

Nehring, Klaus; Puppe, Clemens

Working Paper

Multi-dimensional social choice under frugal information: The Tukey median as Condorcet winner ex ante by

KIT Working Paper Series in Economics, No. 160

Provided in Cooperation with:

Karlsruhe Institute of Technology (KIT), Institute of Economics (ECON)

Suggested Citation: Nehring, Klaus; Puppe, Clemens (2023) : Multi-dimensional social choice under frugal information: The Tukey median as Condorcet winner ex ante by, KIT Working Paper Series in Economics, No. 160, Karlsruher Institut für Technologie (KIT), Institut für Volkswirtschaftslehre (ECON), Karlsruhe

This Version is available at:

<https://hdl.handle.net/10419/271096>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Impressum

Karlsruher Institut für Technologie (KIT)
Fakultät für Wirtschaftswissenschaften
Institut für Volkswirtschaftslehre (ECON)

Kaiserstraße 12
76131 Karlsruhe

KIT – Die Forschungsuniversität in der Helmholtz-Gemeinschaft

Working Paper Series in Economics
No. 160, March 2023

ISSN 2190-9806

econpapers.wiwi.kit.edu

Multi-Dimensional Social Choice under Frugal Information: The Tukey Median as Condorcet Winner *Ex Ante* *

Klaus Nehring[†] Clemens Puppe[‡]

March 2023

Abstract

We study a voting model with partial information in which the evaluation of social welfare must be based on information about agents' top choices plus qualitative background conditions on preferences. The former is elicited individually, while the latter is not. The social evaluator is modeled as an imprecise Bayesian characterized by a set of priors over voters' complete ordinal preference profiles. We apply this 'frugal aggregation' model to multi-dimensional budget allocation problems and propose a solution concept of 'ex-ante' Condorcet winners. We show that if the social evaluator has symmetrically ignorant beliefs over profiles of quadratic preferences, the ex-ante Condorcet winners refine the set of Tukey medians (Tukey, 1975).

Keywords: Social choice under partial information; participatory budgeting; frugal aggregation; ex-ante Condorcet approach; Tukey median.

JEL classification: D71

*This paper supersedes and significantly extends material that first appeared in the unpublished working paper Nehring and Puppe (2022). That paper has been split into two separate papers of which the present is the first. Versions of this work have been presented at the Online Social Choice Seminar Series, May 2021, the D-TEA conference in Paris, July 2022 and in seminars at Université Libre de Bruxelles, TU Munich, Nuffield College, and in the joint Microeconomics Seminar of the University and the ETH Zürich. We are grateful to the audiences for helpful feedback and comments. Special thanks to Jérôme Lang for his detailed feedback and pointers to a number of relevant contributions in the computer science literature. All errors are our own.

[†]Dept. of Economics, University of California at Davis, kdnehring@ucdavis.edu.

[‡]Dept. of Economics & Management, Karlsruhe Institute of Technology, clemens.puppe@kit.edu.

1 Introduction

Many economic and political decisions involve the allocation of resources under a budget constraint. Examples are the allocation of public goods, the redistribution across classes of beneficiaries, the allocation of tax burden, the choice of intertemporal expenditure streams, or the macro-allocation between expenditure, tax receipts and net debt. Here we explore the possibility of taking these decisions collectively by voting. This will be done in a somewhat more general and abstract multi-dimensional setting in which alternatives are elements of a convex subset of a Euclidean space and preferences are convex. The budget allocation problem is the special case in which the set of alternatives is a budget hyperplane.

Standard approaches to preference aggregation and voting assume ordinal or even cardinal preference information as their input. Their application to budget allocation problems poses substantial difficulties. First, at the foundational level, except for the one-dimensional case with two public goods and single-peaked preferences (Black, 1948; Arrow, 1951/63), one is faced with generic impossibility results under all reasonable domain restrictions (Kalai *et al.*, 1979; Le Breton and Weymark, 2011) just as in spatial voting models (Plott, 1967). In particular, in higher dimensions there is no hope to generally find a Condorcet winner even if all agents have well-behaved preferences. Indeed, the indeterminacy of majority voting is generic and can be severe; for example, generically every alternative can be the outcome of a dynamic (non-strategic) majority vote for an appropriate agenda (McKelvey, 1979). Thus, from the point of view of ordinal social choice theory, it is not even conceptually clear what allocations an optimal voting rule should aim at.

Second, at a pragmatic level, a basic problem already arises from the sheer number of alternatives which grows exponentially in the number of dimensions (i.e. alternative uses of the public resource). Articulating and communicating a complete ordering over the set of all alternatives for each agent (whether citizen or representative) is often simply infeasible. Clearly, much is to be said for making the task of the voter as easy as possible. Here, we take a minimalist approach by assuming that only voters' preference tops are individually elicited. As a collateral benefit, as detailed below, the parsimony of the informational basis entailed by our approach allows one to overcome the foundational indeterminacies of the classic preference aggregation model.

Tops-Only Information

We focus on the tops-only information assumption firstly because of its simplicity and prevalence in practice. Note that knowledge of voters' tops is indeed required to arrive at a reasonable decision in the most elementary instances of aggregation, namely those of unanimity.

Elicitation of tops is simple in that it requires every voter only to determine what she must in order to arrive at a choice all by herself. Reliance on tops-only elicitation thus addresses a fundamental tension in the standard ordinal aggregation framework, as the elicitation of a complete ordinal ranking requires much more cognitive effort on part of the individuals than would be required for solo decision making, while individual incentives to figure out ones own preferences are greatly reduced due the diluted impact of a voter on the final choice.¹

The Social Evaluator as an Imprecise Bayesian

This paper aims at determining which of the feasible social choices (here: allocations) represent social welfare optima in the light of the available information. Due to the lack of knowledge of the profile of complete preferences underlying an elicited profile of tops, the social evaluator faces a decision problem under uncertainty. Besides the individually elicited tops, the social evaluation may also be based on background knowledge about the structure of voters preferences, such as preference convexity. Probabilistic judgments may play a role as well. So the social evaluator will be modeled as an 'imprecise Bayesian' whose epistemic state is described as a set of admissible probability measures ('priors') over profiles of ordinal preferences compatible with a given profile of top choices and the available background information.

Within this framework, one might want to postulate the evaluator to have precise probabilistic beliefs (i.e. unique priors). But this approach has limited appeal here. In particular, on what evidential basis is the evaluator to make the manifold subjective judgments required for a precise Bayesian approach? Are there any sound reference models to sensibly describe ignorance priors over a state space of profiles of ordinal preferences on a continuous, multi-dimensional domain? Indeed, *whose* subjective probability is supposed to serve as the basis of the evaluation? If the social evaluator

¹This theme of 'rational ignorance' goes back to Down's classic treatment (Downs, 1957, pp. 244-246, 266-271).

was understood as a social planner (‘bureaucrat’), one may think of the required judgmental input as reflecting the planner’s expertise; but in a voting context, the social evaluator is naturally viewed as representing ‘the group’ at a constitutional stage at which individual preference profiles are unknown.

Instead of assuming a precise Bayesian prior, alluding to the notion of ‘fast and frugal heuristics’ due to Gigerenzer and Goldstein (1996), in our ‘frugal’ approach we rely on a qualitative specification of the social evaluator’s beliefs reflecting minimally demanding informational assumptions. Our aim is to show that even from these minimalist premises, attractive and credible choice implications can be derived.

The Ex-Ante Condorcet Approach

To determine ‘ex-ante’ optimal social choices, we propose a novel ex-ante Condorcet (EAC) approach. The EAC approach relies on ex-ante comparisons between arbitrary pairs of alternatives. These comparisons are based on the interval of expected majority counts consistent with the evaluator’s imprecise set of priors. A simple yet fundamental observation shows that the pairwise comparisons can be made in canonical manner independently of subjective attitudes of pessimism vs. optimism, or ambiguity aversion vs. ambiguity proneness. The EAC approach then uses this ex-ante majority relation to select an ex-ante Condorcet winner if it exists, and settles for some Condorcet extension rule – left unspecified here – if not. Remarkably, in the models at the center of this paper, ex-ante Condorcet winners do exist and can be characterized explicitly.

The Plain Convex Model

An obvious starting point in the context of public resource allocation is to assume knowledge of preference convexity (together with knowledge of the tops), and complete ignorance about anything else. We shall refer to this as the *plain convex model* of the evaluator’s beliefs. The plain convex model is very successful in the one-dimensional (two goods) case in which convexity is tantamount to single-peakedness of ex-post preferences. As the ex-post Condorcet winner is the median of voters’ tops, it is known ex-ante and equal to the ex-ante Condorcet winner.²

²We use the ex-ante vs. ex-post metaphor purely for conceptual purposes in order to describe the epistemic state of the social evaluator, without any assumption of an ex-post stage in real time at which the actual profile of (‘ex-post’) preferences is observed.

But in the multi-dimensional case (at least three competing uses of resources), convexity by itself loses much of its bite. In particular, with tops in general position, convexity does not permit any significant novel inferences about preferences beyond those available from knowledge the tops; by consequence, all tops are ex-ante Condorcet winners (Proposition 3). This appears quite counterintuitive and unsatisfactory, since any notion of centrality of the ex-ante Condorcet winner is lost, in stark contrast to the one-dimensional case.

Looking more closely, this negative result indeed hinges on extreme cases involving special ex-post profiles which appear unlikely a priori. Heuristically, one would want to rule out such cases and obtain more plausible majority intervals by assuming that preferences over pairs depend on the preference tops in a regular manner.

Symmetric Quadratic Models

To execute this formally, we assume a parametric form of convex preferences, namely quadratic preferences. A particular quadratic form \mathcal{Q} describes the substitution-complementation structure of a quadratic preference ordering in terms of the cross-partial of the utility function. Notably, assuming quadraticity does not help by itself to overcome the counterintuitive implications of the plain convex model, for the expected majority counts remain the same as the plain convex model (Fact 4.1).

Yet things change significantly once it is assumed that the evaluator's beliefs are symmetric in the sense that, for each admissible prior, the marginal distribution over quadratic forms is the same across voters irrespective of their top. Heuristically, symmetry expresses the idea that the evaluator lacks any grounds a priori to form different probabilistic beliefs about the unknown quadratic preference structure of different voters; in particular, the knowledge of voters' tops does not form such a ground. In addition, we also assume that the social evaluator is completely ignorant about the preference structure for each voter in isolation, just as in the plain convex model. These assumptions define the class of *symmetrically ignorant quadratic (s.i.q.) models* of the evaluator's beliefs. The main result of the paper, Theorem 1, shows that in any s.i.q. model ex-ante Condorcet winners exist and coincide, when unique, with the classical Tukey median (Tukey, 1975). When not unique, the EAC winners coincide with a well-defined refinement of the set of Tukey medians. The Tukey median is a well studied coordinate-free generalization of ordinary medians to multiple dimensions, see Small (1990) for a classic survey and Rousseeuw and Hubert (2017)

for comprehensive treatment.

Related Literature

To the best of our knowledge, the present EAC approach and its application to the ‘frugal aggregation’ model of budget allocation are new to the literature. But there are, of course, related approaches in the literature. Indirectly, the Tukey median has been studied in the social choice literature inasmuch as it is equivalent to the outcome of the minimax voting rule in standard spatial voting with Euclidean preferences (Kramer, 1977; Demange, 1982; Caplin and Nalebuff, 1988). This model can be viewed as a degenerate frugal model in which voters preferences conditional on their top are known. But with this additional, sub-top preference information, the Tukey median is no longer welfare optimal as we shall argue in Section 6.2.

Most work of theoretical interest in the problem of incomplete information as studied here has come from the computer science literature, see Boutilier and Rosenschein (2016) for an overview.³ One strand explores the implications of partial knowledge of complete (ex-post) preference profiles for inferences about the outcome of standard social choice rules and criteria, e.g. via the notions of ‘possible’ vs. ‘necessary’ winners (Konczak and Lang, 2005). A rather small strand in the literature adopts a decision-theoretic ex-ante approach as the present paper does. Some papers seek solutions that maximize expected welfare based on some utilitarian welfare criterion and a probability distribution over profiles, frequently uniform. Others argue for the modeling of the social evaluator’s epistemic state in terms of a set of possible profiles, as we do, and propose to apply classical criteria of decision making under ignorance such as maximin or minimax regret (Lu and Boutilier, 2011). In the highly complex state spaces associated with the epistemic models studied here, it may be very difficult to execute these approaches if that is possible at all. Significantly, the two quoted strands share the major conceptual limitation of having to rely on an interprofile-comparable standard of aggregate welfare ex post. Thus, they in fact suppose that the Arrovian problems of coherent aggregation and interpersonal non-comparability have been solved or assumed away, e.g. by assuming strong forms of utilitarian aggregation ex post.

