_j} a_k(\succeq)$ downwards, which is worse for him than announcing truthfully, since $\Phi(\succeq)$ is closer to $a_i(\succeq)$. If he announces $\hat{\succeq}_i$ with peak greater than or equal to $\Phi(\succeq)$, then this doesn't change the argument that attains $q_j - \min_{k \in \mathcal{C}_j} a_k(\succeq)$. Hence, vote-by-committee is strategy proof.

While the result nests Theorem 1, the proof does not proceed by reducing to the case where the rule is anonymous. Note that the only step of the argument for Theorem 1 which invokes anonymity is the step where we prove that the rule is a q-rule on two-peaked extreme profiles. The argument that shows

[11] Therefore, the argument for necessity of strategyproof-ness reduces to the necessity argument (which we omitted) for q-rules.

that the rule is determined by its values on these profiles proceeds *verbatim* as in Theorem 1. Hence, when we relax anonymity we need to compute the structure of the rule on two-peaked extreme profiles. By fiat, we set the quota of a singleton committee equal to 0.

Step 1: Characterization on two-peaked extreme profiles.
We will use variation in the family of two-peaked (extreme) profiles to extract committees and committee-dependent quotas. Fix a labeling of the agents, $i = 1, \ldots, N$, and consider profiles $\succeq^1 := (b, a, \ldots, a), \succeq^2 := (a, b, a, \ldots, a), \ldots, \succeq^N := (a, a, \ldots, a, b)$. Consider the values $\Phi(\succeq^i)$. If any $\Phi(\succeq^i) = b$, then voter i forms a singleton committee (with quota $q_i \equiv 0$). Let Π_1 denote the set of all voters who form singleton committees, and consider voters $i \in N \backslash \Pi_1$. We inductively elicit committees by considering consecutive sequences of profiles $\succeq^i$ where two voters have peaks at b, then sequences where three voters have peaks at b, and so on. To this end, for each voter i let $\mathcal{C}_i(2)$ denote the set of all (minimal) decisive coalitions of pairs of voters (where i is in the pair). That is, we have $\Phi(\succeq^{i,j}) = b$ where voters i, j have peak at b and all others have peak at a. Moreover, if we replace either $\succeq_i$ or $\succeq_j$ with a profile peaked at a, then the map Φ selects a – this is the minimality condition. For each pair $\{i, j\} \in \mathcal{C}_i(2)$ create a committee $\{i, j\} =: \pi_2(i, j)$ with quota $q_{i,j} = 1/2$. Do this for each voter i.

Now proceed to minimal decisive coalitions of size 3. For each triple $\{i, j, k\} \in \mathcal{C}_i(3)$ (extending notation in the obvious way) create a committee $\{i, j, k\} =: \pi_3(i, j, k)$ with quota $q_{i,j,k} = 1/3$. Inductively proceed so that for each voter i and each $\{i, i_1, \ldots, i_{k-1}\} \in \mathcal{C}_i(k)$ we create a committee $\{i, i_1, \ldots, i_{k-1}\} =: \pi_k(i, i_1, \ldots, i_{k-1})$ with quota $q = \frac{1}{k}$. Discarding repeats from the list we obtain a collection of committees $\mathcal{C} = \{\mathcal{C}_i\}$ along with a collection of committee-dependent quotas q_i. We claim that the corresponding vote-by-committee rule generated by the pair $(\mathcal{C}, \{q_i\})$ recovers the family $\{\Phi_A\}$ for two-peaked (extreme) profiles. Let $\Phi_{(\mathcal{C}_i, \{q_i\})}$ denote the rule just constructed and let Φ be the given (family of) SCF's. Fix a two-peaked profile $\succeq$ (with peaks at a and b with $a < b$). If $\Phi(\succeq) = b$, then there is some minimal coalition (equiv. committee) of voters $\mathcal{C}_i$ such that $\succeq_j$ has peak at b for all $j \in \mathcal{C}_i$. It follows that, since $q_i - \min_{j \in \mathcal{C}_i} a_j(\succeq) = b$, $\Phi(\succeq) = \Phi_{(\mathcal{C}_i, \{q_i\})}(\succeq)$ in this case. If $\Phi(\succeq) = a$, then by construction of the coalitions $\mathcal{C}_i$ we must have $q_i - \min_{j \in \mathcal{C}_i} a_j(\succeq) = a, \forall \mathcal{C}_i$. Hence, $\Phi(\succeq) = \Phi_{(\mathcal{C}_i, \{q_i\})}(\succeq)$ in this case as well.

