

**Topological Methods of Preference and Judgment
Aggregation**

by

Alexander M. Jakobsen

MA, Economics, University of Calgary, 2009

BA, Economics, University of Calgary, 2007

BSc, Computer Science, University of Calgary, 2007

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Master of Science

in

THE FACULTY OF GRADUATE STUDIES

(Mathematics)

The University Of British Columbia

(Vancouver)

June 2011

© Alexander M. Jakobsen, 2011

Abstract

Arrow's Impossibility Theorem is a classical result in social choice theory (a branch of economic theory), which states that any system of rules for combining ("aggregating") individual preference relations into a single representative relation results in a "dictatorship" where the combined preference only reflects the wishes of a single individual (provided that the aggregation rule satisfies two basic criteria). Since the 1980s, this result has been reformulated and understood using algebraic topology. The topological approach offers some geometric intuition as to why Arrow's theorem holds, and can also be used to find alternative hypotheses which may escape the dictatorship outcome. A thorough examination of such topological models constitutes the main body of this thesis.

Recently, social choice theory has been generalized (resulting in a field called "judgment aggregation"), and results analogous to Arrow's theorem have been established. The second part of this thesis introduces this field of study, and studies how some of the techniques from topological social choice theory can be extended to understand dictatorship outcomes in the theory of judgment aggregation. Although the analysis is restricted to a rather simple case, it nonetheless highlights the potential for a more general topological model of judgment aggregation, and exposes the main challenges that must be overcome in constructing such a theory.

Table of Contents

Abstract	ii
Table of Contents	iii
List of Tables	iv
List of Figures	v
Acknowledgments	vi
1 Introduction	1
2 Classical Social Choice Theory	3
2.1 Arrow's Impossibility Theorem	4
3 Topological Social Choice Theory	10
3.1 Continuous Social Choice	11
3.1.1 Chichilnisky and Heal (1983)	12
3.2 Discrete Social Choice via Nerves	17
3.2.1 Baryshnikov (1993)	18
4 Judgment Aggregation	29
Bibliography	38

List of Tables

Table 4.1	The doctrinal paradox	30
-----------	---------------------------------	----

List of Figures

Figure 3.1	The nerve $N_{\mathcal{P}}$ for $n = 3$	20
Figure 4.1	The nerve N_X for $X = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$	34

Acknowledgments

I would like to start by thanking my supervisor, Alejandro Adem, for all of his patience and support during my graduate studies in mathematics, and also for encouraging me to take up the topic of topological social choice theory for my thesis. It has been a pleasure to study problems in the (usually empty) intersection of algebraic topology and economic theory, and I am grateful that he granted me the freedom to do so.

I also want to thank all of my new friends here in the UBC mathematics department. There are too many to name them all, but I would particularly like to thank Michele Klaus, Jose Gomez, and Cihan Okay, with whom I have had many helpful discussions.

Last but not least, I would also like to thank my parents, Bent and Birte Jakobsen, as well as the rest of my (quickly growing!) family, for all of their support during my many years as a student. I am afraid I will require their patience and support in that regard for a while yet, as I will be taking up PhD study in economics this fall. It is comforting to know that wherever I am, my family is always with me!

Chapter 1

Introduction

This thesis studies various topological approaches to so-called *aggregation problems* in economic theory. We will be interested in two strands of aggregation theory: that of preferences, and that of “judgments”. The theory of preference aggregation is now a classic component of economic theory. It asks how individual preference relations (orderings) might be combined to give a single ordering which captures the overall wishes of the group of individuals. Of course, there are certain properties that one would like such an aggregation rule to have, and these will be discussed in detail. We also discuss some motivating examples of how this theory might be useful in practical situations (such as the design of voting systems), but will quickly encounter the famous Impossibility Theorem due to Kenneth Arrow. This is a rather negative result, as it states that any aggregation system satisfying two basic (and desirable) criteria will be “dictatorial”. The topological approaches to studying this result yield interesting insights as to why this dictatorship outcome holds true, and also offers a useful framework for identifying aggregation systems which are non-dictatorial.

The other type of aggregation problem studied in this thesis is that of judgment aggregation. This is a relatively new branch of theory (having been formulated less than ten years ago), and is currently a very active area of research. In contrast to preference aggregation, judgment aggregation revolves around the problem of aggregating *choices* that individuals make. Specifically, there is some common set of alternatives, and each individual chooses some subset (subject to the constraint

that only some subsets are allowed). In this setting, the aggregation problem asks whether some particular subset can be taken to represent the overall choices of the individuals. This seemingly simple setup actually encompasses a large variety of situations, including that of preference aggregation (we will show that preference aggregation is a special case of judgment aggregation). As the name suggests, judgment aggregation problems were originally motivated by problems in the design of legal systems, and there is also much interplay between judgment aggregation models and the theory of formal systems and languages.

We begin in section 2 by presenting the classic theory of preference aggregation (also called social choice theory), including a simple proof of Arrow's Impossibility Theorem and a brief discussion of how this outcome can be avoided. Next, we present the "continuous" approach to social choice theory. This is a model of preference aggregation analogous to Arrow's setup, but not formally equivalent. We will see that the problem of finding nice aggregation maps can be fully characterized in this setting using the tools of algebraic topology. Next, we present a topological approach to Arrow's theorem; that is, we will give a valid proof of Arrow's theorem by employing methods of algebraic topology. This approach is interesting in its own right, and also reveals useful geometric insights into the classic aggregation problem; insights which can be used to (correctly) formulate conjectures about ways to escape the dictatorship outcome. Finally, in chapter 4 we show how these techniques might be extended to the theory of judgment aggregation. Although we do not present a complete theory of topological judgment aggregation, the analysis we give is encouraging and seems to indicate that a fairly general theory is possible. It also provides some insight into how such a theory might look and what the relevant obstacles to its construction might be.

Chapter 2

Classical Social Choice Theory

Suppose a group of k individuals must decide how to rank a (finite) number of alternatives. Each individual (which we will also refer to as an *agent* or a *voter*) has his or her own personal choice of how to rank the alternatives, and different individuals may have very different orderings. A central question of social choice theory (also called *preference aggregation*) is whether such individual rankings may be combined in a way which “reasonably” represents the overall wishes of the entire collection of agents.

A classic example is that of voting systems, where the alternatives to be ranked are political candidates, and the agents are the general populace, each with their own preferences regarding the desirability of the different candidates running for office. This is the most commonly used example, but it is not ideal since usually the only relevant characteristic of an agent’s ordering is which candidate is ranked at the very top (that is, you vote for one candidate only; you do not necessarily submit your entire ordering of the candidates).

One might also imagine a situation where a government has proposed a number of public projects, and wishes to ascertain how important the different projects are to the public. Once such a ranking has been achieved, the government can set about implementing the most desirable project. This differs from the voting example in that the government may face various constraints at different points in time. The purpose of gathering individual rankings (and formulating a “social ranking”) is to have a reasonable idea of what people want. Then, once uncer-

tainty has been resolved regarding, say, financial constraints, the government can set about implementing the most desirable project which is actually feasible (given the constraints), and this may or may not coincide with the highest-ranked project in the social ranking.

Alternatively, a group of colleagues (graduate students, perhaps) might be collaborating on a series of projects (joint papers; homework assignments; etc), each of which requires their joint effort. They must decide the order in which to complete these projects, and each of them may have different preferences regarding the order. Can a “fair” ordering be achieved?