By contrast, the EAC approach introduced here rests on an evaluation of decisions

³We thank Jérôme Lang who pointed us to the pertinent literature.

in pairs of alternatives taking the full state space (set of possible profiles) into account. In such pairwise comparisons, the majority criterion carries over naturally to the ex-ante stage, without raising new issues of interpersonal comparison, and allowing a tractable characterization in many cases. These pairwise comparisons need then be put together to obtain a coherent rationale for an ex-ante evaluation of complex choices such as budget allocations. At this juncture, Arrowian style issues of coherent aggregation might arise in principle. It is a rather remarkable finding of this paper that, in the models studied here, these problems do not arise.

With respect to the focal application to the allocation of public budgets, there is also an important recent literature on ‘participatory budgeting’ with intended application to cities and local communities (Shah, 2007). Participatory budgeting schemes have been put into practice at various scales in many places around the world. The ballots are typically very parsimonious, often taking the form of a set of projects approved.⁴ Again, most of the theoretical contributions come from the computer science community, with a focus on indivisibilities and on ‘proportionality’ considerations to ensure that the interest of different local subcommunities are fairly represented (Aziz and Shah, 2020). By contrast, our focus is on continuous divisible budgets, and on finding allocations that best satisfy the aggregate interest (in parallel with most of standard voting theory).

Overview of Paper

The remainder of this paper is organized as follows. In the next Section 2, we introduce the general EAC approach. The subsequent two sections apply it to the budget allocation problem under the assumption that the social evaluator knows the profile of voters’ top alternatives (the ‘frugal aggregation’ model). Section 3 studies the plain convex model which assumes in addition knowledge of convexity of voters’ preferences but complete ignorance about anything else. Proposition 3 shows that, generically, the ex-ante Condorcet winners coincide with the voters’ tops in the plain convex model. By contrast, our main result, Theorem 1 in Section 4, demonstrates that in the symmetric quadratic model the ex-ante Condorcet winner coincide with (a refinement of) the Tukey median. Section 5 extends this result to the case of prob-

⁴See, for instance, the open source project ‘Stanford Participatory Budgeting Platform’ (<https://pbstanford.org>) which offers guidance and allows municipalities, cities and other institutions to run participatory budgeting elections online.

abilistic uncertainty about voters' tops. Sections 6.1 and 6.2 discuss the robustness of our analysis with respect to the specific epistemic assumptions about the social evaluator. Section 6.3 offers some considerations on the use of the Tukey median as a voting mechanism played by self-interested voters. While the Tukey median is frequently manipulable – unsurprisingly in view of strong impossibility results such as Zhou (1991) –, we point out two ways in which it is notably attractive in terms of its incentive properties, especially in contrast to salient alternatives such as the multi-dimensional mean.

2 Condorcet Winners, Ex-Ante

We envisage a social evaluator who has to choose from a universe of alternatives X on behalf of a group of $n \in \mathbb{N}$ voters under uncertainty about their preferences. The social evaluator is modeled as an ‘imprecise’ Bayesian decision maker, i.e. his epistemic state is described by a set of probability distributions over ‘admissible’ profiles $\succ = (\succ_1, \dots, \succ_n)$ of true (‘ex-post’) preferences.

Concretely, denote by π a probability measure over profiles $(\succ_1, \dots, \succ_n)$ of complete preference orderings over X , and by Π a non-empty set of admissible such priors.⁵ The social evaluator is completely ignorant as to which probability distribution in Π is the most appropriate and therefore needs to take into account all of them.

Often one will be interested in cases in which the priors in Π satisfy specific additional properties. For instance, an important special case in the following will involve $X \subseteq \mathbb{R}^L$ and the assumption that all priors are concentrated over profiles of *convex* preferences.

For all distinct $x, y \in X$, a prior $\pi \in \Pi$ induces an expected support count $m_\pi(x, y)$ of votes for x against y , i.e.

$$m_\pi(x, y) := E_\pi [\#\{i : x \succ_i y\}],$$

where E_π denotes the expectation operator with respect to the probability distribution π . Thus, a set of priors induces an interval $m_\Pi(x, y)$ of expected support counts in

⁵To make this fully rigorous, one needs to specify a measure space on the set of profiles. For our purposes, the essential property is that, for each agent i and all alternatives x and y , the ‘event’ that agent i prefers x to y represents a measurable set.

the vote of x against y ,

$$m_{\Pi}(x, y) := [m_{\Pi}^{-}(x, y), m_{\Pi}^{+}(x, y)],$$

where

$$m_{\Pi}^{-}(x, y) := \inf_{\pi \in \Pi} m_{\pi}(x, y), \quad (2.1)$$

$$m_{\Pi}^{+}(x, y) := \sup_{\pi \in \Pi} m_{\pi}(x, y). \quad (2.2)$$

The family of these intervals will be what matters in our analysis. In deciding ex-ante on a hypothetical choice between x and y , it is natural to base this choice on a comparison of the intervals $m_{\Pi}(x, y)$ and $m_{\Pi}(y, x)$. Due to the imprecision of priors, the intervals $m_{\Pi}(x, y)$ and $m_{\Pi}(y, x)$ may well overlap in general. But due to the additivity of the complementary vote counts for x against y and for y against x , a comparison of the lower and upper expected counts must yield the same result. This evidently holds if preferences are known to be strict ex-post. To guarantee it more generally, the following regularity condition is needed which ensures that possible indifferences play a negligible role; this condition is satisfied in all applications considered in the following, and we maintain it throughout. Say that a set of priors Π is *regular* if for all priors $\pi \in \Pi$ and all pairs $x, y \in X$ of distinct alternatives, there exists a prior π' such that $\pi'(x \sim_i y) = 0$ for all $i = 1, \dots, n$, and $m_{\pi'}(x, y) \leq m_{\pi}(x, y)$. Thus, regularity guarantees that, for any pair $x, y \in X$, the minimal/infimal expected support for x against y is realized by a prior for which all indifferences between x and y have zero probability.

Proposition 1. *Let Π be regular. For all θ and all distinct $x, y \in X$,*

$$m_{\Pi}^{-}(x, y) \geq m_{\Pi}^{-}(y, x) \iff m_{\Pi}^{+}(x, y) \geq m_{\Pi}^{+}(y, x). \quad (2.3)$$

(Proof in appendix.)

By Proposition 1, an *unambiguous* balance of uncertainties ex-ante is possible; in contrast to the classical theory of decision making under ignorance (Luce and Raiffa, 1957), there is no need or even meaningful role for an evaluator's degree of pessimism vs. optimism (ambiguity aversion vs. ambiguity proneness in more modern

terminology).⁶

The **ex-ante majority relation** R_Π (for regular Π) is now defined as follows. For all distinct $x, y \in X$,

$$\begin{aligned} xR_\Pi y &: \iff m_\Pi^-(x, y) \geq m_\Pi^-(y, x) \\ &\iff m_\Pi^+(x, y) \geq m_\Pi^+(y, x). \end{aligned} \tag{2.4}$$

The maximal elements with respect to the ex-ante majority relation are referred to as the **ex-ante Condorcet winners**, i.e.

$$\text{CW}(\Pi) := \{x \in X \mid xR_\Pi y \text{ for all } y \in X\}.$$

An aggregation rule is called **ex-ante Condorcet consistent** if it selects all ex-ante Condorcet winners (if there are any).

In the following, we will refer to a set of priors Π as a **model** (of the evaluator's epistemic state). Moreover, we will say that two models are *equivalent* if they induce the same expected majority intervals. Note that, trivially, sets of priors with the same convex hull are equivalent, but the converse need not be true. Evidently, two equivalent models induces the same set of ex-ante Condorcet winners, i.e. $\text{CW}(\Pi') = \text{CW}(\Pi)$ whenever Π' and Π are equivalent.

3 The Plain Convex Model

In the rest of this paper, we will study the case in which X is a convex subset of \mathbb{R}^L for some $L \in \mathbb{N}$, and all preferences in any profile are convex. For our purposes, the following notion of convex preference will be useful. A weak order \succsim on $X \subseteq \mathbb{R}^L$ is *convex* if, (i) for all $x, y, z, w \in X$, $y = t \cdot x + (1 - t) \cdot z$ for some $0 \leq t \leq 1$, $x \succsim w$ and $z \succsim w$ jointly imply $y \succsim w$, and (ii) for all $x, y, z \in X$, $y = t \cdot x + (1 - t) \cdot z$ for some $0 < t < 1$, and $x \succ z$ jointly imply $y \succ z$.⁷

An important economic application is the *budget allocation problem* in which X takes the form of a budget hyperplane. Concretely, consider a group of agents that

⁶Nor is there a conflict – possibly even threatening an Arrow-like impossibility – between axioms of choice consistency and of independence; see Milnor (1954); Arrow (1960); Nehring (2000, 2009).

⁷Observe that (ii) is clearly implied by but significantly weaker than strict convexity. For instance, linear preferences satisfy both conditions (i) and (ii) but are not strictly convex.

has to collectively decide on how to allocate a fixed budget, normalized to unity, to a number L of public goods. Assuming given prices, the problem is fully determined by specifying the expenditure shares. The corresponding allocation problem can thus be modeled as the choice of an element of the following $(L - 1)$ -dimensional polytope:

$$X := \left\{ x \in \mathbb{R}^L \mid \sum_{\ell=1}^L x^\ell = 1 \text{ and } x^\ell \geq 0 \text{ for all } \ell = 1, \dots, L \right\}, \quad (3.1)$$

where $x = (x^1, \dots, x^L)$. Convex preferences are entirely standard in this context.

Other applications include the spatial voting model in which the coordinates represent different issues and alternatives represent political positions on these issues (Downs, 1957), or the collective choice of design of projects positioned in a characteristics space in the sense of Lancaster (1966).

The model of all priors with convex preferences on X without any further restriction is referred to as the **plain convex model** and denoted by Π_{co} .

3.1 Certainty about Tops

To simplify the task of the social evaluator, we assume first that the evaluator knows the top choices of voters. Concretely, denote by $\theta = (\theta_1, \dots, \theta_n)$ the profile of the voters' top alternatives which we assume to be unique. The epistemic state of the social evaluator will now be denoted by Π_{co}^θ to indicate the knowledge of θ . Here, the set Π_{co}^θ is assumed to consist only of priors π that are *compatible* with the top profile θ in the sense that every profile $\succ = (\succ_1, \dots, \succ_n)$ in the support of π has $\theta = (\theta_1, \dots, \theta_n)$ as the corresponding top profile.

3.2 The One-Dimensional Case: Median Voting

In the one-dimensional case, our notion of preference convexity is equivalent to the standard notion of single-peakedness, and the choice of the median top(s) constitutes the unique ex-ante Condorcet consistent aggregation rule; specifically, we have the following result. For every profile $\theta = (\theta_1, \dots, \theta_n)$, denote by θ_{med} the unique median if n is odd, and by $[\theta_{\text{med}^-}, \theta_{\text{med}^+}]$ the median interval if the number of voters is even.

Proposition 2. *Suppose that $X \subseteq \mathbb{R}$, and let $\theta = (\theta_1, \dots, \theta_n)$ be a profile of tops in X . Then,*

$$\text{CW}(\Pi_{\text{co}}^\theta) = \begin{cases} \{\theta_{\text{med}}\} & \text{if } n \text{ is odd} \\ [\theta_{\text{med}^-}, \theta_{\text{med}^+}] & \text{if } n \text{ is even} \end{cases}.$$

(Proof in appendix.)

Thus, in the one-dimensional case the ex-post and ex-ante Condorcet criterion give the same result under single-peakedness. The reason is, evidently, that under knowledge of single-peakedness any given top uniquely determines the preference on both sides of the top, and that is all what is needed to apply the Condorcet criterion.