Step 2: Extension to all profiles.
Let $\succeq$ denote any profile and let $\succeq'$ the two-peaked companion to this profile, where we put $a_i(\succeq') := \Phi(\succeq)$ if $a_i(\succeq) \leq \Phi(\succeq)$ and $a_i(\succeq') := \min\{a_i(\succeq) : a_i(\succeq) > \Phi(\succeq)\}$. Fixing the committee structure on two-peaked extremes that we found, denoted as $(\mathcal{C}, \{q_i\})$, we want to check that $\Phi(\succeq)$ can be recovered via a vote-by-committee rule using the same committee structure $(\mathcal{C}, \{q_i\})$. Formally,

let $\Phi_{(\mathcal{C},\{q_i\})}$ denote the choice function induced by this committee structure. We have checked in the previous step that $\Phi(\succeq) = \Phi_{(\mathcal{C},\{q_i\})}(\succeq)$ whenever $\succeq$ is a two-peaked profile. Here we extend the equality to all profiles by verifying, in turn, $\Phi(\succeq) \geq \Phi_{(\mathcal{C},\{q_i\})}(\succeq)$ and $\Phi(\succeq) \leq \Phi_{(\mathcal{C},\{q_i\})}(\succeq)$. For the first inequality, it suffices to check that committee-by-committee we have $q_i - \min_{j\in\mathcal{C}_i} a_i(\succeq) \leq \Phi(\succeq)$. To show this note that it suffices to check the inequality on the two-peaked companion $\succeq'$. But this is obvious by virtue of the equality $\Phi(\succeq) = \Phi_{(\mathcal{C},\{q_i\})}(\succeq)$. For the reverse inequality, note that if there is a committee $\mathcal{C}_i$ where $q_i - \min a_j(\succeq) > \Phi(\succeq)$, then for the associated two-peaked companion profile $\succeq'$ we would obtain $\Phi_{(\mathcal{C}_i,\{q_i\})}(\succeq') > \Phi(\succeq)$. This contradicts the hypothesis that Φ and $\Phi_{(\mathcal{C}_i,\{q_i\})}$ agree on two-peaked profiles.

Step 3: Conclusion.
We now show that the equality between Φ and $\Phi_{(\mathcal{C}_i,\{q_i\})}$ holds on any choice set $A \subseteq X$. Given any profile $\succeq_A$, let $\succeq_X$ denoted the profile supported on X which extends $\succeq_A$ such that peaks agree, viz. $a_i(\succeq_X) = a_i(\succeq_A)$. By sequential consistency, we know that $\Phi(\succeq_A) = \Phi(\succeq_X)$ (we suppress the subscript on the SCF, which denotes the ambient menu of options). By the preceding part we know that $\Phi(\succeq_X) = \Phi_{(\mathcal{C}_i,\{q_i\})}(\succeq_X)$. Hence, since the set of peaks under $\succeq_A$ and $\succeq_X$ are the same we recover $\Phi(\succeq_A)$ with the same committee structure $(\mathcal{C}_i, \{q_i\})$. □

Two comments on this result. First, note that committees and committee-dependent quotas are not uniquely pinned down by the aggregate choices Φ_A. To see this most starkly, consider the rule that picks the $\leq$-maximal element from each menu A. This can be recovered by two distinct committee structures, (i) $\mathcal{C}_i \equiv \{i\}$, $\mathcal{C} \equiv \cup_{i\in N} \mathcal{C}_i, q_i \equiv 0$, and (ii) $\mathcal{C}_i \equiv \{1, 2, \ldots, N\}, q_i = 1$. Hence, we cannot identify the committee-rule parameters from the aggregate choice data. This, at least partially, explains why we cannot immediately deduce Theorem 1 from Theorem 2 by imposing anonymity. Anonymity is consistent with there being a representation of the family Φ_A in which there is more than one committee. Hence, we cannot deduce the "axioms imply representation" direction by using the representation given in Theorem 2, a separate argument is required.

The second comment regards the construction in Theorem 2. Note that committees consist of minimal decisive coalitions, viz. in the sense that when all committee members have the same peak and everyone outside the committee has a lower peak, then the committee's preference determines the aggregate choice. Thus, while the max-min structure of the vote-by-committee rule might suggest a more democratic aggregate choice procedure, the construction itself shows that the choice rule behaves like an oligarchy. This points us to some subtle connections between the vote-by-committee rules in this paper and the class of committee voting rules characterized in Barberá et al. (1991).