2.1 Arrow’s Impossibility Theorem

To answer these sorts of questions, social choice theory begins by fixing a finite set X of *alternatives* (to be ranked) which, for convenience, we will often assume to be the set $[n] = \{1, \dots, n\}$. There is also a finite collection of k agents, each of which is represented by some *preference relation* \succsim over X . Such relations are assumed to be *complete* (for all $i, j \in X$, at least one of $i \succsim j$ or $j \succsim i$ is satisfied) as well as transitive. The statement $i \succsim j$ should be interpreted as “alternative i is weakly preferred over alternative j ”. If both $i \succsim j$ and $j \succsim i$, we write $i \sim j$ and say that the agent is *indifferent* between i and j . Obviously \sim is an equivalence relation on X . If it happens that $i \succsim j$ but not $j \succsim i$, then the individual is said to *strictly prefer* i over j . Corresponding to every preference relation \succsim is a *strict part* \succ , where $i \succ j$ if and only if i is strictly preferred over j . We take \mathcal{R} to be the entire collection of preference relations on X , and let $\mathcal{P} \subset \mathcal{R}$ denote the set of strict preference relations (those without any indifferences; or equivalently, those preference relations which coincide with their strict part). Letting \mathcal{A} be either \mathcal{R}^k or \mathcal{P}^k , we define a *profile* of preferences to be a k -tuple of preference relations belonging to \mathcal{A} . Finally, a *social choice function* (also called an *aggregation map*) is a function $f : \mathcal{A} \rightarrow \mathcal{R}$; that is, it is a rule which assigns to each k -tuple $\vec{\succsim} = (\succsim_1, \dots, \succsim_k)$ of (arbitrary or strict, depending on \mathcal{A}) preference relations a *social ranking* (or *aggregated preference*) $f(\vec{\succsim}) \in \mathcal{R}$. For such a function, we call \mathcal{A} the *domain of preferences*.

There are various properties that one might want an aggregation map f to have,

one of which is the following:

Definition 2.1.1 (Pareto Property). *An aggregation map $f : \mathcal{A} \rightarrow \mathcal{R}$ satisfies the Pareto property (we also say f is Paretian) if for all $x, y \in X$ and all profiles $\vec{\succsim} = (\succsim_1, \dots, \succsim_k) \in \mathcal{A}$, we have that $\succsim := f(\vec{\succsim})$ satisfies $x \succ y$ whenever $x \succ_i y$ for every agent i .*

The Pareto criterion is one of the most basic properties we might expect a reasonable aggregation map to satisfy. It says that if every single agent strictly prefers an alternative x over another alternative y , then the social preference relation must also rank x strictly above y . It can be interpreted as an “agreement” property: if everyone agrees about the ranking between two alternatives, then there is no reason why the social ranking between those two alternatives should not satisfy each individual’s ranking. A slightly more controversial property is the following:

Definition 2.1.2 (Independence of Irrelevant Alternatives). *An aggregation map $f : \mathcal{A} \rightarrow \mathcal{R}$ satisfies the Independence of Irrelevant Alternatives (IIA) axiom if for all $x, y \in X$ and all profiles $(\succsim_1, \dots, \succsim_k), (\succsim'_1, \dots, \succsim'_l) \in \mathcal{A}$ (denoted $\vec{\succsim}$ and $\vec{\succsim}'$, respectively), we have that*

$$\succsim_i|_{\{x,y\}} = \succsim'_i|_{\{x,y\}} \text{ for all } i \Rightarrow \succsim|_{\{x,y\}} = \succsim'|_{\{x,y\}},$$

where $\succsim := f(\vec{\succsim})$ and $\succsim' := f(\vec{\succsim}')$.

In words, this says that the social ranking between two alternatives $x, y \in X$ should only depend on how these two alternatives (and not some other alternative z) are ranked by the individuals. That is, if $\vec{\succsim}$ and $\vec{\succsim}'$ are two profiles of preferences where no individual has changed their ranking of x and y , then the corresponding social rankings should rank x and y the same way. This seems harmless at first, but we will see that it is in fact quite a powerful axiom.

One property which any desirable map f should *not* satisfy is the following:

Definition 2.1.3 (Dictatorship). *An aggregation map $f : \mathcal{A} \rightarrow \mathcal{R}$ is Dictatorial if there is an agent d such that for all $x, y \in X$ and all preference profiles $\vec{\succsim} = (\succsim_1, \dots, \succsim_l) \in \mathcal{A}$, we have that $\succsim := f(\vec{\succsim})$ satisfies $x \succ y$ whenever $x \succ_d y$. (If this happens, agent d is the “dictator”).*

A dictatorial map is one where the social ranking always agrees with the strict part of the dictator's ranking. That is, agent d is free to submit any preference relation \succsim_d , but if (for any $x, y \in X$) $x \succ_d y$, then the corresponding social ranking will put x strictly above y as well, regardless of how all other agents rank x and y . This is obviously an undesirable property, since it conflicts with practically any sense of "fairness": all other voters could have $y \succ x$, but the social ranking will place x strictly above y as long as $x \succ_d y$. The only situation where other voters could (potentially) influence the social ranking of two alternatives is when the dictator is indifferent between those alternatives.

At this point, it is worth asking whether this setup (that is, attempting to find non-dictatorial aggregation maps which are Paretian and satisfy IIA) is "reasonable"; does it successfully capture real situations? In the voting example above, the Paretian property is reasonable, but the IIA axiom seems unnecessary since, as we have mentioned, the only issue of relevance in most voting systems is who each individual's "most preferred" candidate is. However, the other two examples do seem to warrant both the Pareto and IIA axioms. There is also the issue of transitive preferences: it is well documented in psychology and experimental economics that, in various situations, individuals may exhibit preferences which fail to be transitive. Yet, there are also many situations where transitive preferences *do* make sense (including all three examples at the start of this section). There are other potential shortcomings as well which we do not discuss here. The purpose of this discussion, however, is to point out that an economic model differs from a purely mathematical model in that it consists of a formal model together with an interpretation. As we have seen, any formal model permits several interpretations, some of which may or may not raise doubts about basic assumptions in the formal model. Thus, the "reasonableness" of a theoretical model cannot be evaluated by examining only its formal properties, nor by examining its properties in the context of any single interpretation. No model will ever be completely reasonable in each of its conceivable interpretations; we will be content knowing that there are, at least, many interesting situations where this model is quite reasonable, and focus our attention on the formal properties. Of these, the most famous is the following theorem due to Arrow [1]:

Theorem 2.1.1 (Arrow’s Impossibility Theorem). *Suppose $|X| \geq 3$ and $\mathcal{A} \in \{\mathcal{R}^k, \mathcal{P}^k\}$. If $f : \mathcal{A} \rightarrow \mathcal{R}$ is Paretian and satisfies the IIA condition, then f is dictatorial.*

Most proofs of Arrow’s theorem are rather long and complicated (the topological proof we present in section 3.2 is no exception). There is, however, a fairly brief proof given by Geanakoplos [6], which we now present.

Lemma 2.1.2 (Geanakoplos, “Extremal Lemma”). *Let $b \in X$. Suppose \succsim_i is a profile where, for each $i = 1, \dots, k$, \succsim_i either has b at the very top (meaning $b \succ_i x$ for every $b \neq x \in X$) or at the very bottom ($x \succ_i b$ for every $b \neq x \in X$). Then $\succsim := f(\vec{\succsim})$ also ranks b at the very top or the very bottom.*

Proof. Suppose (toward a contradiction) that there are distinct $a, b, c \in X$ and that \succsim satisfies $a \succ b \succ c$. We modify \succsim to give a profile \succsim' by making the following changes to each individual profile \succsim_i :

- if $b \succ_i c \succ_i a$, make no changes;
- if $b \succ_i a \succ_i c$ or $b \succ_i c \sim_i a$, move c so that $b \succ'_i c \succ'_i a$;
- if $a \succ_i c \succ_i b$ or $a \sim_i c \succ_i b$, move c so that $c \succ'_i a \succ'_i b$.