3.3 The Multi-Dimensional Case: Generic Plurality Rule

In the multi-dimensional case, a result similar to Proposition 2 holds if the top profile is contained in a one-dimensional subspace; but in general, in the plain convex model the ex-ante Condorcet winners essentially coincide with the plurality winners.

In the following, we say that a set of points $Y \subseteq \mathbb{R}^L$ is *in general position* if no three elements of Y are collinear. The crucial observation for the plain convex model is that, if θ, x, y are not collinear, then there exist convex preferences \succ and \succ' with top θ such that $x \succ y$ and $y \succ' x$. This implies the following characterization of the ex-ante Condorcet winners in the plain convex model. For its formulation, it will be useful to identify profiles of individual tops with *type profiles* of tops with different counts. Specifically, we denote by $\theta = (\theta_1; p_1, \dots, \theta_m; p_m)$ the anonymous profile in which the fraction p_i of all voters has top θ_i , where $0 < p_i \leq 1$ and $\sum_i p_i = 1$; in that context, we also refer to θ_i as the *type* of voter i and assume without of loss of generality that the θ_i are pairwise distinct.

Proposition 3. *Consider a type profile $(\theta_1; p_1, \dots, \theta_m; p_m)$ such that $\{\theta_1, \dots, \theta_m\} \subseteq X$ are in general position. If p_{i^*} is maximal among $\{p_1, \dots, p_m\}$, then $\theta_{i^*} \in \text{CW}(\Pi_{\text{co}}^\theta)$. Moreover, if p_{i^*} is uniquely maximal among $\{p_1, \dots, p_m\}$, then*

$$\text{CW}(\Pi_{\text{co}}^\theta) = \{\theta_{i^*}\}.$$

(Proof in appendix.)

This is somewhat paradoxical. Intuitively it would appear that preference convexity contains substantial information beyond knowledge of the tops but Proposition 3

appears to contradict this. What is amiss?

Example 1. Consider a set of voters with pairwise distinct tops in a set U . In addition, suppose that two voters are concentrated at a point x outside U (see Figure 1). If all tops in U plus the point x are in general position then, according to Proposition 3, x is the unique ex-ante Condorcet winner. Indeed, for any point $z \neq x$, $m_{\Pi_{\text{co}}}^-(x, z) = 2$ while $m_{\Pi_{\text{co}}}^-(z, x) \leq 1$, or equivalently, $m_{\Pi_{\text{co}}}^+(x, z) \geq n - 1$. Note that the expected majority intervals are extremely wide, and the ex-ante Condorcet winner is left to ‘grasp for straws’ in picking the optimal alternative that happens to be the top of two voters rather than just of one. Nonetheless, if the epistemic state of the social evaluator is literally that of complete ignorance within Π_{co}^θ , then the ex-ante preference for x over any other alternative z seems defensible.

However, this rationale is not very robust. Consider in particular the comparison of x to y where y is sufficiently close to x and ‘between’ x and U as shown in Fig. 1. Note that for x to be preferred to y by some voter with top θ_i in U , i ’s preference must be very special; for instance, geometrically, only rather special ellipses with center at θ_i that include x will not include y .




Figure 1: Illustration of Proposition 3

The conceivable convex preference for x against y of a voter with top in U , on which the conclusion in Example 1 hinges, seems very unlikely a priori. It would therefore be desirable to capture this intuition by an appropriate specification of somewhat more precise evaluator’s beliefs. The challenge is to describe these beliefs in a qualitative manner that is weak enough to be acceptable on slim information while at the time sufficiently strong to have substantive implications. This task is at the heart of this paper and is taken on in the next section.

4 Symmetrically Ignorant Quadratic Models

Our proposal for modeling the evaluators beliefs in a more appropriate and specific manner involves two key features: First, we conceptually separate voters' tops from the substitution vs. complementation structure of preferences, and secondly, we assume that this substitution vs. complementation structure is 'ex-ante independent' of the tops. The first feature allows one to model preferences as quadratic; the second feature means that knowledge of the tops is not informative ex-ante for the substitution vs. complementation structure described by the quadratic preferences.

Specifically, say that a preference \succsim on X is *quadratic* if it can be represented ordinarily by a utility function of the form

$$u_{\theta_i}(x) = -(x - \theta_i)^T \cdot \mathcal{Q}_i \cdot (x - \theta_i), \quad (4.1)$$

for some $\theta_i \in X$ and a positive definite, symmetric $L \times L$ matrix \mathcal{Q}_i . Geometrically, the representation in (4.1) means that the indifference curves are ellipsoids generated from circles with center θ_i by a common affine transformation. The special case in which the quadratic form \mathcal{Q}_i is the identity matrix \mathcal{I} corresponds to the case of *Euclidean preferences* which has been extensively studied in the literature on spatial voting (Austen-Smith and Banks, 1999).

The cross-partial derivatives given by \mathcal{Q}_i capture the specific pattern of complementarities and/or substitutabilities between different goods. Quadratic preferences can thus also be viewed as (second-order) Taylor approximations of arbitrary smooth convex preferences. Denote by $\Pi_{\text{quad}} \subseteq \Pi_{\text{co}}$ the model consisting of all sets of priors over profiles of quadratic preferences on X , the *plain quadratic model*. Evidently, for all tops $\theta_i \in X$ and all $x, y \in X$ such that θ_i, x, y are not collinear, there exist quadratic preferences \succsim_i, \succsim'_i both with top θ_i such that $x \succ_i y$ and $y \succ'_i x$. By consequence, we have:

Fact 4.1. *The models Π_{quad} and Π_{co} are equivalent. In particular, the two models induce the same ex-ante majority relation and $\text{CW}(\Pi_{\text{quad}}) = \text{CW}(\Pi_{\text{co}})$.*

Thus, the plain quadratic model can be viewed as a parametrized version of the plain convex model. In particular the 'generic plurality' conundrum posed by Example 1 continues to apply to the plain quadratic model. But the great boon of the quadratic model is that it allows for a clear separation between the preference top and the

preference *structure* (described by the quadratic form \mathcal{Q}_i). This will be the key in our proposed resolution of the puzzle posed by Example 1.

Specifically, the epistemic state of the evaluator is given by a set of priors Π with state space $X_1 \times \dots \times X_n \times \mathcal{Q}_1 \times \dots \times \mathcal{Q}_n$ where X_i is the set of possible tops for voter i and \mathcal{Q}_i the set of possible quadratic forms (symmetric and positive definite $L \times L$ matrices) for voter i . For every prior $\pi \in \Pi$ and all $i = 1, \dots, n$, denote by π_{X_i} and $\pi_{\mathcal{Q}_i}$ the marginal distributions induced by π on X_i and \mathcal{Q}_i , respectively.

In the remainder of this section, we will impose the following conditions on a model Π . For all $x \in X$, denote by $\delta(x)$ the degenerate probability distribution that puts unit mass on x ; similarly, for all $\mathcal{Q} \in \mathcal{Q}$, denote by $\delta(\mathcal{Q})$ the degenerate probability distribution that puts unit mass on \mathcal{Q} .

1. **Concentration on Quadratic Preferences.** $\Pi \subseteq \Pi_{\text{quad}}$.
2. **Tops Certainty.** For all $\pi \in \Pi$ and all i , $\pi_{X_i} = \delta(\theta_i)$ for some $\theta_i \in X$.
3. **Symmetry.** For all $\pi \in \Pi$ and all i, j , $\pi_{\mathcal{Q}_i} = \pi_{\mathcal{Q}_j}$.
4. **Complete Ignorance of Marginals.** For all i and all $\mathcal{Q} \in \mathcal{Q}$, there exists $\pi \in \Pi$ such that $\pi_{\mathcal{Q}_i} = \delta(\mathcal{Q})$.

A model Π satisfying Assumptions 1 to 4 will be called **symmetrically ignorant quadratic**, or **s.i.q.** for short. Assumption 2 means that all voters' tops are known; therefore Assumptions 1 and 2 can be summarized as requiring $\Pi \subseteq \Pi_{\text{quad}}^\theta$ in our previous notation, where $\theta = (\theta_1, \dots, \theta_n)$ is the known profile of voters tops. Assumption 3 means that an individual's top (or any other observable individual characteristic) does not contain any information on the distribution of the individual's preferences by itself. Finally, Assumption 4 assumes in effect complete ignorance about each agent's \mathcal{Q}_i .

The plain quadratic model satisfies all assumptions except Symmetry. The 'regularizing' effect of the Symmetry assumption can be illustrated in Example 1.

Example 1 (cont.) Consider again the situation depicted in Fig. 1 above, but now suppose that the epistemic state of the social evaluator is described by a symmetric quadratic model Π rather than by the plain convex model. The minimal expected majority count for x against y is still 2, since it is evidently possible to find a symmetric prior such that all voters in U prefer y to x , i.e. $m_{\Pi}^-(x, y) = 2$. For example, one may

take the prior that assumes with certainty that all preferences are Euclidean. What about $m_{\Pi}^{-}(y, x)$? As before, one can assign quadratic forms $(\mathcal{Q}_1, \dots, \mathcal{Q}_n)$ to the tops such that all voters with top in U prefer x to y . But, as is evident from Fig. 1, these quadratic forms generally have to be distinct for different voters; the prior assuming this profile with certainty is therefore not symmetric. Any symmetric prior must thus be properly probabilistic; for example, a symmetric prior might assign equal probability $1/n!$ to each of the permutations of the profile $(\mathcal{Q}_1, \dots, \mathcal{Q}_n)$. But for any such prior the expected majority count for y against x will be at least 3, i.e. $m_{\Pi}^{-}(y, x) \geq 3$, as a key argument in the proof of our main result shows; for the geometric intuition behind this argument, see Figure 3 below. Hence, for any symmetric quadratic model Π we obtain $m_{\Pi}^{-}(y, x) > m_{\Pi}^{-}(x, y)$, and thus $yP_{\Pi}x$, where P_{Π} denotes the asymmetric part of the ex-ante majority relation R_{Π} ; in other words, x is not an ex-ante Condorcet winner.

At one extreme, there exists a unique largest (most imprecise) s.i.q. model consisting of *all* symmetric priors. Note that it does not impose any additional knowledge, i.e. probability one restrictions, beyond the plain quadratic model; it can thus be viewed as a regularized version of that model.

At the other extreme, there is also a unique smallest (most precise) s.i.q. model, as follows. Call a prior *uniform* if it puts all mass on profiles of the form $(\mathcal{Q}, \dots, \mathcal{Q})$ for some quadratic form \mathcal{Q} , and denote by Π_{unif} the **uniform (quadratic) model** consisting of all uniform priors; moreover, denote by Π_{exunif} the **extremal uniform model** consisting of all uniform priors of the form $\delta(\mathcal{Q}, \dots, \mathcal{Q})$, i.e. all priors that put unit mass on some single profile of the form $(\mathcal{Q}, \dots, \mathcal{Q})$. Evidently, the extremal uniform model satisfies Assumptions 1 to 4; conversely, combining Assumptions 3 and 4 also shows that any s.i.q. model contains the extremal uniform model.⁸

There is a wide range of intermediate specifications. For example, the quadratic forms \mathcal{Q}_i can be assumed to be drawn i.i.d. from some unknown distribution. In the subjectivist tradition, this is captured (and slightly generalized in the finite case) by assuming that Π consists of all exchangeable priors in the sense of de Finetti (1931).⁹ Finally, a s.i.q. model Π may also incorporate beliefs in possibly learnable correlations

⁸By the preceding observation, the class of all s.i.q. models forms a bounded lattice partially ordered by set inclusion.

⁹Specifically, in our context a prior π is *exchangeable* if, for all events $E \subseteq \mathcal{Q}_1 \times \dots \times \mathcal{Q}_n$ and all permutations σ of agents, $\pi(E) = \pi(\sigma(E))$, where $\sigma(E)$ is the event obtained from E by applying σ .

between the tops and the quadratic forms; it only excludes prior information about what these correlations are.