3.1 Relation to Barberá et al. (1991)

Fix a finite set X and let 2^X denote the set of (possibly empty) subsets of X. Assume each individual voter is endowed with a menu preference, i.e. a binary relation on 2^X which is, as in our paper, assumed to be a strict order. Fixing N to be the set and (by abuse) also the cardinality of the set of voters let the space of strict menu orders be denoted as $\mathcal{P}^N$. A SCF is then defined as, $\Phi : \mathcal{P}^N \to 2^X$. We recall some definitions from Barberá et al. (1991).

Definition 7. A *committee* is a collection of coalitions $W \subseteq N$, denoted $\mathcal{W}$.

Definition 8. An SCF $\Phi : \mathcal{P}^N \to 2^X$ is a *vote-by-committee* rule if, fixing a profile $\succeq \in \mathcal{P}^N$, there is – for each $x \in X$ – a coalition $W_x \in \mathcal{W}$ such that the following property holds: (let $A(\succeq_i)$ denote the $\succeq_i$-maximal element in 2^X)

$$x \in \Phi(\succeq) \Leftrightarrow \{i \in N : x \in A(\succeq_i)\} \in \mathcal{W}$$

Interpret the aggregate choice as a set of candidates, or perhaps a single candidate who is (implicitly) identified with a set of attributes. For each individual voter, fixing his preference $\succeq_i$ over menus of attributes, consider the set of all voters (denoted $\mathcal{W}_x(\succeq)$) who rank x as a top attribute, i.e. voters i for whom $x \in A(\succeq_i)$. From a cooperative games perspective, we imagine that the aggregate menu is determined via an implicit bargaining process between the voters. The set $\Phi(\succeq)$ should then consist precisely of those attributes x for which *some* (minimal) coalition of voters who are decisive on issue x rank x within there respective top menu of attributes, $A(\succeq_i)$. This is the interpretation of the vote-by-committee rule described above.

The comparison with our vote-by-committee rule is more transparent when we use the construction in our proof of Theorem 2. There we showed that the (*prima facie*) more restrictive class of committee structures, $(\mathcal{C}, \{q_i\})$, where $\mathcal{C}$ was the set of all minimal decisive coalitions of voters and the quotas q_i were set equal to the (resp.) quorums $q_i = \frac{1}{|\mathcal{C}_i|}$. Our definition of "decisive" is a little more restrictive than usual. First, we only apply the concept to two-peaked profiles $\succeq$ where voters have peak at a or b (the extreme values of the peaks is not a necessary restriction). We say a coalition, $\mathcal{C}_i$, is decisive (and, implicitly, minimal) if preferences $\succeq_j$ for all the voters $j \in \mathcal{C}_i$ have peak at b, then the SCF selects b, whereas if even one voter's peak flips to a, the SCF selects a instead. Since the policy space is one-dimensional in our model, the aggregate policy is determined by (i) the decisive coalition whose common ideal point lies furthest to the right, i.e. $\leq$-maximal, on the policy dimension and (ii) when there is no decisive coalition, the ideal points of the vetoers (i.e. within each candidate decisive coalition this is the voter with the lowest ideal point).

The characterization result in Barberá et al. (1991) applies to strict menu orders $\mathcal{P}$ that satisfy a condition called *separability*, viz. a menu order $\succeq \in \mathcal{P}$ is called separable if $A \cup \{x\} \succ A$ (if $x \notin A$ and $\{x\} \succ \emptyset$). From the viewpoint of the decision theory literature on menu choice this means that the preference $\succeq$ exhibits a strict *preference for flexibility*, after Kreps (1979). Denoting strict menu orders which exhibit separability as $\mathcal{P}^S$ (one of) the main results in Barberá et al. (1991) is the following. In addition to strategy-proofness it invokes just one more (standard) restriction known as *citizen sovereignty*. This condition requires that the SCF is surjective onto its range.

Theorem 3 (Barberá et al. (1991))**.** *An SCF $\Phi : (\mathcal{P}^S)^N \to 2^X$ is a vote-by-committee rule if and only if it is strategy-proof and satisfies citizen sovereignty.*

The authors also show that this collapses to a quota rule[12] when the SCF is anonymous and neutral. On account of the domain restriction to separable menu preferences, there is no transparent way to map our result into the Barberá et al. (1991) framework, or vice-versa. Since the policy space is one-dimensional in our model, the menu preference associated to a preference on alternatives $x \in X$ is defined by the value function $U(A) := \max_{x \in A} u(x)$. Note that no menu preferences generated by a value function $U(\cdot)$ can be separable, hence the formal intersection of the two characterizations is the empty-set. Moreover, as Barberá et al. (1991) show, there is no larger domain of preferences that can accommodate strategy-proof vote-by-committee rules in our sense and also in the sense of Barberá et al. (1991).