Notice that $\succsim|_{\{a,b\}} = \succsim'|_{\{a,b\}}$ and $\succsim|_{\{b,c\}} = \succsim'|_{\{b,c\}}$ (that is, no agents change their rankings between a and b , or between b and c). So, by IIA, $\succsim' := f(\vec{\succsim}')$ also satisfies $a \succ' b \succ' c$ (hence $a \succ' c$). However, each order \succsim'_i has $c \succ'_i a$, so that the Pareto property gives $c \succ' a$, a contradiction. \square

Proof of Arrow’s Theorem (Geanakoplos). The proof is given in three steps.

Step 1. Let $b \in X$. We claim there is a voter $d^* = d(b)$ who, for some profile $\vec{\succsim}$ where $f(\vec{\succsim})$ ranks b at the very bottom, can change his preference (yielding a modified profile $\vec{\succsim}'$) so that $f(\vec{\succsim}')$ ranks b at the very top. To see this, let $\vec{\succsim}$ be a profile where every voter has b at the very bottom of their ranking. By the Pareto property, $\succsim := f(\vec{\succsim})$ ranks b at the very bottom as well. Now let individual voters successively move b from the bottom to the top, leaving other relative rankings unchanged. Let d^* be the first agent for which the corresponding social ranking of b changes (such an agent exists because if all agents have b at the very top,

then so does the corresponding social ranking). Let \succsim^1 denote the profile where agents $1, \dots, d^* - 1$ have moved b to the top, and let \succsim^2 be the profile where agents $1, \dots, d^*$ have moved b to the top. By the lemma, $\succsim^2 := f(\succsim^1)$ ranks b at the very top.

Step 2. Agent $d^* = d(b)$ is a dictator over any pair $\{a, c\}$ with $a \neq b \neq c$. To see this, define \succsim^3 from \succsim^2 by having d^* move a above b (thus voter d^* now ranks a at the very top, so that $a \succ_{d^*}^3 b \succ_{d^*}^3 c$), and let the other voters arrange their relative rankings of a and c arbitrarily (but leaving b in its extreme position). The individual rankings of a and b in \succsim^3 agree with those in \succsim^1 , and \succsim^1 ranks b at the very bottom, so we have $a \succ^3 b$. The individual rankings of b and c agree with those in \succsim^2 , we also have $b \succ^3 c$. Thus $a \succ^3 c$. Since we chose a arbitrarily from $\{a, c\}$ (that is, we could perform the same procedure above with c in place of a), this proves the claim.

Step 3. We prove d^* is a dictator over every pair $\{a, b\}$ (hence d^* is a dictator in the usual sense). Take some $c \notin \{a, b\}$, and place c at the bottom of each social ranking (just as in Step 1). Applying Step 2 to this case, we see that there is a voter $d(c)$ who is a dictator over any pair $\{x, y\}$ with $c \notin \{x, y\}$. In particular, $\{a, b\}$ is such a pair. Since d^* is a dictator for $\{a, b\}$, and since dictators must obviously be unique, this proves that $d(c) = d^*$; that is, d^* is a dictator. \square

Arrow's theorem is a rather negative result, especially because the hypotheses of the theorem are quite basic and could easily apply to many different situations. Moreover, it is not obvious which requirements should be changed to allow non-dictatorial aggregation maps. The Pareto property is clearly quite desirable, and the IIA axiom is also quite natural. Another alternative, and the one which has received the most attention in the literature, is to restrict the domain \mathcal{A} of possible preferences. A suitable restriction is found by considering *single-peaked preferences*, which we now describe.

Given a set $[n] = \{1, \dots, n\}$ of alternatives, fix a linear order \geq on $[n]$ (this means \geq is reflexive, transitive, and complete; we also let $>$ denote the strict part of the relation). A preference relation \succsim on $[n]$ is said to be *single-peaked with respect to* \geq if there is an $x \in [n]$ so that for all $y, z \in [n]$, $x \geq z > y \Rightarrow z \succ y$ and $y > z \geq x \Rightarrow z \succ y$. The idea is that if we line up the alternatives in $[n]$ according

to the order \geq , then the preference relation has some maximal element $x \in [n]$ so that \succsim is “increasing” as we approach x from either side. For example, suppose $n = 3$ and that \geq is the strict order $1 > 2 > 3$. The set \mathcal{P}_{\geq} of strict preference relations which are single-peaked with respect to \geq must satisfy either $1 \succ 2 \succ 3$, $1 \prec 2 \succ 3$, or $1 \prec 2 \prec 3$. The second form does not put any restrictions on the relative rankings of 1 and 3, so that both $2 \succ 3 \succ 1$ and $2 \succ 1 \succ 3$ qualify. Thus $\mathcal{P}_{\geq} = \{1 \succ 2 \succ 3, 2 \succ 3 \succ 1, 2 \succ 1 \succ 3, 3 \succ 2 \succ 1\}$.

Interestingly, single-peaked domains do escape the dictatorship result. In fact, if k is odd and preferences are strict, pairwise majority voting gives the desired aggregation map. We do not prove these results here, but remark that, at first glance, there is no obvious reason why single-peaked domains are a suitable restriction for escaping Arrow’s theorem. Indeed, the condition might seem somewhat random or miraculous. We will see in section 3.2, however, that the topological approach to Arrow’s theorem provides appropriate geometric intuition to anticipate single-peaked domains quite easily.

Chapter 3

Topological Social Choice Theory

At first glance, social choice theory does not appear to have any connection to topology whatsoever. Indeed, the central ingredients of the theory are quite primitive. Thus far, we have only encountered “combinatorial” arguments in the theory, in that all proofs are given from first principles without any need to introduce more advanced mathematical concepts, let alone the formidable tools of algebraic topology. How is algebraic topology able to help us understand problems of social choice?

One possibility is to alter the classical model -that is, to investigate slightly different aggregation problems- so that the model may be more readily susceptible to a topological analysis. Initiated by Chichilnisky [3], this so-called “continuous social choice” model was in fact the first approach taken to understanding results such as Arrow’s Impossibility Theorem in a topological framework, and it has received the most attention in the literature. In section 3.1, we present the main results of this theory, due to Chichilnisky and Heal [4] (an extended version of Chichilnisky’s 1980 paper).

One may well wonder if the classical theory can be understood in a topological setting without modifying the aims of the original model. Can Arrow’s theorem, as originally stated, be understood through topological means? As we shall see, a discrete form of Arrow’s theorem (that is, one in which the choice space X is finite) permits a very natural and interesting topological formulation. A hint as to why this is so is the IIA assumption, which can be interpreted as a sort of “sta-

bility” condition: if one does not modify a profile of preferences “too much” (in that the individual rankings of some pair are unchanged in the profile), then the resulting social preference will not change the ranking of that pair. This suggests that the aggregation map f may be continuous in some topology. Indeed, section 3.2 presents a remarkable paper by Baryshnikov [2], where the IIA axiom is exploited to construct a nontrivial topological space (via nerves) to represent agent preferences, transform f to a continuous map, and examine issues of social choice. Hence, we will have faithfully formulated the classical problem in a useful topological setting. It is Baryshnikov’s technique that we will extend in chapter 4 to the theory of judgment aggregation.