It turns out that the ex-ante Condorcet winners in the s.i.q. models are *Tukey medians* (Tukey, 1975) of a particular kind. For all $x \in X$, denote by \mathcal{H}_x the family of all Euclidean half-spaces that contain x (i.e. the family of all sets of the form $\{y \in X : a \cdot y \geq a \cdot x\}$ for some non-zero vector $a \in \mathbb{R}^L$). For all profiles $\theta = (\theta_1, \dots, \theta_n)$ and all half-spaces H , denote $\theta(H) := \#\{i : \theta_i \in H\}$, and define the *Tukey depth* of x at the profile θ by

$$\mathfrak{d}(x; \theta) := \min_{H \in \mathcal{H}_x} \theta(H).$$

Intuitively, the Tukey depth measures the ‘centrality’ of x with respect to the profile of tops: the larger $\mathfrak{d}(x; \theta)$ the more tops θ_i are guaranteed to lie in every direction viewed from x , and $\mathfrak{d}(x; \theta) = 0$ means that x can be separated from the entire set of tops θ by a hyperplane. Denote by $\mathfrak{d}^*(\theta) := \max_{x \in X} \mathfrak{d}(x; \theta)$ the maximal Tukey depth over X . The *Tukey median rule* selects, for every profile θ , the alternatives that attain this maximal depth:

$$T(\theta) := \arg \max_{x \in X} \mathfrak{d}(x; \theta) = \{x \in X \mid \mathfrak{d}(x; \theta) = \mathfrak{d}^*(\theta)\}.$$

Our main result involves the following refinement. For all profiles θ and all x , denote by $\mathcal{H}_x^* := \{H \ni x : \theta(H) = \mathfrak{d}^*(\theta)\}$. A Tukey median $x \in T(\theta)$ is *strict* if, for no $y \in T(\theta)$, $\mathcal{H}_y^* \subsetneq \mathcal{H}_x^*$. The set of **strict Tukey medians** is denoted by $T^*(\theta)$.

Theorem 1. *For all profiles θ and every symmetrically ignorant quadratic model $\Pi \subseteq \Pi_{\text{quad}}^\theta$, $\text{CW}(\Pi)$ is non-empty. Moreover,*

$$\text{CW}(\Pi) = T^*(\theta).$$

The proof of Theorem 1 (provided in the appendix) proceeds in a series of steps. First, it is shown that all s.i.q. models are equivalent. The argument relies crucially on both the symmetry assumption and the EAC solution concept. It allows to focus on the characterization of the analytically convenient uniform model. This simplifies matters greatly since the uniform model is characterized by strong ex-post restrictions on profiles. In particular, profiles of preferences with a common quadratic form are intermediate preferences in the sense of Grandmont (1978). More specifically, for any two alternatives x and y , the tops in a profile of preferences with a common quadratic

form \mathcal{Q} that prefer x to y are separated from those preferring y to x by a hyperplane through the midpoint between x and y , see Lemma A.2 in the appendix; Figure 2 shows these hyperplanes for selected common quadratic forms (for the identity matrix $\mathcal{Q} = \mathcal{I}$, the hyperplane is perpendicular to the straight line through x and y).




Figure 2: Intermediate preferences with separating hyperplane

From this one can show that the ex-ante majority relation of the uniform model coincides locally with the comparison of alternatives in terms of their *relative Tukey depth*: for all distinct $x, y \in X$, let

$$xR_{\delta}y \iff \min_{H \in \mathcal{H}_{x,y \notin H}} \theta(H) \geq \min_{H \in \mathcal{H}_{y,x \notin H}} \theta(H). \quad (4.2)$$

The ex-ante majority relation does not coincide globally with the relative Tukey depth relation (4.2) since the half-spaces separating the underlying tops in the quadratic model must go through the midpoint between x and y . Nonetheless, the set of local maxima of this relation is shown to coincide with the set of global maxima, which in turn coincides with the set of strict Tukey medians. Finally, the existence of strict Tukey medians is shown by an appeal to the Hausdorff maximal principle.

Example 1 (cont.) In Example 1, the Tukey depth of x relative to y is evidently $\min_{H \in \mathcal{H}_{x,y \notin H}} \theta(H) = 2$. Conversely, the Tukey depth of y relative to x is obtained by looking at the straight line ∂H through x and y : the tops that support y against x must at least contain the tops in $U \cap H$, or the tops in $U \cap H^c$. As can be inferred from Fig. 3, we therefore have $\min_{H \in \mathcal{H}_{y,x \notin H}} \theta(H) = 3$, and hence $yP_{\delta}x$, where P_{δ} denotes the asymmetric part of R_{δ} . It follows from the arguments provided in the proof of Theorem 1 in the appendix that we thus also obtain $yP_{\Pi}x$ for any s.i.q. model Π , as

claimed above.




Figure 3: The Tukey depth of y relative to x is equal to 3

To illustrate the relation between absolute and relative Tukey depth, consider the following example.

Example 2. Suppose that there are five voters whose tops $\theta_1, \dots, \theta_5$ form a pentagon as shown in Figure 4. The (strict) Tukey median (and hence by Theorem 1 also the ex-ante Condorcet winners of any s.i.q. model) is given by the points in the inner convex pentagon marked in red.¹⁰ Fig. 4 also shows a point y and its associated upper contour set with respect to the relative Tukey depth relation given by (4.2) in blue color. Note in particular that the points x and y have the same (absolute) Tukey depth but different relative depth, to wit $xP_\delta y$.




Figure 4: Absolute versus relative Tukey depth

¹⁰This can be verified from the following observations. First, any line passing through the inner red pentagon has at least two tops on either side; on the other hand, for any point outside the inner pentagon there is a Euclidean half-space containing that point and at most one top. In particular, the maximal Tukey depth is $\mathfrak{d}^*(\theta) = 2$; all points in the convex hull of the tops that are not in this inner pentagon have depth one, and all points outside the convex hull of the tops have depth zero.

While every Tukey median is strict in Example 2, it is an open question if this is the case generally. It must be the case whenever Tukey medians are unique (because strict Tukey medians always exist). Demange (1982) has in fact shown such uniqueness whenever voters' tops are continuously distributed with a convex support. Using this result, we obtain the following 'continuous' version of Theorem 1.

Theorem 1' *Suppose that voters' tops are distributed according to a continuous measure θ with convex support. Then, the strict Tukey median set $T^*(\theta)$ consists of a single point, and for every s.i.q. model $\Pi \subseteq \Pi_{\text{quad}}^\theta$,*

$$\text{CW}(\Pi) = T(\theta) = T^*(\theta).$$

(Proof in appendix.)

5 Uncertainty about Tops

We have so far assumed that the only individuating information about individual preferences concerns their tops, and that this information is perfect (tops assumed to be known by the social evaluator). We now extend the frugal aggregation approach maintaining the first assumption while abandoning the second. Formally, we now assume that the evaluator has a precise prior over the top of each voter. Such models may be of interest when individual tops are elicited by a vote or a poll, and when there are doubts whether they should be taken at face value, for instance for incentive reasons. Obviously, for concrete applications this needs to be developed further by specifying how the evaluator's probabilistic beliefs over tops are themselves formed.

Specifically, we adapt our assumptions on the epistemic state Π of the evaluator as follows.

1. **Concentration on Quadratic Preferences.** $\Pi \subseteq \Pi_{\text{quad}}$.
- 2a. **Tops Probabilism.** For all $\pi, \pi' \in \Pi$ and all i , $\pi_{X_i} = \pi'_{X_i} =: \mu_i$.
- 2b. **Independence.** For all i and all $\theta_i, \theta'_i \in \text{supp } \mu_i$, $\pi_{\mathcal{Q}_{\theta_i}} = \pi_{\mathcal{Q}_{\theta'_i}}$.
3. **Symmetry.** For all $\pi \in \Pi$ and all i, j , $\pi_{\mathcal{Q}_i} = \pi_{\mathcal{Q}_j}$.
4. **Complete Ignorance of Marginals.** For all i , $\theta_i \in \text{supp } \mu_i$ and all $\mathcal{Q} \in \mathcal{Q}$, there exists $\pi \in \Pi$ such that $\pi_{\mathcal{Q}_{\theta_i}} = \delta(\mathcal{Q})$.

Assumption 2a says that all priors in Π agree on the distribution of tops, i.e. the uncertainty about tops is probabilistic rather than imprecise. Assumption 2b adds that any top θ_i in the support of the marginal top distribution μ_i induces the same marginal distribution $\pi_{\mathcal{Q}_{\theta_i}}$ over quadratic forms. The marginal top distributions can take the form of finite or continuous measures; in the latter case, in order to apply Theorem 1', we need to assume that all μ_i have a common convex support.¹¹ Denoting by $\mu = (\mu_1, \dots, \mu_n)$ the profile of the marginal distributions over tops, and by Π_{quad}^μ the set of all quadratic priors that induce the marginal distribution profile μ , we can summarize Assumptions 1 and 2a by requiring $\Pi \subseteq \Pi_{\text{quad}}^\mu$. With slight abuse of terminology, we continue calling a model satisfying these modified assumptions *symmetrically ignorant quadratic (s.i.q.)* since no confusion can arise.

To adapt our main result to the situation in which the social evaluator is uncertain about the voters' tops but has a unique prior μ over the profile of the distribution of tops, denote by $\bar{\mu}$ the average distribution of tops defined by

$$\bar{\mu} := \sum_{i=1}^n \frac{1}{n} \cdot \mu_i.$$

Associate with each i an 'ex-ante subpopulation' with distribution of tops μ_i and relative size $1/n$; these combine to a total ex-ante population with distribution of tops $\bar{\mu}$ and quadratic forms still unknown as in Theorem 1. Independence (Assumption 2b) ensures symmetry within each subpopulation, while Symmetry (Assumption 3) ensures symmetry across subpopulations. Therefore, the argument of Theorem 1 applies and yields the strict Tukey median with respect to $\bar{\mu}$ as the ex-ante Condorcet winners; note that, in the following result, Tops Probabilism (Assumption 2a) is indispensable as it is necessary to even define the characterized set $T^*(\bar{\mu})$.

Theorem 2. *For all profiles $\mu = (\mu_1, \dots, \mu_n)$ such that the μ_i are either finite or continuously distributed with a common convex support, and for every symmetrically ignorant quadratic model $\Pi \subseteq \Pi_{\text{quad}}^\mu$, $\text{CW}(\Pi)$ is non-empty. Moreover,*

$$\text{CW}(\Pi) = T^*(\bar{\mu}).$$

(Proof in appendix.)

¹¹It might be possible to generalize the result to arbitrary or arbitrary continuous probability distributions, but this would require additional arguments.

6 Discussion

How robust is our proposed solution concept, the (strict) Tukey median, with respect to the precise epistemic assumptions of the symmetrically ignorant quadratic model? We now argue that the solution is indeed remarkably robust with respect to assuming *less* knowledge, but not necessarily with respect to assuming *more* knowledge about the underlying preferences. The latter should, of course, come as no surprise, since additional available information generally has to be accounted for in the search for welfare optima.

6.1 Less Informative Beliefs: Hedging Quadraticity

The argument for the Tukey median as ex-ante Condorcet winner relies crucially on the assumption of quadratic preferences and symmetrically ignorant beliefs. But if these assumptions cannot be taken for granted, it seems sensible for the social evaluator to hedge the commitment to the s.i.q. model by mixing it with the plain convex model. Formally, assume that the epistemic state of the social evaluator is described by a ‘mixture’ of models, as follows.

For $\beta \in [0, 1]$, define the *mixture* of the models Π and Π' by

$$\beta\Pi + (1 - \beta)\Pi' := \{\beta\pi + (1 - \beta)\pi' \mid \pi \in \Pi \text{ and } \pi' \in \Pi'\}.$$

Concretely, consider a distribution θ of tops, any s.i.q. model $\Pi \subseteq \Pi_{\text{quad}}^\theta$ and the mixture $\beta\Pi + (1 - \beta)\Pi_{\text{co}}^\theta$. With continuously distributed tops, the effect of this mixing on the outcome selection is clearcut and striking: as long as $\beta > 0$, there is none, i.e. the Tukey median continues to be the normative optimum!