4 Conclusion

This paper conducts a normative analysis of dynamic voting rules. Our motivation comes from aggregation procedures that are not, by nature, one-shot, i.e. there are often multiple stages, preferences are aggregated in each stage, and an aggregate decision is taken at each stage. In such situations, we consider two desiderata on the (dynamic) aggregation rule: (i) it induces truthful reporting of preferences in each stage and (ii) conditional on truthful reporting across all stages, the criteria used to aggregate preferences in different stages are consistent with one another. We use an axiom inspired by Sen's α from choice theory, called "sequential consistency", which intends to capture the internal consistency of the aggregation process and provide two characterization results. The first shows that when we require the aggregation procedure to be anonymous,

[12] In the Barberá et al. (1991) setting a quota rule (with quota Q) is a vote-by-commitee rule where membership of a given choice x in the aggregate menu $\Phi(\succeq)$ is determined by whether at least Q voters have x in their resp. most preferred menu $A(\succeq_i)$.

the only class of rules which are both (dynamically) strategy-proof and sequentially consistent are q-rules, i.e. the SCF selects the ideal point of the voter whose ideal lies at the q-th N-tile – where N is the population size. When we relax anonymity, the class of strategy-proof and internally consistent rules expands to a class that we dub "vote-by-committee" rules. A vote-by-committee rule determines the aggregate choice in two stages. First, voters are grouped into committees and a choice is determined within each committee using a q-rule (where the q will typically depend on the committee). Second, an aggregate choice is determined from the reported set of committee choices.

What we have dubbed a "vote by committee" rule can be interpreted as the analogue of the committee voting rules introduced in Barberá et al. (1991), when applied to the domain of single-peaked preferences.[13] An alternative interpretation of the committees, suggested by that paper, is that each committee is an implicit representation of a simple game (in the sense of cooperative game theory). Under this interpretation and when the social decision is binary, the threshold q corresponds to the voter with the lowest ideal point since he(she) is the one with the (within-committee) veto power. Vote by committee rules also nest common voting rules, e.g. majority, super-majority voting, and are themselves nested within the larger class of generalized median voter schemes, see e.g. Barberá (2001). Under anonymity, the choice of N-tile, q, provides a different parameterization (and interpretation) of the class of majority voting rules. Previous results on generalized median voter schemes suggest that all strategy-proof rules look like majority (i.e. median) voter rules when we add enough "phantom" voters. When the decision process has multiple stages, our result provides a different interpretation of these rules. We don't interpret the voting process as being influenced by phantoms, so that strategy-proof rules can look non-majoritarian. However, the deviation from non-majoritarian rule is "consistent" – all such families of rules are q-rules, with the same q being applied at every stage of the decision process. In this sense, our results provide a normative foundation for not just strict majoritarian rule, but also for the (discrete) spectrum of q-rules that lie between unanimity (e.g. all voters have veto power) and super-majoritarian (e.g. only large coalitions have veto power) voting procedures.

[13] The domain of preferences in Barberá et al. (1991) is the set of profiles of *menu preferences*, so that each voter has a ranking on candidates (equivalently, issues) but also on subsets of candidates. Accordingly, the SCF selects a subset of candidates as a function of reported profiles of menu preferences.

5 Appendix

Proof of Lemma 1. (1). The extensive-form is defined by (i) the transition process $p(\cdot,\cdot)$, (ii) the stage reporting mechanism Φ_A, (iii) cardinal utilities over terminal choices.[14], and (iv) a per-period transfer function $r_i^t(\cdot,\cdot)$. We pick up from the discussion in section 2 and show truth-telling is a backwards induction solution to the extensive-form. Assuming the result for time $t+1$ we obtain that truthful reporting is an SPNE if and only if for each reporting stage t it solves,

$$\max_{\hat{\succeq}\in\mathcal{R}_A} p(\emptyset|(A,\Phi_A(\hat{\succeq},\succeq_{-i})))\cdot u_i(\Phi_A(\hat{\succeq},\succeq_{-i}))+r_i^t(\hat{\succeq},\succeq_{-i}) \\ +\sum_B p(B|(A,\Phi_A(\hat{\succeq},\succeq_{-i})))\cdot U_i^{t+1}(\succeq|B),$$