3.1 Continuous Social Choice

To understand the setup for the continuous social choice model, it is necessary to introduce some ideas and terminology from microeconomic theory. Given a rational preference relation \succsim on a set X , a *utility function* is a map $u : X \rightarrow \mathbb{R}$ such that for all $x, y \in X$, $u(x) \geq u(y) \Leftrightarrow x \succsim y$; that is, it is an order-preserving embedding of (X, \succsim) in the reals. The level sets of u , given by $\{x \in X \mid u(x) = c\}$ for all choices of $c \in u(X)$, are called *indifference curves*. It is easy to see that if $I \subseteq X$ is an indifference curve, then $x \sim y$ for all $x, y \in I$, so that the collection of indifference curves inherits a strict linear ordering from \succsim . Of course, any given relation \succsim has infinitely many utility representations (simply compose u with any strictly increasing function $f : \mathbb{R} \rightarrow \mathbb{R}$ to form a new utility function). It is common in microeconomic theory to consider the choice space X to be some subset of \mathbb{R}^n (often the positive orthant), with the interpretation that each axis measures the quantity of some particular good. Under various technical assumptions which we will not discuss here, a utility function u can then be realized as a continuously differentiable function.

Chichilnisky and Heal [4] begin by noting that if the gradient $\nabla(u)$ of a utility function $u : X \rightarrow \mathbb{R}$ does not vanish on X , then the relation which \succsim represents may be characterized by normalizing $\nabla(u(x))$ to length 1 at every point $x \in X$. Since u is a real-valued utility function, the gradient $\nabla(u(x))$ indicates the “most preferred” direction at a point $x \in X$; however, we are interested only in the *ordinal* properties

of u (that is, the ordering \succsim which u represents), so that by normalizing to length 1 we are ignoring information about the “intensity” of preferences. Obviously a nowhere-vanishing gradient is required to carry out this construction; this means the function u must have no critical points, which (interpreting the underlying economic framework) means there is no *satiation* (so “more is always better”).

For convenience, we take the choice space X to be the closed unit ball in \mathbb{R}^n , for some choice of n . Let $V(X)$ denote the space of all C^1 vector fields on X , equipped with the C^1 topology. Motivated by the above discussion, we define a *preference* to be a map $p : X \rightarrow \mathbb{R}^n$ such that for every $x \in X$, $\|p(x)\| = 1$ and there exists a real-valued utility function $u : X \rightarrow \mathbb{R}$ so that $p(x)$ is locally the gradient of u .¹ Our objective is to examine problems of social choice for various collections \mathcal{P} of such preferences (viewed as subspaces of $V(X)$).

Given a space $\mathcal{P} \subseteq V(X)$ of preferences and a collection of k individuals, a *social choice rule* is a map $\phi : \mathcal{P}^k \rightarrow \mathcal{P}$. If $\phi(p_1, \dots, p_k) = \phi(p_{\sigma(1)}, \dots, p_{\sigma(k)})$ for every permutation σ of the agents, ϕ is said to be *anonymous*. If $\phi(p, \dots, p) = p$ for every $p \in \mathcal{P}$, then ϕ is *unanimous*. Note that unanimity is weaker than the Pareto criterion. Of course, ϕ is required to be continuous.

3.1.1 Chichilnisky and Heal (1983)

In this section, we present the main results of Chichilnisky and Heal [4]. In particular, we give a complete characterization of the continuous topological choice problem for the case where \mathcal{P} is a CW complex of finite-type.²

Theorem 3.1.1 (C.H. Theorem 1). *Suppose \mathcal{P} (as a subspace of $V(X)$) is a connected CW complex of finite type. Then there exists a continuous, unanimous, and anonymous social choice rule $\phi : \mathcal{P}^k \rightarrow \mathcal{P}$ for every $k \geq 2$ if and only if \mathcal{P} is contractible.*

Proof. (\Leftarrow). If \mathcal{P} is contractible, then each homotopy group $\pi_i(\mathcal{P})$ is trivial. We may therefore extend the identity map $i : \mathcal{P} \rightarrow \mathcal{P}$ to a map $r : V(X) \rightarrow \mathcal{P}$ by results

¹That is, p is a continuously differentiable, locally integrable vector field. Note that $p(x)$ is not, strictly speaking, *equal* to the gradient $\nabla(u(x))$ (since we normalize $\|p(x)\| = 1$), but $p(x)$ is assumed to be in the *direction* of $\nabla(u(x))$.

²By this, we mean a CW complex with a finite number of cells in each dimension. Such CW complexes are sometimes called *parafinite* in the topological social choice literature.

of obstruction theory (see, for example, section 8.4 of Spanier [8]). This means r is a retract of $V(X)$ onto \mathcal{P} . Let $f: \mathcal{P}^k \rightarrow V(X)$ be given by $f(p_1, \dots, p_k) = \frac{1}{k} \sum_{i=1}^k p_i$ (this is well-defined because $V(X)$, the space of all vector fields on X , is convex). Define $\phi: \mathcal{P}^k \rightarrow \mathcal{P}$ by $\phi = r \circ f$. Then, if $p \in \mathcal{P}$, we have $\phi(p, \dots, p) = r \circ f(p, \dots, p) = r(p) = p$, so that ϕ is unanimous. We also have $\phi(p_{\sigma(1)}, \dots, p_{\sigma(k)}) = r \circ f(p_{\sigma(1)}, \dots, p_{\sigma(k)}) = r \circ f(p_1, \dots, p_k) = \phi(p_1, \dots, p_k)$, so that ϕ is anonymous, as required.

(\Rightarrow). We will prove that all homotopy groups $\pi_i(\mathcal{P})$ are trivial; for then the constant map $f: \mathcal{P} \rightarrow \{p_0\}$ ($p_0 \in \mathcal{P}$) induces an isomorphism $f_*: \pi_n(\mathcal{P}) \rightarrow \pi_n(\{p_0\})$ for each n . Since \mathcal{P} is connected, we may then apply Whitehead's theorem to conclude f is a homotopy equivalence. Hence \mathcal{P} is homotopy equivalent to $\{p_0\}$, which means \mathcal{P} is contractible.

To see that the homotopy groups are trivial, suppose (toward a contradiction) that some minimal index i satisfies $\pi_i(\mathcal{P}) \neq 0$.

Case 1: $i = 1$. Let $k \geq 2$; by hypothesis, we have a continuous, anonymous, unanimous map $\phi: \mathcal{P}^k \rightarrow \mathcal{P}$ inducing a homomorphism of groups $\phi_*: \pi_1(\mathcal{P}^k) \rightarrow \pi_1(\mathcal{P})$. It is clear by definition that ϕ_* inherits the unanimity property of ϕ , so that $\phi_*(x, \dots, x) = x$ for every $x \in \pi_1(\mathcal{P})$; similarly, ϕ_* must also be anonymous. Hence

$$\begin{aligned} \phi_*(x, \dots, x) &= \phi_*((x, e, \dots, e) \cdot (e, x, e, \dots, e) \cdot \dots \cdot (e, \dots, e, x)) \\ &= \phi_*(x, e, \dots, e) \cdot \dots \cdot \phi_*(e, \dots, e, x) \\ &= \phi_*(x, e, \dots, e) \cdot \dots \cdot \phi_*(x, e, \dots, e) \\ &= \phi_*(x, e, \dots, e)^k. \end{aligned}$$