Proposition 4. *Suppose that θ is continuously distributed with convex support, and let $\Pi \subseteq \Pi_{\text{quad}}^\theta$ be any symmetrically ignorant quadratic model. Then, for all $\beta > 0$,*

$$\text{CW}(\beta\Pi + (1 - \beta)\Pi_{\text{co}}^\theta) = T(\theta).$$

To see this, consider two distinct alternatives x and y . The tops of voters who prefer x to y with probability one under the plain convex model are all located on the line through x and y and therefore have mass zero under a continuous distribution;

in other words $m_{\Pi_{\text{co}}^{\theta}}^{-}(x, y) = 0$. Thus,

$$\begin{aligned} m_{\beta\Pi+(1-\beta)\Pi_{\text{co}}^{\theta}}^{-}(x, y) &= \beta \cdot m_{\Pi}^{-}(x, y) + (1 - \beta) \cdot m_{\Pi_{\text{co}}^{\theta}}^{-}(x, y) \\ &= \beta \cdot m_{\Pi}^{-}(x, y). \end{aligned} \tag{6.1}$$

By (6.1), the mixed model $\beta\Pi+(1-\beta)\Pi_{\text{co}}^{\theta}$ induces the same ex-ante majority relation as the s.i.q model Π , hence also the same ex-ante Condorcet winner, for all $\beta > 0$.

In the finite case, there is no exact counterpart to Proposition 4. Indeed, an ex-ante Condorcet winner might well fail to exist in the mixed model. But one can still expect appropriate Condorcet extensions to deliver an approximate version of the Tukey median rule in the finite case since the underlying rationale – the large size of the expected majority intervals under plain convexity – is transparent and robust.

6.2 Informative Beliefs about Marginals

Another possible issue in applying the Tukey median in a particular setting is based on a concern that it may rest on too little information. We have already discussed that Symmetry (Assumption 3) accommodates a wide range of beliefs about the correlation of quadratic forms across voters. So here we ask about the possible implication of relaxing Complete Ignorance of Marginals (Assumption 4) by considering sets of priors Π such that the induced set of marginals over quadratic forms is strictly contained in $\Delta(\mathcal{Q})$.

By a preliminary consideration, note that, even if the set of priors can be described simply and plausibly, it may be very difficult to characterize the induced expected majority intervals analytically or computationally; furthermore, an ex-ante Condorcet winner may well not exist, and some appropriate (ex-ante) Condorcet extension would need to be applied. So even if the Complete Ignorance of Marginals assumption is seen as too agnostic entailing a misspecification in the set of priors, the Tukey median may retain much of its appeal as a useful informationally conservative and tractable ‘approximation’ or ‘stand-in’ for the intractable normative optimum corresponding to the actual set of beliefs.

Setting tractability aside, to get a better grip on the possible opportunity loss of unused information, focus on the polar opposite of the Complete Ignorance of Marginals condition above, namely certain knowledge of the individual quadratic forms \mathcal{Q}_i which, by Symmetry, must then coincide with some common quadratic

form \mathcal{Q} . This is a limiting case of our frugal aggregation framework in which the top reveals the entire preference ordering; in fact, the quest for a frugal optimum boils down to a question of standard ordinal aggregation of complete preferences on a restricted domain. If \mathcal{Q} is the unit matrix, we are in the classical spatial model in which preferences are assumed to be Euclidean. (Note that the aggregation problems for general \mathcal{Q} can be reduced to a Euclidean aggregation problem by a change of coordinates via an appropriate affine transformation of the space of alternatives).

In the case of a known common quadratic form \mathcal{Q} , welfare optima are naturally obtained as the maxima of the program

$$\arg \max_{x \in X} \sum_{i=1}^n f(u_i(x)), \quad (6.2)$$

where f is a common transformation and the $u_i(\cdot)$ are given as in (4.1) with the common quadratic form \mathcal{Q} . The common transform f can be pinned down naturally by appeal to the Condorcet principle which selects Condorcet winners whenever they exist. While it is well-known that their set is generically empty in this setting (McKelvey, 1979), they do exist if all tops are collinear in which case the Condorcet winner coincides with the standard median on a line. This forces f to be the square root function. In the case of Euclidean preferences, this means that the welfare optima minimize the sum of the Euclidean distances to the tops. In general, the utilitarian welfare optimum (6.2) is given by the ‘geometric median’ with respect to the quadratic form \mathcal{Q} . Concretely, for all profiles θ and all quadratic forms \mathcal{Q} , let

$$\text{Med}_{\mathcal{Q}}(\theta) := \arg \max_{x \in X} \sum_{i=1}^n -\sqrt{(x - \theta_i)^T \cdot \mathcal{Q} \cdot (x - \theta_i)}. \quad (6.3)$$

which we refer to as the *geometric \mathcal{Q} -median*. The geometric median is another classic multi-dimensional median; see, e.g., Vardi and Zhang (2000) for its basic properties.

With the geometric medians as a normative yardstick under full knowledge, one can make the question of opportunity loss more precise by comparing the Tukey and geometric medians. It is easy to see that, without any restrictions on the profile of tops, different \mathcal{Q} -medians may be far apart, and so may be a \mathcal{Q} -median and the Tukey median.

On the other hand, if the tops are sufficiently ‘bunched together,’ there may be

fairly tight restrictions on how far apart the two medians can be. Indeed, there are strong results on the relationship between the Tukey median and the coordinate-wise *mean* in such settings. (Note that the mean is the welfare optimum under a different ordinal transform f , namely the identity.) Specifically, Caplin and Nalebuff (1988, 1991) show that if the distribution of tops has a log-concave density, the mean has Tukey depth of at least $1/e$; they also argue that the set of points with this property is ‘small,’ hence that the mean and the Tukey median must be close together.

It is an interesting – and challenging – question for future research to determine if a similar result holds for the geometric median. If it did, it would provide additional support for the Tukey median as a frugal aggregator, inasmuch as it would bound the potential loss from underspecification tightly.¹²

The literature has approached the spatial model as an instance of general-purpose ordinal aggregation rules applied to a specific domain of profiles, focusing on different standard Condorcet extension rules. Most prominent among them is the min-max (‘Simpson-Kramer’) solution, see Kramer (1977); Demange (1982); Caplin and Nalebuff (1988). Remarkably, the minmax solution under Euclidean preferences coincides with the Tukey median, and this equality generalizes to all uniform profiles of quadratic preferences. By consequence, the minmax rule ignores the non-top preference information entirely (even though it is available) and thus fails to exploit the metric structure of Euclidean resp. uniformly quadratic preference profiles. By contrast, the Tukey median as ex-ante Condorcet solution in the s.i.q. models cannot be criticized for failing to use such information, since this information is not available in these models ex hypothesis.

6.3 The Tukey Median as a Voting Mechanism?

In this paper, we have studied the (strict) Tukey median as a criterion of normative evaluation under restricted (‘frugal’) information. A different though not unrelated question concerns the suitability of the Tukey median as a voting mechanism in which voters choose ‘top’ messages in a self-interested fashion.

An initial question is whether honesty is consistently in voters’ self-interest, and

¹²In support of this possibility, the Tukey median is more similar to the geometric median than to the mean as the former agree exactly on collinear profiles, while the latter does not. On the other hand, the Tukey median shares with the mean the feature of being equivariant under affine transformation which the geometric median does not.

the unsurprising answer is that it is not, and cannot be. In a variation of the classical impossibility result on strategy-proof social choice, this has been as shown for rich domains of convex preferences (including rich domains of quadratic preferences) in particular by Zhou (1991).

So strategy-proofness is a moot issue. Beyond signaling the value of future research, with the following remarks we aim to provide a few indications why the Tukey median holds promise also as a parsimonious voting mechanism employed by self-interested voters.¹³ Specifically, we suggest the following two notable points. Voters will frequently have an incentive to ‘manipulate’ (i.e. to depart from their true top). But in some cases, such manipulations are arguably permissible, indeed even desirable from an impartial social-evaluator point of view; in that sense, strategy-proofness is not even an appropriate ideal in a frugal setting. In other cases, the impartial evaluation of manipulation might be more ambiguous, but the risk to the interests of other voters remains limited using the Tukey mechanism.

Understood as a game form, a voting rule now takes messages as inputs not the true tops. (The latter, now denoted by $\bar{\theta}_i$, are not directly observed and remain in the background). At a given profile of other voters’ messages θ_{-i} , call a message by i ‘consonant’ with that profile if it enhances the Tukey depth of the profile (then necessarily by 1), and ‘dissonant’ if it does not. Formally, θ_i is *consonant* at θ_{-i} if $\mathfrak{d}^*(\theta_i, \theta_{-i}) = \mathfrak{d}^*(\theta_{-i}) + 1$; and it is *dissonant* if $\mathfrak{d}^*(\theta_i, \theta_{-i}) = \mathfrak{d}^*(\theta_{-i})$. Moreover, for $m \geq 1$ denote by

$$T^{[-m]}(\theta) := \{x \in X \mid \mathfrak{d}(x; \theta) \geq \mathfrak{d}^*(\theta) - m\}.$$

Fact 6.1. (i) If θ_i is consonant at θ_{-i} , $T(\theta_i, \theta_{-i}) \subseteq T(\theta_{-i})$.

(ii) If $\theta_i \in T(\theta_{-i})$, then θ_i is consonant at θ_{-i} and $\theta_i \in T(\theta_i, \theta_{-i})$.

(iii) If θ_i is dissonant at θ_{-i} , $T(\theta_{-i}) \subseteq T(\theta_i, \theta_{-i}) \subseteq T^{[-1]}(\theta_{-i})$

To give a normative angle to these properties of the Tukey median rule as a mechanism, based on the main argument of the paper, take $T(\cdot)$ as normatively

¹³In the following remarks, we focus for simplicity on the non-strict Tukey median. We also allow the mechanism being defined as a correspondence, leaving the ultimate selection of an alternative unspecified. This corresponds to a valuable strand in the strategy-proofness literature starting with Kelly (1977); Gärdenfors (1979); a fuller analysis would presumably want to consider stochastic or deterministic single-valued selections as well.

optimal if all voters messages are truthful, and focus on a potential manipulation by voter i , assuming all others' messages to be truthful. If i 's message is consonant, then, according to Fact 6.1(i), the final choice set will select from within $T(\theta_{-i})$; by hypothesis, the elements of $T(\theta_{-i})$ are equally optimal (from the impartial point of view of the social evaluator) for the subset of voters other than i . So, after restricting the social choice to $T(\theta_{-i})$, i 's interests are arguably aligned with those of the social evaluator. Obtaining a maximal element (or subset) in $T(\theta_{-i})$ may require a departure from i 's true top, but why should such manipulation bother the social evaluator?¹⁴

On the other hand, if i 's message is dissonant, then some of the chosen alternatives may now conflict with the aggregate interests of the other voters as represented by $T(\theta_{-i})$. However, this departure is minimal since the depth is still 'almost' maximized. Note also that, for any θ , the sets $\{T^{[-m]}(\theta)\}_{m \geq 1}$ form a nested family of closed convex sets. Unless the number of voters n is small, the elements of $T^{[-1]}(\theta)$ will thus stay 'close' to $T(\theta)$. In this sense, the risk of manipulation by any single voter stays bounded. Contrast this with other tops-only mechanisms, such as the mean rule on \mathbb{R}^L . At any profile θ_{-i} , under the mean rule voter i can achieve his true top $\bar{\theta}_i$ by selecting $\theta_i = n\bar{\theta}_i - \sum_{j \neq i} \theta_j$, overriding the input of the other voters completely.

One can extend the preceding bounded risk argument to coalitional manipulations by noting that, for any subset of voters J ,

$$T(\theta_J, \theta_{-J}) \supseteq T^{[-|J|]}(\theta_J).$$

So to completely override the voters outside J by moving outside the convex hull of their tops, it takes a coalition of size at least $\mathfrak{d}^*(\theta)$.¹⁵ Such coalitions need to be sizeable since by a fundamental result (Donoho and Gasko, 1992, Prop. 2.3), for any profile of tops θ ,

$$\mathfrak{d}^*(\theta) \geq \frac{n}{L+1}.$$

We conclude with two illustrating examples.

¹⁴At first blush, it might seem that consonant messages are likely to be the exception, and dissonant ones the rule. Not so! Take any profile θ with n voters, choose a number $m \leq n$ with uniform probability, a random subset $J \subseteq \{1, \dots, n\}$ of size m and a random voter $j \in J$. Then θ_j is consonant with θ_{J-j} with probability equal to $\mathfrak{d}^*(\theta)$, as can be established by a double counting argument.