where $U_i^{t+1}(\succeq|B)$ denotes the continuation utility (when the option set is B) under the hypothesis that truth-telling is an SPNE from time $t+1$ onwards and $r_i^t(\hat{\succeq},\succeq_{-i})$ denotes the time t transfer from player i when all other players report $\succeq_{-i}$ and he reports $\hat{\succeq}$. The fact that Φ_A is strategy-proof implies that, under the associated direct mechanism, truth-telling is a (strict) dominant strategy equilibrium when reporting lasts just one round and the game ends. Hence, given report $\succeq_{-i}$ find the profile $\succeq^*$ that maximizes $u_i(\Phi_A(\hat{\succeq},\succeq_{-i}))$. By the fact that truth-telling is strictly dominant, we have $\succeq^*=\succeq_i$ (i.e. the maximum occurs at agent i's true preference. Put $a^*:=\Phi_A(\succeq^*,\succeq_{-i})$ and $a':=\Phi_A(\succeq',\succeq_{-i})$. Also let $U_{t+1}(\succeq,\succeq_{-i}|B)$ denote the continuation utility when (i) the pool in stage $t+1$ is B and (ii) the continuation strategy is for all agents to truth-tell from time $t+1$ thereafter. Now we define a transfer function. Let $\mathbf{1}_E$ denote the indicator function on an event and let $E=\emptyset_a$ denote the event that the game ends after the stage t choice of $a\in A$ is made. Put

$$\varepsilon(\succeq^*,\succeq_{-i}):=\max_{a'}\sum_B[p(B|(A,a'))-p(B|(A,a^*))]U_{t+1}(\succeq|B)$$

$$r_t^i(\hat{\succeq},\succeq_{-i})=\begin{cases}\mathbf{1}_{(E=\emptyset_a)^c}[u(a^*)+\frac{1}{(1-p(\emptyset|(A,a^*)))}\cdot\max\{\varepsilon(\succeq^*,\succeq_{-i}),0\}], & \text{if } \hat{\succeq}=\succeq^*,\ p(\emptyset|(A,a^*)\neq 0,\\ \max\{\varepsilon(\succeq^*,\succeq_{-i}),0\}, & \text{if } \hat{\succeq}=\succeq^*,\ p(\emptyset|(A,a^*))=0,\\ 0, & \text{if } \hat{\succeq}\neq\succeq^*.\end{cases}$$

We check that this induces truth-telling from player i in round t (the same construction applies to all other players), assuming truth-telling is a SPNE in the continuation game (given any history). This is obvious if $p(\emptyset|(A,a^*))=0$. If

[14]Note that players report only ordinal preferences, whereas to compute payoffs and verify that truthful reporting is a best-response we need to make an *ad hoc* specification of cardinal preferences.

the probability of acceptance is non-zero, then the payoff from truth-telling, i.e. choosing $\hat{\succeq} = \succeq^*$, is (net of transfers)

$$u(a) + \text{max}_{a'} \sum_B p(B|(A,a'))\cdot U_{t+1}(\succeq|B).$$

Now observe that,

$$\begin{aligned} &u(a) + \text{max}_{a'} \sum_B p(B|(A,a'))\cdot U_{t+1}(\succeq|B) \\ &\quad \geq p(\emptyset|(A,a'))u(a') + (1 - p(\emptyset|(A,a'))\cdot \sum_B p(B|(A,a'))\cdot U_{t+1}(\succeq |B). \end{aligned}$$

where on the LHS we have the (net of transfers) payoff from truth-telling and on the RHS we have the (net of transfers) payoff from announcing any $\hat{\succeq}_i \neq \succeq_i^*$ where $\Phi_A(\hat{\succeq}_i, \succeq_{-i}) = a'$. Hence, $\hat{\succeq} = \succeq^*$ is a best-response in period t. $\square$

References

BARBERÁ, S. (2001): "An introduction to strategy-proof social choice functions," *Social Choice and Welfare*, 18, 619–653.

BARBERÁ, S., H. SONNENSCHEIN, AND L. ZHOU (1991): "Voting by Committees," *Econometrica*, 59, 595–609.

CHAMBERS, C. P. (2007): "Ordinal Aggregation and Quantiles," *Journal of Economic Theory*, 137, 416–431.

KREPS, D. M. (1979): "A Representation Theorem for 'Preference for Flexibility'," *Econometrica*, 47, 565–577.

MOULIN, H. (1980): "On strategy-proofness and single peakedness," *Public Choice*, 35, 437–455.

SCHMEIDLER, D. AND H. SONNENSCHEIN (1978): "Two Proofs of the Gibbard-Sattherthwaite Theorem on the Possibility of a Strategy-Proof Social Choice Function," in *Decision Theory and Social Ethics: Issues in Social Choice*, ed. by H. Gottinger and W. Leinfellner, D. Reidel Publishing Company.