So, for all $x, y \in \pi_1(\mathcal{P})$, we have

$$\begin{aligned} x \cdot y &= \phi_*(x, e, \dots, e)^k \cdot \phi_*(e, y, e, \dots, e)^k \\ &= \phi_*(x^k, y^k, e, \dots, e) \\ &= \phi_*(y^k, x^k, e, \dots, e) \\ &= y \cdot x. \end{aligned}$$

Thus $\pi_1(\mathcal{P})$ is abelian. By the Hurewicz theorem, $H_1(\mathcal{P})$ is isomorphic to the

abelianization of $\pi_1(\mathcal{P})$; hence $H_1(\mathcal{P}) \cong \pi_1(\mathcal{P})$. From cellular homology (and its equivalence to singular homology), we know that $H_1(\mathcal{P})$ is generated by at most d elements, where d is the number of 1-cells of \mathcal{P} . By assumption, \mathcal{P} has finitely many cells in each dimension, so that $H_1(\mathcal{P})$ (hence $\pi_1(\mathcal{P})$) is finitely generated. Let x be a generator of any of the (free part) groups \mathbb{Z} in this decomposition. From the calculation above (employing additive notation), we have $x = \phi_*(x, \dots, x) = k\phi_*(x)$; as k was arbitrary, this must hold for every $k \geq 2$, forcing $x = 0$. Similarly, if $x \in \mathbb{Z}_{p_j}$ is a generator for some j , the same condition implies $x = 0$. Hence all generators of $\pi_1(\mathcal{P})$ are zero, implying $\pi_1(\mathcal{P}) = 0$, a contradiction.

Case 2: $i \geq 2$. By the Hurewicz Isomorphism Theorem and minimality of i , we have $\pi_i(\mathcal{P}) \cong H_i(\mathcal{P})$. As above, this implies $\pi_i(\mathcal{P})$ is finitely generated. So, the proof for Case 1 above carries through to this case as well (simply let ϕ_* denote the induced homomorphism on the i^{th} homotopy group instead of the first group), and we conclude that $\pi_i(\mathcal{P}) = 0$.

Hence all homotopy groups of \mathcal{P} are trivial. □

Remark. The proof above is the one given by Chichilnisky and Heal. However, it can be simplified and shortened by observing that all homotopy groups of a connected, finite-type CW complex must be finitely generated. This can be shown using a simplicial approximation argument (see section 2.C of Hatcher [7]).

The key idea behind this theorem is that if \mathcal{P} is a contractible subspace of $V(X)$, then we may perform an “averaging” operation in $V(X)$ and compose this with a deformation retract $r : V(X) \rightarrow \mathcal{P}$ to yield a continuous, anonymous, and unanimous choice rule (provided \mathcal{P} is a CW complex of finite type). In fact, the contractibility hypothesis ensures that any continuous, anonymous, and unanimous choice rule must be “equivalent” to a rule derived from such a convex averaging procedure, because contractibility of \mathcal{P} implies that any two maps $\phi, \psi : \mathcal{P}^k \rightarrow \mathcal{P}$ are homotopic. To see this, suppose $h : [0, 1] \times \mathcal{P} \rightarrow \mathcal{P}$ is a homotopy, where $h_0 = id_{\mathcal{P}}$ and $h_1(p) = p_0$ for all $p \in \mathcal{P}$, where p_0 is a point of \mathcal{P} . Take $H :$

$[0, 1] \times \mathcal{P}^k \rightarrow \mathcal{P}$ given by

$$H_t(p) = \begin{cases} h_{2t} \circ \phi(p) & \text{for } 0 \leq t \leq \frac{1}{2} \\ h_{2(1-t)} \circ \psi(p) & \text{for } \frac{1}{2} \leq t \leq 1 \end{cases}.$$

Clearly H is continuous with $H_0 = \phi$ and $H_1 = \psi$, so that ϕ and ψ are homotopic. Hence, we have the following theorem:

Theorem 3.1.2 (C.H. Theorem 2). *If \mathcal{P} (as a subspace of $V(X)$) is a connected, contractible CW complex of finite type, then any continuous, anonymous, and unanimous social choice rule ψ is homotopic to a convex averaging rule ϕ .*

Proof. By Theorem 3.1.1 and its proof, we have a continuous, anonymous, and unanimous rule $\phi : \mathcal{P}^k \rightarrow \mathcal{P}$ which is given by convex averaging. If $\psi : \mathcal{P}^k \rightarrow \mathcal{P}$ is any other continuous, anonymous, and unanimous rule, then the above argument shows that ψ is homotopic to ϕ . \square

Corollary 3.1.3 (C.H. Corollary 1). *Suppose \mathcal{P} is a connected CW complex of finite type and let $k \geq 2$. Then \mathcal{P} is contractible if and only if \mathcal{P}^k admits $\Delta(\mathcal{P}^k) = \{(p, \dots, p) \in \mathcal{P}^k \mid p \in \mathcal{P}\} \cong \mathcal{P}$ as a deformation retract.*

Proof. If \mathcal{P} is a deformation retract of \mathcal{P}^k , then \mathcal{P} and \mathcal{P}^k are homotopy equivalent, so that for each i we have $\pi_i(\mathcal{P}) \cong \pi_i(\mathcal{P}^k)$. But $\pi_i(\mathcal{P}^k) \cong \bigoplus_{j=1}^k \pi_i(\mathcal{P})$, so that $\pi_i(\mathcal{P}) \cong \bigoplus_{j=1}^k \pi_i(\mathcal{P})$. Since the homotopy groups $\pi_i(\mathcal{P})$ are finitely generated, this implies $\pi_i(\mathcal{P}) = 0$. As in the proof of Theorem 3.1.1, this implies \mathcal{P} is contractible.

Conversely, suppose $r : \mathcal{P} \rightarrow \{p_0\}$ ($p_0 \in \mathcal{P}$) is a deformation retract. The map $(id_{\mathcal{P}}, r, \dots, r) : \mathcal{P}^k \rightarrow \mathcal{P}^k$ is continuous and sends \mathcal{P}^k to the set

$$Q = \{(p, p_0, \dots, p_0) \mid p \in \mathcal{P}\},$$

which is homeomorphic to \mathcal{P} . \square

Combining the results of Theorems 3.1.1 and 3.1.2, we arrive at an alternate characterization of the social choice problem: a continuous, unanimous, and anonymous choice rule exists for each $k \geq 2$ if and only if for every $k \geq 2$, \mathcal{P}^k retracts onto the “unanimous” profiles $\Delta(\mathcal{P}^k) = \{(p, \dots, p) \in \mathcal{P}^k \mid p \in \mathcal{P}\}$.

We may also use the result of Theorem 3.1.1 to examine the social choice problem when \mathcal{P} is not connected.

Corollary 3.1.4 (C.H. Corollary 2). *Suppose \mathcal{P} (as a subspace of $V(X)$) is a CW complex of finite type. Then a continuous, anonymous, and unanimous choice rule exists for each $k \geq 2$ if and only if each connected component of \mathcal{P} is contractible.*

Proof. (\Rightarrow). Let $k \geq 2$, and suppose $\phi : \mathcal{P}^k \rightarrow \mathcal{P}$ is continuous, anonymous, and unanimous. If C is a connected component of \mathcal{P} (hence a finite type CW complex as well), then the map $\tilde{\phi} = \phi|_{C^k} \rightarrow \mathcal{P}$ inherits the unanimity property of ϕ , and so $\tilde{\phi}$ maps any unanimous profile $(p, \dots, p) \in C^k$ to $p \in C$. But $\tilde{\phi}$ is also continuous, so by connectedness of C^k , we have $\tilde{\phi}(C^k) \subseteq C$. It is also clear that $\tilde{\phi}$ inherits anonymity from ϕ . Thus, for every $k \geq 2$, we have a continuous, anonymous, and unanimous map $C^k \rightarrow C$, so that Theorem 3.1.1 implies that C must be contractible.