¹⁵These considerations are formally analogous to the robustness analysis of multi-dimensional statistical estimators in terms of 'breakpoints,' see Donoho and Gasko (1992). The concern with outliers in the statistical setting translates into a concern with coalitional manipulations here.

Example 3. Suppose the (true) tops of 4 voters are in general position as, for instance, in the profile displayed in Figure 5.




Figure 5: Four points in general position

Here, no top is in the convex hull of the others, and the line segments $[\bar{\theta}_1, \bar{\theta}_3]$ and $[\bar{\theta}_2, \bar{\theta}_4]$ intersect at z . Then $\mathfrak{d}^*(\bar{\theta}) = 2$. Note that, for any voter, any message is consonant; moreover, any voter i can obtain any of the alternatives in $\text{co}\{\bar{\theta}_{-i}\}$ as a singleton. For instance, voter 1 can obtain any $y \in \text{co}\{\bar{\theta}_2, \bar{\theta}_3, \bar{\theta}_4\}$ as a singleton by choosing $\theta_1 = y$. With convex preferences, the preference maximum in $\text{co}\{\bar{\theta}_2, \bar{\theta}_3, \bar{\theta}_4\}$ lies on the line segment $[\bar{\theta}_2, \bar{\theta}_4]$, yet it would be entirely accidental if it was equal to z . So voter 1 would typically have an incentive to manipulate; but, as argued above, the evaluator has good reason to concur since voter 1 reveals useful private information.¹⁶ The thrust of our argument here is simply to counter the common intuition that manipulations are generally detrimental.

Example 4. Now consider a profile with three voters. Again, suppose that the true tops are in general position (see Figure 6), and consider the options of voter 1.

Any message θ_1 outside the line through $\bar{\theta}_2$ and $\bar{\theta}_3$ is dissonant, and entails the choice of $\text{co}\{\theta_1, \bar{\theta}_2, \bar{\theta}_3\}$. It is not clear how voter 1 would profit from a dissonant departure from the truth. Voter 1 may have a stronger motive for choosing a consonant message in the segment $[\bar{\theta}_2, \bar{\theta}_3]$, depending on his actual preferences. While she forfeits the chance at his top alternative, she avoids the risk of obtaining the worst in $\text{co}\{\bar{\theta}_1, \bar{\theta}_2, \bar{\theta}_3\}$ which, by preference convexity, is either $\bar{\theta}_2$ or $\bar{\theta}_3$. Again, there is no evident reason for the social evaluator to object.

¹⁶Evidently, the same reasoning applies to each of the four voters, so all of them might manipulate. Predicting what will happen is a matter of game theoretic analysis which is notoriously difficult in the context of voting.




Figure 6: Three points in general position

Evidently, a detailed study of the Tukey median as a voting mechanism is a rich topic far beyond the scope of this paper. The above observations are meant to suggest the promise of that study, and of the Tukey median as a voting mechanism.

References

- ARROW, K. J. (1951/63). *Social Choice and Individual Values*. New York: Wiley.
- (1960). Decision theory and the choice of a level of significance for the t -test. In I. Olkin, S. G. Churye, W. Hoeffding, W. Madow and H. B. Mann (eds.), *Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling Theories*, Stanford University Press, pp. 70–78.
- AUSTEN-SMITH, D. and BANKS, J. (1999). *Positive Political Theory I: Collective Preference*. Ann Arbor: Michigan University Press.
- AZIZ, H. and SHAH, N. (2020). Participatory budgeting: Models and approaches. In Rudas and Gabor (eds.), *Pathways between Social Sciences and Computational Social Science: Theories Methods and Interpretations*, Springer.
- BLACK, D. S. (1948). On the rationale of group decision-making. *J. Political Economy*, **56**, 23–34.
- BOUTILIER, C. and ROSENSCHEIN, J. S. (2016). Incomplete information and communication in voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang and A. Procaccia (eds.), *Handbook of Computational Social Choice*, 10, Cambridge: Cambridge University Press, pp. 223–257.
- CAPLIN, A. and NALEBUFF, B. (1988). On 64% majority rule. *Econometrica*, **56**, 787–814.
- and — (1991). Aggregation and social choice: A mean voter theorem. *Econometrica*, **59**, 1–23.
- DE FINETTI, B. (1931). Funzione caratteristica di un fenomeno aleatoria. *Atti della R. Accademia Nazionale dei Lincei, Serie 6. Memorie, Classe die Scienze Fisiche, Matematiche e Naturale*, **4**, 251–299.
- DEMANGE, G. (1982). A limit theorem on the minmax set. *Journal of Mathematical Economics*, **9**, 145–164.
- DONOHO, D. L. and GASKO, M. (1992). Breakdown properties of location estimates based on halfspace depth and projected outlyingness. *Annals of Statistics*, **20**, 1803–1827.

- DOWNES, A. (1957). *An Economic Theory of Democracy*. New York: Harper.
- GÄRDENFORS, P. (1979). On definitions of manipulation of social choice functions. In J.-J. Laffont (ed.), *Aggregation and Revelation of Preferences*, Amsterdam: North-Holland.
- GIGERENZER, G. and GOLDSTEIN, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, **104**, 650–669.
- GRANDMONT, J.-M. (1978). Intermediate preferences and the majority rule. *Econometrica*, **46**, 317–330.
- KALAI, E., MULLER, E. and SATTERTHWAITTE, M. A. (1979). Social welfare functions when preferences are convex, strictly monotonic, and continuous. *Public Choice*, **34**, 87–97.
- KELLY, J. S. (1977). Strategy-proofness and social choice functions without single-valuedness. *Econometrica*, **45**, 439–446.
- KONCZAK, K. and LANG, J. (2005). Voting procedures with incomplete preferences. In *Proceedings of the Multidisciplinary IJCAI-05 Workshop on Advances in Preference Handling*, Palo Alto, CA: AAAI, pp. 124–129.
- KRAMER, G. H. (1977). A dynamical model of political equilibrium. *Journal of Economic Theory*, **16**, 310–334.
- LANCASTER, K. J. (1966). A new approach to consumer theory. *Journal of Political Economy*, **74**, 132–157.
- LE BRETON, M. and WEYMARK, J. A. (2011). Arrovian social choice theory on economic domains. In K. J. Arrow, A. Sen and K. Suzumura (eds.), *Handbook of Social Choice and Welfare, Volume 2*, Amsterdam: North-Holland, pp. 191–299.
- LU, T. and BOUTILIER, C. (2011). Robust approximation and incremental elicitation in voting protocols. In *Proceedings of the 22nd IJCAI*, Palo Alto, CA: AAAI, pp. 287–293.
- LUCE, R. and RAIFFA, H. (1957). *Games and Decisions*. New York: Wiley.

- MCKELVEY, R. (1979). General conditions for global intransitivities in formal voting models. *Econometrica*, **47** (5), 1085–1112.
- MILNOR, J. (1954). Games against nature. In R. M. Thrall, C. H. Coombs and R. L. Davis (eds.), *Decision Processes*, Wiley, pp. 49–59.
- NEHRING, K. (2000). A theory of rational choice under ignorance. *Theory and Decision*, **48**, 205–240.
- (2009). Imprecise probabilistic beliefs as a context for decision-making under ambiguity. *Journal of Economic Theory*, **144**, 1054–1091.
- and PUPPE, C. (2022). Condorcet solutions in frugal models of budget allocation. *KIT Working Paper Series in Economics # 156*.
- PLOTT, C. (1967). A notion of equilibrium and its possibility under majority rule. *American Economic Review*, **57**, 787–806.
- ROUSSEEUW, P. J. and HUBERT, M. (2017). Computation of robust statistics: Depth, median, and related measures. In *Handbook of Discrete and Computational Geometry*, Boca Raton: CRC Press.
- SHAH, A. (2007). *Participatory Budgeting*. Washington D.C.: The World Bank.
- SMALL, C. G. (1990). A survey of multidimensional medians. *International Statistical Review*, **58**, 263–277.
- TUKEY, J. W. (1975). Mathematics and the picturing of data. In *Proceedings of the ICM*, Vancouver: Canadian Log Builder Assoc., pp. 523–531.
- VARDI, Y. and ZHANG, C.-H. (2000). The multivariate l_1 -median and associated data depth. *Proc. Nat. Acad. Sci.*, **97**, 1423–1426.
- ZHOU, L. (1991). Impossibility of strategy-proof mechanisms for economies with pure public goods. *Review of Economic Studies*, **58**, 107–119.

Appendix A: Proofs

Proof of Proposition 1. Given the pair $x, y \in X$, let $\Pi' \subseteq \Pi$ be the subset of all priors $\pi \in \Pi$ such that, for all $i = 1, \dots, n$, $\pi(x \sim_i y) = 0$. By the regularity assumption, we have $m_{\Pi}^-(x, y) = m_{\Pi'}^-(x, y)$ and $m_{\Pi}^-(y, x) = m_{\Pi'}^-(y, x)$, and therefore also $m_{\Pi}^+(x, y) = m_{\Pi'}^+(x, y)$ and $m_{\Pi}^+(y, x) = m_{\Pi'}^+(y, x)$. By construction, we have for all $\pi \in \Pi'$,

$$m_{\pi}(x, y) + m_{\pi}(y, x) = n.$$

This implies

$$\begin{aligned} m_{\Pi'}^-(x, y) + m_{\Pi'}^+(y, x) &= n \text{ and} \\ m_{\Pi'}^-(y, x) + m_{\Pi'}^+(x, y) &= n, \end{aligned}$$

therefore

$$m_{\Pi'}^-(x, y) + m_{\Pi'}^+(y, x) = m_{\Pi'}^-(y, x) + m_{\Pi'}^+(x, y),$$

and hence (2.3). \square

For the following proofs, the following observation will be useful. Denote by Π_{exco} the ‘extremal’ convex model, i.e. the submodel of the plain convex model which only contains sets of priors that put all mass on a *single* profile of convex preferences.

Lemma A.1. *The two models Π_{exco} and Π_{co} are equivalent.*

Proof. Consider any pair of distinct alternatives $x, y \in X$. Let π^0 be a minimizer of the support count for x against y under the model Π_{co} , i.e. $m_{\Pi_{\text{co}}}^-(x, y) = m_{\pi^0}(x, y)$. Furthermore, let \succsim^0 be a profile of convex preferences in the support of π^0 such that $\#\{i : x \succ_i y\}$ is minimal among all profiles in the support of π^0 . Let $\delta(\succsim^0)$ be the prior that puts all mass on \succsim^0 ; then, $m_{\delta(\succsim^0)}(x, y) \leq m_{\pi^0}(x, y)$. But since $\delta(\succsim^0)$ is an admissible prior under the model Π_{co} , we have in fact $m_{\delta(\succsim^0)}(x, y) = m_{\pi^0}(x, y)$; this implies the desired result. \square

Proof of Proposition 2. Using Lemma A.1, the proof is straightforward from well-known properties of single-peaked preferences. \square

Proof of Proposition 3. Again using Lemma A.1, it is sufficient to prove the statement for the model $\Pi_{\text{exco}} \subseteq \Pi_{\text{co}}^{\theta}$. Consider any $x \neq \theta_{i^*}$; by assumption, there is at most one

$\theta_j \neq \theta_{i^*}$ on the straight line through x and θ_{i^*} . Since preference convexity entails no restriction in the comparison of x and θ_{i^*} for tops outside that straight line, and since θ_{i^*} has largest popular support, this implies $m_{\Pi_{\text{exco}}}^-(\theta_{i^*}, x) \geq m_{\Pi_{\text{exco}}}^-(x, \theta_{i^*})$, i.e. θ_{i^*} is an ex-ante majority winner against x ; if θ_{i^*} has uniquely largest popular support, we even have $m_{\Pi_{\text{exco}}}^-(\theta_{i^*}, x) > m_{\Pi_{\text{exco}}}^-(x, \theta_{i^*})$. Since x was chosen arbitrarily, the result follows. \square

Proof of Theorem 1

The proof of Theorem 1 is given by means of a series of auxiliary results. First, Proposition 5 shows that all s.i.q. models are equivalent to the uniform quadratic model hence, by an argument analogous to that given in the proof of Lemma A.1, also to the extremal uniform model Π_{exunif} . The main subsequent steps are summarized in two further propositions: Proposition 6 shows that the ex-ante Condorcet winners of the extremal uniform model coincide with the maximizers of the relative Tukey depth. Finally, Proposition 7 demonstrates that the set of maximizers of the relative Tukey depth is non-empty and coincides with the strict Tukey median.