(\Leftarrow). Let $k \geq 2$ and suppose each connected component of \mathcal{P} is contractible. Since \mathcal{P} has finitely many cells in each dimension, \mathcal{P} must have only finitely many connected components C_α , which we enumerate C_1, \dots, C_n . Each component is itself a CW complex of finite type, so that by Theorem 3.1.1 we have a continuous, anonymous, and unanimous map $\phi_\alpha : C_\alpha^k \rightarrow C_\alpha$ for each α . Fix an element p_α^* in each component C_α , and define a map $\phi : \mathcal{P}^k \rightarrow \mathcal{P}$ as follows: if $p = (p_1, \dots, p_k) \in \mathcal{P}^k$ belongs to some C_α^k , set $\phi(p) = \phi_\alpha(p)$. Otherwise, suppose $p_i \in C_{\alpha_i}$, and let α be the smallest such α_i (according to our enumeration C_1, \dots, C_n). Consider the profile $p' = (p'_1, \dots, p'_k)$, where $p'_i = p_i$ if $p_i \in C_\alpha$, and $p'_i = p_\alpha^*$ otherwise. Then set $\phi(p) = \phi_\alpha(p')$. It is easy to see that ϕ must be anonymous and unanimous. To see that ϕ is continuous, let $p = (p_1, \dots, p_k) \in \mathcal{P}^k$ and let V be an open neighborhood of $\phi(p)$. We must find an open neighborhood U of p for which $\phi(U) \subseteq V$. By continuity of the maps ϕ_α , this can clearly be done if $p \in C_\alpha^k$ for some α . So, suppose (without loss of generality) that p_1, \dots, p_m belong to the minimal component α , and that $m < k$ (so that all other p_i belong to components C_{α_i} with $\alpha_i > \alpha$). By definition, then, $\phi(p) = \phi_\alpha(p_1, \dots, p_m, p_\alpha^*, \dots, p_\alpha^*)$. So $\phi(p) = \phi(p_1, \dots, p_m, p_\alpha^*, \dots, p_\alpha^*)$ as well. By continuity of ϕ_α , there is a neighborhood U_α of $(p_1, \dots, p_m, p_\alpha^*, \dots, p_\alpha^*)$ with $\phi_\alpha(U_\alpha) \subseteq V$. Hence there are open sets $W_1, \dots, W_k \subseteq C_\alpha$ with $(p_1, \dots, p_m, p_\alpha^*, \dots, p_\alpha^*) \in W_1 \times \dots \times W_k \subseteq U_\alpha$, so $\phi_\alpha(W_1 \times \dots \times W_k) \subseteq V$ as well. Now define $U = W_1 \times \dots \times W_m \times C_{\alpha_1} \times \dots \times C_{\alpha_{k-m}}$, where

C_{α_i} is the connected component containing p_{m+i} . Then C_{α} is still the minimal component for any $q = (q_1, \dots, q_k) \in U$, and so $\phi(q) = \phi_{\alpha}(q_1, \dots, q_m, p_{\alpha}^*, \dots, p_{\alpha}^*) \in V$ since $(q_1, \dots, q_m, p_{\alpha}^*, \dots, p_{\alpha}^*) \in W_1 \times \dots \times W_k$. Since U is an open neighborhood of p , this completes the proof. \square

We conclude this section with a simple example. Consider the set of all linear preference relations on the unit ball $B \subset \mathbb{R}^n$. This means each utility function $u : B \rightarrow \mathbb{R}$ has the form $u(x_1, \dots, x_n) = a_1x_1 + \dots + a_nx_n$ for some constants a_i , not all 0. Since such a function can be identified with its gradient vector (a_1, \dots, a_n) , the set of all such linear functions corresponds to the set of all nonzero vectors of \mathbb{R}^n . Recall, however, that we normalize gradient vectors to length 1 since we are only interested in ordinal information. So, in this case our preference space will consist of all unit vectors of \mathbb{R}^n ; that is, $\mathcal{P} = S^{n-1}$. Obviously this space is not contractible. If, however, we remove a single point (so that we are ignoring all linear utility functions with gradient a multiple of this particular unit vector), the resulting space will be contractible, so we will have continuous, anonymous, and unanimous social choice rules $\mathcal{P}^k \rightarrow \mathcal{P}$ for every $k \geq 2$. An even wilder restriction would be to consider only those points on the sphere where each coordinate is positive, which will obviously be a contractible space as well. Is such a restriction reasonable? This, of course, depends on the interpretation. If the x_i indicate quantities of (desirable) goods, then it might be reasonable that all coefficients a_i are positive.

3.2 Discrete Social Choice via Nerves

One consequence of the IIA axiom is that it allows the aggregation problem to be studied in terms of sets of preference relations, rather than individual relations. For example, suppose f is an IIA social choice function, and that we wish to understand how f ranks alternatives 1 and 2 for different profiles. The IIA axiom tells us that the social ranking of 1 and 2 only depends upon the rankings of 1 and 2 submitted by the voters, and not on how any other alternatives are ranked. That is, each voter could submit a different preference relation, but so long as the rankings of 1 and 2 are unchanged, the social ranking of 1 and 2 will be the same. So, to study the relevant properties of such a function f (for example, dictatorship issues), it will

suffice to study how f behaves on sets of profiles where the ranking of pairs is fixed. This is precisely the approach taken by Baryshnikov [2]. We will see that Arrow's theorem may be understood topologically by studying the intersection properties of these sets; in fact, the main motivation for Baryshnikov's approach is to develop an alternate proof of Arrow's theorem which lends itself well to geometric intuition.

3.2.1 Baryshnikov (1993)

As in the original framework, let $[n]$ denote the set $\{1, \dots, n\}$ of alternatives, and let \mathcal{R} denote the set of all rational preference relations over $[n]$. Assume there are $k \geq 2$ individuals. We will be interested in the set $\mathcal{P} \subset \mathcal{R}$ of *strict* relations (those without any indifferences). To see why, consider the following lemma:

Lemma 3.2.1 (Baryshnikov Lemma 1). *Suppose $f : \mathcal{R}^k \rightarrow \mathcal{R}$ is a social choice rule satisfying the Pareto property and IIA. Then $f(\mathcal{P}^k) \subseteq \mathcal{P}$.*

Since we are interested in proving Arrow's theorem, this lemma shows that it will suffice to prove that any map $f : \mathcal{P}^k \rightarrow \mathcal{P}$ satisfying the Pareto property and IIA is dictatorial. For if a non-dictatorial, Pareto, and IIA map $f : \mathcal{R}^k \rightarrow \mathcal{R}$ exists, then $f|_{\mathcal{P}^k}$ would be a non-dictatorial map $f : \mathcal{P}^k \rightarrow \mathcal{P}$ satisfying IIA and the Pareto property.

Proof. Suppose toward a contradiction that there are distinct elements $a, b \in [n]$ and some profile $\succsim = (\succsim_1, \dots, \succsim_k) \in \mathcal{P}^k$ with $\succsim = f(\succsim)$ satisfying $a \sim b$. Since $n \geq 3$, there is some element $c \in [n]$ with $c \notin \{a, b\}$. Let \succsim' be a profile which agrees with \succsim on $\{a, b\}$, but where each individual \succsim'_i ranks c immediately below b (by this we mean $a \succsim'_i b \Rightarrow a \succsim'_i b \succsim'_i c$, and $b \succsim'_i a \Rightarrow b \succsim'_i c \succsim'_i a$). Since \succsim' agrees with \succsim on $\{a, b\}$, the IIA axiom tells us $\succsim' = f(\succsim')$ satisfies $a \sim' b$ as well. Similarly, let \succsim'' agree with \succsim on $\{a, b\}$, but suppose each individual \succsim''_i ranks c immediately above b . Letting $\succsim'' = f(\succsim'')$, we once again have $a \sim'' b$. The above construction, together with the Pareto property, tells us that $b \succ' c$, so that $a \sim' b \succ' c$, giving $a \succ' c$. But the Pareto property also implies $c \succ'' b$, so that $c \succ'' b \sim'' a$ and so $c \succ'' a$. But, as is easily verified, \succsim' and \succsim'' agree on $\{a, c\}$. The IIA axiom, then, implies that \succsim' and \succsim'' must agree on $\{a, c\}$, contradicting our calculation above. \square

With this lemma in place, we will henceforth restrict attention to maps $f : \mathcal{P}^k \rightarrow \mathcal{P}$. As alluded to above, we will focus on sets of preferences where the ranking of two alternatives is fixed. For $i, j \in [n]$ with $i < j$, define $U_{ij}^+ = \{\succ \in \mathcal{P} \mid i \succ j\}$ and $U_{ij}^- = \{\succ \in \mathcal{P} \mid j \succ i\}$ (since we are restricting to preferences in \mathcal{P} , we could equally well define $U_{ij}^- = \mathcal{P} \setminus U_{ij}^+$). Although we have only defined these sets for $i < j$, for notational convenience we will take U_{ji}^σ to mean the set $U_{ij}^{\bar{\sigma}}$ (where σ and $\bar{\sigma}$ are opposite signs in $\{+, -\}$). Let $\mathcal{U} = \{U_{ij}^\sigma \mid i < j, \sigma \in \{+, -\}\}$ denote the entire collection of sets U_{ij}^σ . We will be interested in the topology of the nerve of \mathcal{U} (more specifically, its geometric realization $N_{\mathcal{P}}$).

Theorem 3.2.2 (Baryshnikov Proposition 1). *The simplices of maximal dimension in $N_{\mathcal{P}}$ are in one-to-one correspondence with orders in \mathcal{P} ; these simplices have dimension $\frac{(n+1)(n-2)}{2}$, and $N_{\mathcal{P}}$ coincides with the union of these simplices.*

Proof. We know that $\binom{n}{2}$ is the number of ways of choosing a set $\{i, j\}$ of two elements from $[n]$. In a strict linear order \succ , any two elements of $[n]$ must be comparable, so every set $\{i, j\}$ must be accounted for (that is, either U_{ij}^+ or U_{ij}^- must contain \succ). So, provided that the appropriate σ 's are chosen for the sets U_{ij}^σ , the order \succ will be the only element in the intersection of these U_{ij}^σ 's. Therefore \succ will correspond to a simplex of dimension $\binom{n}{2} - 1 = \frac{(n+1)(n-2)}{2}$. The intersection of more than $\binom{n}{2}$ sets must be empty, because some set $\{i, j\}$ will be represented twice (and $U_{ij}^+ \cap U_{ij}^- = \emptyset$). Therefore, the simplices corresponding to orders are of maximal dimension in $N_{\mathcal{P}}$. Finally, note that any nonempty intersection S of less than $\binom{n}{2}$ sets corresponds to a (strict) partial order on $[n]$. But any strict partial order can be extended to a linear order³, so that S belongs to at least one of the simplices of maximal dimension. Hence $N_{\mathcal{P}}$ is the union of the maximal simplices. \square

The correspondence between preference relations and maximal simplices of $N_{\mathcal{P}}$ given by theorem 3.2.2 will be quite useful throughout this section. For ex-

³This is precisely the statement of Szpilrajn's Extension Theorem, which is the main theorem given in Szpilrajn [9].




Figure 3.1: The nerve $N_{\mathcal{U}}$ for $n = 3$.

ample, consider the case $n = 3$; the geometric realization of $N_{\mathcal{U}}$ is given in figure 3.1.

There are six cells of dimension $2 = \frac{(n+1)(n-2)}{2}$, each of which clearly corresponds to a strict order on $[n]$ (we use the notation (pqr) to represent the order $p \succ q \succ r$). Notice that for $n = 3$, the nerve $N_{\mathcal{U}}$ is homotopy equivalence to $S^1 = S^{n-2}$. In fact, this relationship holds for every $n \geq 3$:

Theorem 3.2.3 (Baryshnikov Theorem 1). *The simplicial complex $N_{\mathcal{U}}$ is homotopy equivalent to the $(n - 2)$ -dimensional sphere S^{n-2} .*

Proof. Our strategy is to specify a topological space M (homotopy equivalent to S^{n-2}) and an open cover \mathcal{V} of M with the exact same intersection properties as \mathcal{U} , so that $N_{\mathcal{U}} = N_{\mathcal{V}}$. The cover \mathcal{V} will be chosen so that it satisfies all requirements of the nerve theorem (see section 4.G of Hatcher [7]), so that we may conclude $N_{\mathcal{U}} = N_{\mathcal{V}} \simeq M \simeq S^{n-2}$. We take $M = \mathbb{R}^n \setminus \Delta$, where $\Delta = \{x_1 = \dots = x_n\}$ is the diagonal of \mathbb{R}^n . For $i < j$, define $V_{ij}^+ = \{(x_1, \dots, x_n) \in \mathbb{R}^n \mid x_i > x_j\}$ and $V_{ij}^- = \{(x_1, \dots, x_n) \in \mathbb{R}^n \mid x_i < x_j\}$. It is clear that the collection \mathcal{V} of all such sets V_{ij}^σ covers M , and that a subcollection of the V_{ij}^σ 's has nonempty intersection if and

only if the corresponding sets U_{ij}^σ (just replace the symbol V with U to get the corresponding set) have nonempty intersection. So $N_{\mathcal{V}} = N_{\mathcal{D}}$. It is also clear that the sets V_{ij}^σ are each convex (hence contractible), so that any intersection of them is either empty or contractible as well. Thus we may apply the nerve theorem to conclude $N_{\mathcal{V}} \simeq M$. Finally, M can be shown to be homotopy equivalent to S^{n-2} as follows. First, we may rotate M to get $M \simeq \mathbb{R}^n \setminus \{x_1 = \dots = x_{n-1} = 0\}$, and then perform a linear homotopy to get $M \simeq S^{n-1} \setminus \{(0, \dots, 0, 1), (0, \dots, 0, -1)\}$. A stereographic projection gives $M \simeq \mathbb{R}^{n-1} \setminus \{0\}$, so that another linear homotopy gives $M \simeq S^{n-2}$. \square

Naturally, we expect a similar topological construction to work for the profile space \mathcal{P}^k . Let $\vec{\sigma} = (\sigma_1, \dots, \sigma_k)$ be a vector of signs in $\{+, -\}$, and for $i < j$ define

$$U_{ij}^{\vec{\sigma}} = \{(\succ_1, \dots, \succ_k) \in \mathcal{P}^k \mid \forall m \succ_m \in U_{ij}^{\sigma_m}\} = U_{ij}^{\sigma_1} \times \dots \times U_{ij}^{\sigma_k}.$$

Clearly, the collection of all such $U_{ij}^{\vec{\sigma}}$ covers \mathcal{P}^k . Let $N_{\mathcal{P}^k}$ denote the nerve of this covering. A similar argument to that of Theorem 3.2.2 establishes the following result:

Theorem 3.2.4 (Baryshnikov Proposition 3). *The simplices of maximal dimension in $N_{\mathcal{P}^k}$ are in one-to-one correspondence with profiles of orders in \mathcal{P}^k ; these simplices have dimension $\frac{(n+1)(n-2)}{2}$, and $N_{\mathcal{P}^k}$ coincides with the union of these simplices.*

One might suspect that the nerve corresponding to the profile space is homotopy equivalent to $(S^{n-2})^k$. This is not the case; indeed, it appears that $N_{\mathcal{P}^k}$ is rather complicated. However, at least up to dimension $n-2$, the homology of $N_{\mathcal{P}^k}$ behaves like such a product. We record this result in the next theorem, but omit the proof.