A key fact about the s.i.q. models is that they are all equivalent; specifically, we have the following result.

Proposition 5. *All symmetrically ignorant quadratic models are equivalent.*

Proof. We show that any s.i.q. model Π is equivalent to the uniform quadratic model Π_{unif} . Consider a fixed pair $x, y \in X$ of distinct alternatives, and a fixed profile θ of tops. Let π be any symmetric prior and consider any fixed voter $h = 1, \dots, n$. Denote by $\tilde{\pi} \in \Pi_{\text{unif}}$ be the unique prior that is concentrated on uniform profiles and satisfies $\tilde{\pi}_{\mathcal{Q}_h} = \pi_{\mathcal{Q}_h}$. By symmetry of π , we have $\pi_{\mathcal{Q}_i} = \pi_{\mathcal{Q}_h}$ for all $i = 1, \dots, n$, and by construction, $\tilde{\pi}_{\mathcal{Q}_i} = \tilde{\pi}_{\mathcal{Q}_h}$ for all $i = 1, \dots, n$; hence, $\tilde{\pi}_{\mathcal{Q}_i} = \pi_{\mathcal{Q}_h}$ for all $i = 1, \dots, n$. This implies

$$\begin{aligned} m_{\tilde{\pi}}(x, y) &= E_{\tilde{\pi}}[\#\{i : x \succ_i y\}] = \sum_{i=1}^n E_{\tilde{\pi}_{\mathcal{Q}_i}}[x \succ_i y] = \\ &= \sum_{i=1}^n E_{\pi_{\mathcal{Q}_i}}[x \succ_i y] = E_{\pi}[\#\{i : x \succ_i y\}] = m_{\pi}(x, y). \end{aligned}$$

In other words, for every prior $\pi \in \Pi$ there exists a uniform prior $\tilde{\pi} \in \Pi_{\text{unif}}$ that induces the same expected majority count for x against y . This implies

$$m_{\Pi_{\text{unif}}}^-(x, y) \leq m_{\Pi}^-(x, y). \quad (\text{A.1})$$

On the other hand, by Assumptions 3 (Symmetry) and 4 (Complete Ignorance of Marginals), every s.i.q. model contains the extremal uniform model Π_{exunif} , hence

$$m_{\Pi}^-(x, y) \leq m_{\Pi_{\text{exunif}}}^-(x, y). \quad (\text{A.2})$$

Finally, by an argument completely analogous to the argument in the proof of Lemma A.1, we have

$$m_{\Pi_{\text{exunif}}}^-(x, y) = m_{\Pi_{\text{unif}}}^-(x, y). \quad (\text{A.3})$$

Combining (A.1), (A.2) and (A.3), we obtain that the arbitrary s.i.q. model Π induces the same intervals of expected majority counts as the uniform quadratic model Π_{unif} .

□

Denote the *relative Tukey depth of x with respect to y* by

$$\mathfrak{d}(x, y; \theta) := \min_{H \in \mathcal{H}_x, y \notin H} \theta(H),$$

so that $x R_{\mathfrak{d}} y \Leftrightarrow \mathfrak{d}(x, y; \theta) \geq \mathfrak{d}(y, x; \theta)$ (cf. (4.2)), as well as

$$S(\theta) := \{x \in X \mid \text{for no } y \in X, y P_{\mathfrak{d}} x\}.$$

Due to Proposition 5 and Lemma A.1, we can concentrate in the remainder of the proof of Theorem 1 on the extremal uniform model Π_{exunif} (with the fixed top profile θ). Our goal is to prove the following result.

Proposition 6. *For all profiles θ and $\Pi_{\text{exunif}} \subseteq \Pi_{\text{quad}}^{\theta}$,*

$$\text{CW}(\Pi_{\text{exunif}}) = S(\theta).$$

One difficulty in showing this is that the ex-ante majority relation of the extremal uniform model does in fact *not* coincide with the relative Tukey depth, as noted in Example 2 above. Nevertheless, their maximal elements coincide.

Again, we need a preliminary result. Observe that, since a preference is quadratic if and only if it is obtained from a Euclidean preference (with circles as indifference curves) by an affine transformation, we have the following result.

Lemma A.2. *Let $x, y \in X$ be any two distinct alternatives, and $\succ = (\succ_1, \dots, \succ_n)$ a uniform profile of quadratic preferences with tops $\theta = (\theta_1, \dots, \theta_n)$. Then, there exists a (Euclidean) half-space $H \subseteq \mathbb{R}^L$ such that the hyperplane ∂H passes through the midpoint between x and y , and*

$$\{\theta_i \mid x \succ_i y\} \subseteq \text{int}(H) \quad \text{and} \quad \{\theta_i \mid y \succ_i x\} \subseteq \text{int}(H^c), \quad (\text{A.4})$$

where H^c is the complement of H in \mathbb{R}^L . Conversely, for any (Euclidean) half-space H that separates x from y such that ∂H passes through the midpoint between x and y , there exists a uniform profile of quadratic preferences that satisfies (A.4).

Proof of Proposition 6. Let $x^* \in \text{CW}(\Pi_{\text{exunif}})$, i.e. $x^* R_{\Pi_{\text{exunif}}} y$ for all $y \in X$. By contradiction, assume that $x^* \notin S(\theta)$. Then, $y P_{\theta} x^*$ for some $y \in X$, i.e.

$$\mathfrak{d}(x^*, y; \theta) < \mathfrak{d}(y, x^*; \theta). \quad (\text{A.5})$$

Let $H^0 \in \mathcal{H}_{x^*}$ be a Euclidean half-space that separates x^* from y and that minimizes the measure $\theta(H)$ among all such half-spaces. Without loss of generality, we may assume that $x^* \in \partial H^0$ and that $\partial H^0 \cap \{\theta\}_{i=1}^n \subseteq \{x^*\}$ (the latter by the fact that $\{\theta\}_{i=1}^n$ is a discrete set). Therefore, we may shift H^0 slightly towards y to \tilde{H}^0 while keeping the mass with respect to θ constant, i.e. such that $\theta(H^0) = \theta(\tilde{H}^0) = \mathfrak{d}(x^*, y; \theta)$. Consider the intersection point w of the straight line L connecting y and x^* with $\partial \tilde{H}^0$, and the point z on L such that w is the midpoint between w and x^* (see Figure 7). By Lemma A.2 we have $m_{\Pi_{\text{exunif}}}^-(x^*, z) = \theta(\tilde{H}^0) = \mathfrak{d}(x^*, y; \theta)$. Moreover, we evidently also have $\mathfrak{d}(x^*, y; \theta) = \mathfrak{d}(x^*, z; \theta)$, and $\mathfrak{d}(z, x^*; \theta) \geq \mathfrak{d}(y, x^*; \theta)$. Thus, using (A.5) and the fact that, for all $w, v \in X$, $m_{\Pi_{\text{exunif}}}^-(w, v) \geq \mathfrak{d}(w, v; \theta)$, we obtain,

$$\begin{aligned} m_{\Pi_{\text{exunif}}}^-(z, x^*) &\geq \mathfrak{d}(z, x^*; \theta) \geq \mathfrak{d}(y, x^*; \theta) > \\ &\mathfrak{d}(x^*, y; \theta) = \mathfrak{d}(x^*, z; \theta) = m_{\Pi_{\text{exunif}}}^-(x^*, z). \end{aligned}$$

i.e. $z P_{\Pi_{\text{exunif}}} x^*$ in contradiction to the initial assumption that $x^* \in \text{CW}(\Pi_{\text{exunif}})$.




Figure 7: ‘Localization’ argument

Conversely, let $x^* \in S(\theta)$, i.e. $x^* R_\theta x$ for all $x \in X$. Consider any fixed $y \in X$ distinct from x^* , and let w denote the midpoint of the line segment connecting x^* and y . Let H^1 be a half-space with $x^* \in H^1$, $w \in \partial(H^1)$ such that $\theta(H^1)$ is minimal among all half-spaces with these two properties. Then, by Lemma A.2,

$$m_{\Pi_{\text{exunif}}}^-(x^*, y) = \theta(H^1). \quad (\text{A.6})$$

Since x^* is in the interior of H^1 and w on its boundary, we have

$$\theta(H^1) \geq \mathfrak{d}(x^*, w; \theta). \quad (\text{A.7})$$

By the assumption $x^* \in S(\theta)$, we have

$$\mathfrak{d}(x^*, w; \theta) \geq \mathfrak{d}(w, x^*; \theta), \quad (\text{A.8})$$

and again by Lemma A.2,

$$\mathfrak{d}(w, x^*; \theta) = m_{\Pi_{\text{exunif}}}^-(y, x^*). \quad (\text{A.9})$$

Combining (A.6) - (A.9), we thus obtain,

$$m_{\Pi_{\text{exunif}}}^-(x^*, y) \geq m_{\Pi_{\text{exunif}}}^-(y, x^*),$$

i.e. $x^* R_{\Pi_{\text{exunif}}} y$. Since y was arbitrarily chosen, we thus obtain $x^* \in \text{CW}(\Pi_{\text{exunif}})$ as desired. \square

It remains to be shown that $S(\theta)$ coincides with the strict Tukey median, and that

these two sets are indeed non-empty.

Proposition 7. *For all θ , the strict Tukey median $T^*(\theta)$ is non-empty and*

$$S(\theta) = T^*(\theta).$$

The proof of Proposition 7 is given through a series of lemmata.

Lemma A.3. *For all $x, y \in X$, $\mathfrak{d}(x; \theta) > \mathfrak{d}(y; \theta)$ implies $xP_{\mathfrak{d}}y$.*

Proof. By assumption there exists a half-space H containing y with $\theta(H) < \mathfrak{d}(x; \theta)$, hence in particular $x \notin H$. Thus,

$$\mathfrak{d}(x, y; \theta) \geq \mathfrak{d}(x; \theta) > \theta(H) \geq \mathfrak{d}(y, x; \theta).$$

□

Note that Lemma A.3 implies $S(\theta) \subseteq T(\theta)$.

Lemma A.4. *For all distinct $x, y \in X$ such that $\mathfrak{d}(x; \theta) = \mathfrak{d}(y; \theta) =: \alpha$, one has $\mathfrak{d}(x, y; \theta) = \alpha$ or $\mathfrak{d}(y, x; \theta) = \alpha$.*

Proof. Let $H \ni x$ be such that $\theta(H) = \alpha$; without loss of generality, we may assume that x is on the boundary $\partial(H)$ of H (otherwise, one may shift the boundary of H to x without increasing $\theta(H)$). For any such H , $\theta(\partial H \setminus \{x\}) = 0$. Indeed, if $\partial H \setminus \{x\}$ contained some voters' tops, an appropriate slight rotation around x to H' would eliminate some of them without including additional ones (by the finiteness of the set $\{\theta_i\}_{i=1}^n$); but this would entail $\mathfrak{d}(x; \theta) < \alpha$, a contradiction.

If $y \notin H$, then $\mathfrak{d}(x, y; \theta) = \alpha$. If $y \in H$, since x is on the boundary of H , H could be changed slightly (by appropriate shift plus slight rotation) to H' with $y \in H'$ eliminating x without including any additional tops (by the argument above). Thus, $\alpha \leq \theta(H') \leq \theta(H) \leq \alpha$, in particular $\theta(H') = \alpha$. In other words, we have constructed H' such that $\theta(H') = \alpha$, $y \in H'$ and $x \notin H'$, hence $\mathfrak{d}(y, x; \theta) = \alpha$. □

Lemma A.5. *For all distinct $x, y \in X$ with $\mathfrak{d}(x; \theta) = \mathfrak{d}(y; \theta) =: \alpha$,*

$$xP_{\mathfrak{d}}y \iff \mathfrak{d}(x, y; \theta) > \alpha \iff \mathcal{H}_x^\alpha \subsetneq \mathcal{H}_y^\alpha, \quad (\text{A.10})$$

where $\mathcal{H}_x^\alpha := \{H \in \mathcal{H} \mid x \in H \text{ and } \theta(H) = \alpha\}$. In particular, the relation P_δ is a strict partial order and

$$S(\theta) = T^*(\theta). \quad (\text{A.11})$$

Proof. The first biconditional in (A.10) follows from Lemma A.4 since $\mathfrak{d}(y, x; \theta) \geq \mathfrak{d}(y; \theta) = \alpha$. Thus, we only need to show that $\mathfrak{d}(x, y; \theta) > \alpha \iff \mathcal{H}_x^\alpha \subsetneq \mathcal{H}_y^\alpha$. If $\mathfrak{d}(x, y; \theta) > \alpha$, there does not exist a half-space H such that $x \in H$, $y \notin H$ and $\theta(H) = \alpha$. Moreover, by Lemma A.4, $\mathfrak{d}(y, x; \theta) = \alpha$, i.e. there exists a half-space H such that $y \in H$, $x \notin H$ and $\theta(H) = \alpha$; hence in fact $\mathcal{H}_x^\alpha \subsetneq \mathcal{H}_y^\alpha$.

Conversely, if $\mathcal{H}_x^\alpha \subsetneq \mathcal{H}_y^\alpha$, there does not exist a half-space H such that $x \in H$, $y \notin H$ and $\theta(H) = \alpha$, hence $\mathfrak{d}(x, y; \theta) > \alpha$.

The equality stated in (A.11) now follows from the definition of the strict Tukey median. \square

We now show that the sets $T^*(\theta)$ and hence $S(\theta)$ are indeed non-empty. To this end, consider the Tukey median set $T(\theta)$, i.e. the depth level set with maximal depth and denote, for all $x \in T(\theta)$, by $\tilde{L}_x(\theta) := \{y \in T(\theta) \mid xR_\delta y\} \setminus \{x\}$ (i.e. the lower contour set of x with respect to R_δ minus the alternative x itself). Moreover, denote the complement of $\tilde{L}_x(\theta)$ in $T(\theta)$ by $\tilde{U}_x(\theta)$, i.e.

$$\tilde{U}_x(\theta) = \{y \in T(\theta) \mid yP_\delta x\} \cup \{x\}$$

(this is the upper contour set of x with respect to P_δ plus the alternative x itself).

Lemma A.6. *For all $x \in T(\theta)$, the sets $\tilde{U}_x(\theta)$ are relative closed in $T(\theta)$.*

Proof. We show that the complementary sets $\tilde{L}_x(\theta)$ are relative open in $T(\theta)$. Consider any pair $x, y \in T(\theta)$ such that $xR_\delta y$ and $x \neq y$, and let α^* be the maximal Tukey depth. We have $\mathfrak{d}(x; \theta) = \mathfrak{d}(y; \theta) = \alpha^*$, and by Lemmas A.4 and A.5, we have $\mathfrak{d}(y, x; \theta) = \alpha^*$. Thus, there exists a half-space H with $\theta(H) = \alpha^*$, $y \in H$, $x \notin H$. Since the voters' tops form a discrete set, we can move the boundary ∂H slightly towards x in a parallel fashion to obtain a half-space H' such that $H \subseteq H'$, $\theta(H') = \theta(H) = \alpha^*$ and $x \notin H'$. Thus, $\mathfrak{d}(y', x; \theta) = \alpha^*$, and hence again by Lemma A.5, $xR_\delta y'$, for all y' in a small neighborhood of y . This shows that $\tilde{L}_x(\theta)$ is relative open in $T(\theta)$. \square

Lemma A.7. *For all profiles θ , $S(\theta)$ is non-empty.*

Proof. Consider chains of upper contour sets, i.e. subsets $\mathcal{C} \subseteq \{\tilde{U}_x(\theta) \mid x \in T(\theta)\}$ totally ordered by set inclusion, and denote by \mathcal{U} the family of all such chains partially ordered by set inclusion. By Zorn's Lemma, there exists a maximal element in \mathcal{U} , i.e. a maximal chain \mathcal{C}^* .

The function $x \mapsto \mathfrak{d}(x; \theta)$ is upper semicontinuous, hence the set $T(\theta)$ of its maximizers is non-empty and closed. Hence, since $T(\theta)$ is clearly also bounded, $T(\theta)$ is a compact set. By Lemma A.6, the elements of \mathcal{C}^* are relative closed in $T(\theta)$, hence as relative closed subsets of the compact set $T(\theta)$ themselves compact.

Consider the directed net (Z, \geq) where $Z := \{x \in T(\theta) \mid \tilde{U}_x(\theta) \in \mathcal{C}^*\}$ and

$$x \geq y \iff \tilde{U}_x(\theta) \subseteq \tilde{U}_y(\theta).$$

By the compactness of $T(\theta)$, the net (Z, \leq) contains a convergent subnet in Z ; let x^* denote its limit. By the orderedness of the chain \mathcal{C}^* , $x \geq y$ implies $x \in \tilde{U}_y(\theta)$; hence by the closedness of $\tilde{U}_y(\theta)$, we have $x^* \in \tilde{U}_y(\theta)$ for all $y \in Z$, and therefore $x^* \in \cap \mathcal{C}^*$.

By Lemma A.5, the relation $P_\mathfrak{d}$ is transitive on $T(\theta)$, hence $\tilde{U}_{x^*}(\theta) \subseteq \tilde{U}_y(\theta)$ for all $y \in Z$, and therefore $\tilde{U}_{x^*}(\theta) \subseteq \cap \mathcal{C}^*$. By the maximality of \mathcal{C}^* , $\tilde{U}_{x^*}(\theta) = \{x^*\}$. By the definition of $\tilde{U}_{x^*}(\theta)$, $x^* \in S(\theta)$, in particular $S(\theta) = T^*(\theta)$ is non-empty. \square

Proof of Proposition 7. By Lemma A.5, we have $S(\theta) = T^*(\theta)$, and by Lemma A.7 $S(\theta)$ is non-empty. This completes the proof of Proposition 7. \square

Proof of Theorem 1. The proof follows from combining Propositions 5, 6 and 7. \square

Remaining Proofs

Proof of Theorem 1'. The first steps in the proof follow closely the proof of Theorem 1. Indeed, many intermediate steps and arguments hold without change in the case of a continuous distribution θ of tops. In particular, Proposition 5 can be shown in an analogous manner, and Lemmas A.2 and A.3 hold without change. Next, we show that

$$\text{CW}(\Pi_{\text{exunif}}) = S(\theta) \tag{A.12}$$

(cf. Proposition 6). As in the proof of Proposition 6, suppose that $x^* \in \text{CW}(\Pi_{\text{exunif}})$ and, by contradiction, $x^* \notin S(\theta)$, i.e. (A.5) for some $y \in X$. As in the proof of Proposition 6, we choose H^0 such that $y \notin H^0$, $x^* \in \partial(H^0)$ and $\theta(H^0) = \mathfrak{d}(x^*, y; \theta)$.

Since θ is continuously distributed, for any positive ε , we can shift H^0 slightly towards y to \tilde{H}^0 as in Fig. 4 above such that $\theta(\tilde{H}^0) < \theta(H^0) + \varepsilon$. If ε is sufficiently small, we obtain

$$\begin{aligned} m_{\Pi_{\text{exunif}}}^-(z, x^*) &\geq \mathfrak{d}(z, x^*; \theta) \geq \mathfrak{d}(y, x^*; \theta) > \\ &\mathfrak{d}(x^*, y; \theta) + \varepsilon > \theta(\tilde{H}^0) = m_{\Pi_{\text{exunif}}}^-(x^*, z). \end{aligned}$$

i.e. $zP_{\Pi_{\text{exunif}}}x^*$ in contradiction to the initial assumption that $x^* \in \text{CW}(\Pi_{\text{exunif}})$.

The converse statement $S(\theta) \subseteq \text{CW}(\Pi_{\text{exunif}})$ follows exactly as in the proof of Proposition 6 above.

By (Demange, 1982, Sect. 2.4.(ii)), the Tukey median set $T(\theta)$ consists of a unique point x^* . In particular, $\mathfrak{d}(x^*, \theta) > \mathfrak{d}(y, \theta)$ for all $y \in X \setminus \{x^*\}$; hence by Lemma A.3, $S(\theta) = T(\theta) = \{x^*\}$. Thus, by (A.12) also $\text{CW}(\Pi_{\text{exunif}}) = \{x^*\} = T(\theta)$. \square

Proof of Theorem 2. Consider any s.i.q. model $\Pi \subseteq \Pi_{\text{quad}}^\mu$. Any prior $\pi \in \Pi$ corresponds to a unique symmetric prior $\bar{\pi} \in \Pi_{\text{quad}}^{\bar{\mu}}$ such that, for all $\theta \in \text{supp } \mu$, $\bar{\pi}_{\mathcal{Q}_\theta} = \pi_{\mathcal{Q}_\theta}$. Denote by $\bar{\Pi}$ the set of all such priors $\bar{\pi}$, i.e. $\bar{\Pi} := \{\bar{\pi} : \pi \in \Pi\}$. As is easily verified, $\bar{\Pi}$ satisfies Assumptions 1-4 of Section 4 (in particular, Tops Certainty) with respect to the distribution $\bar{\mu}$. Moreover, for all distinct $x, y \in X$, we have

$$m_\pi(x, y) = n \cdot m_{\bar{\pi}}(x, y),$$

hence $m_{\bar{\Pi}}^-(x, y) = n \cdot m_{\bar{\Pi}}^-(x, y)$; therefore, the ex-ante majority relations corresponding to Π and $\bar{\Pi}$ coincide and we have $\text{CW}(\Pi) = \text{CW}(\bar{\Pi})$. Since $\bar{\Pi}$ satisfies the assumptions required for Theorem 1 in the case of a finite distribution, and for Theorem 1' in the case of a continuous distribution with convex support, we can conclude that $\text{CW}(\bar{\Pi}) = T^*(\bar{\mu})$. Together with the preceding observation we thus obtain, $\text{CW}(\Pi) = \text{CW}(\bar{\Pi}) = T^*(\bar{\mu})$. \square

Proof of Fact 6.1. Let $N = \{1, \dots, n\}$. Part (i) follows since any x must guarantee $\mathfrak{d}^*(\theta_{-i})$ within $N \setminus \{i\}$ to guarantee $\mathfrak{d}^*(\theta_{-i}) + 1$ in N . Likewise, the second inclusion of (iii) obtains since any x must guarantee $\mathfrak{d}^*(\theta_{-i}) - 1$ within $N \setminus \{i\}$ to guarantee $\mathfrak{d}^*(\theta_{-i})$ in N . The first inclusion of (iii) obtains since any $x \in T(\theta_{-i})$ guarantees $\mathfrak{d}^*(\theta_{-i}) = \mathfrak{d}^*(\theta_i, \theta_{-i})$. Part (ii) is straightforward. \square

Working Paper Series in Economics

recent issues

- No. 160** *Klaus Nehring, Clemens Puppe*: Multi-dimensional social choice under frugal information: the Tukey median as Condorcet winner ex ante, March 2023
- No. 159** *Clemens Puppe, Arkadii Slinko*: Maximal Condorcet domains. A further progress report, December 2022
- No. 158** *Daniel Hoang, Kevin Wiegartz*: Machine learning methods in finance: recent applications and prospects, December 2022
- No. 157** *David Ehrlich, Nora Szech*: How to start a grassroots movement, September 2022
- No. 156** *Klaus Nehring, Clemens Puppe*: Condorcet solutions in frugal models of budget allocation, March 2022
- No. 155** *Clemens Puppe, Jana Rollmann*: Participation in voting over budget allocations. A field experiment, February 2022
- No. 154** *Andranik S. Tangian*: Analysis of the 2021 Bundestag elections. 4/4. The third vote application, January 2022
- No. 153** *Andranik S. Tangian*: Analysis of the 2021 Bundestag elections. 3/4. Tackling the Bundestag growth, January 2022
- No. 152** *Andranik S. Tangian*: Analysis of the 2021 Bundestag elections. 2/4. Political spectrum, January 2022
- No. 151** *Andranik S. Tangian*: Analysis of the 2021 Bundestag elections. 1/4. Representativeness of the parties and the Bundestag, January 2022
- No. 150** *Marta Serra-Garcia and Nora Szech*: Choice architecture and incentives increase COVID-19 vaccine intentions and test demand, April 2021