Theorem 3.2.5 (Baryshnikov Proposition 4). *The (singular) homology groups of $N_{\mathcal{P}^k}$ are zero in positive dimensions up to dimension $n-3$, and the $(n-2)$ -th group is isomorphic to \mathbb{Z}^k .*

To complete the transformation of the original social choice model into a topological one, we must determine how a map $f : \mathcal{P}^k \rightarrow \mathcal{P}$ transforms into a map

between nerves $N_{\mathcal{P}^k} \rightarrow N_{\mathcal{P}}$. This is straightforward; the IIA axiom ensures that for each $i < j$ and each $\vec{\sigma}$, f maps the vertex $U_{ij}^{\vec{\sigma}}$ to exactly one vertex of $N_{\mathcal{P}}$ (either U_{ij}^+ or U_{ij}^-). We also know that if $(\succ_1, \dots, \succ_k) \in \mathcal{P}^k$ is a profile (corresponding to some simplex in $N_{\mathcal{P}^k}$), then $f(\succ_1, \dots, \succ_k)$ is a relation \succ (corresponding to some simplex in $N_{\mathcal{P}}$). Thus f , being defined on the vertices of $N_{\mathcal{P}^k}$, may be extended linearly to yield a simplicial map $N_{\mathcal{P}^k} \rightarrow N_{\mathcal{P}}$; we will simply call this map f . Naturally, f is continuous because it is a well-defined simplicial map. We will see that the dictatorship property is closely related to the way f maps generators (basis elements) of $H_{n-2}(N_{\mathcal{P}^k}) \cong \mathbb{Z}^k$ to elements of $H_{n-2}(N_{\mathcal{P}}) \cong \mathbb{Z}$. To understand this relation, we must first give a useful characterization of the generators of $H_{n-2}(N_{\mathcal{P}})$, followed by those of $H_{n-2}(N_{\mathcal{P}^k})$.

Suppose a graph g is a simple unoriented loop of length n with vertices in $[n]$ (for example, g might be the graph $1 - 2 - 3 - \dots - n - 1$). If we choose some orientation for the arrows (yielding an oriented graph \vec{g}), we may identify a corresponding set of vertices U_{ij}^{σ} , where an arrow $1 \rightarrow 2$ corresponds to U_{12}^+ , the arrow $1 \leftarrow 2$ corresponds to U_{12}^- , and so on. This gives a collection of n vertices (one for each edge of g). The simplex $S(\vec{g})$ spanned by these vertices may or may not belong to the nerve $N_{\mathcal{P}}$; indeed, if the oriented graph is $1 \rightarrow 2 \rightarrow \dots \rightarrow n \rightarrow 1$ (we will call such a graph an *oriented cycle*), then the simplex will *not* belong to the nerve (the intersection of the corresponding sets will be empty since \vec{g} induces relations which do not satisfy transitivity). However, such a simplex will belong to the “total” simplex Δ_{tot} (it contains every simplex spanned by any set of vertices in \mathcal{U}). So, if we let the chain $h(\vec{g}) = \partial S(\vec{g})$ denote the boundary of $S(\vec{g})$, we must have that $h(\vec{g})$ is a cycle in Δ_{tot} (that is, $\partial h(\vec{g}) = 0$; this is clear since $\partial^2 = 0$). In fact, $h(\vec{g})$ is a cycle in $H_{n-2}(N_{\mathcal{P}})$, and represents a generator whenever \vec{g} is an oriented loop:

Theorem 3.2.6 (Baryshnikov Proposition 2). *Suppose \vec{g} is an oriented graph on $[n]$ whose unoriented support g is a simple loop of length n . Then the cycle $h(\vec{g})$ is an $(n-2)$ -dimensional cycle of $N_{\mathcal{P}}$ (that is, $h(\vec{g})$ represents some class of $H_{n-2}(N_{\mathcal{P}})$). If \vec{g} is an oriented cycle, then $h(\vec{g})$ represents a generator of $H_{n-2}(N_{\mathcal{P}})$; otherwise, $h(\vec{g})$ is trivial in $H_{n-2}(N_{\mathcal{P}})$.*

Proof. Since $S(\vec{g})$ is spanned by n vertices, $S(\vec{g})$ is a simplex of dimension $n-1$;

hence the boundary $h(\vec{g})$ is a sum of $(n-2)$ -dimensional simplices. Each simplex in this sum corresponds to the simplex given by \vec{g} , but with one arrow removed; hence, the vertices corresponding to each $(n-2)$ -dimensional simplex in the sum give a partial order on $[n]$, which may be extended to a linear order. But then every such simplex must belong to $N_{\mathcal{D}}$, because the simplices of $N_{\mathcal{D}}$ correspond to linear orders on $[n]$. Hence $[h(\vec{g})] \in H_{n-2}(N_{\mathcal{D}})$. Similarly, the order corresponding to \vec{g} may be extended to a linear order whenever \vec{g} is not an oriented cycle. But then $h(\vec{g})$ is the boundary of some maximal simplex in $N_{\mathcal{D}}$, so that $h(\vec{g})$ is trivial in $H_{n-2}(N_{\mathcal{D}})$ if \vec{g} is not an oriented cycle. Finally, suppose \vec{g} is an oriented cycle. For each arrow $i \rightarrow j$ in \vec{g} , we have a half-space $W_{ij}^+ = V_{ij}^+ = \{(x_1, \dots, x_n) \in \mathbb{R}^n \mid x_i > x_j\}$. Consider the collection \mathcal{W} of all such half-spaces given by \vec{g} . Clearly $\cup \mathcal{W} \subseteq M$ (recall $M = \mathbb{R}^n \setminus \Delta$ from the proof of theorem 3.2.3); if there is some $\vec{x} = (x_1, \dots, x_n) \in M$ which does not belong to any of the W_{ij}^+ , then \vec{x} belongs to the complement of each W_{ij}^+ . Supposing (without loss of generality) that \vec{g} is the oriented cycle $1 \rightarrow 2 \rightarrow \dots \rightarrow n \rightarrow 1$, this implies $x_1 \leq x_2 \leq \dots \leq x_n \leq x_1$; that is, we have $\vec{x} = (x, \dots, x) \in \Delta$, contradicting the fact that $\vec{x} \in M$. Hence \mathcal{W} is a cover of M . Since $\mathcal{W} \subset \mathcal{V}$, the nerve theorem yields a commutative diagram

$$\begin{array}{ccc}
 N_{\mathcal{W}} & \xrightarrow{\iota} & N_{\mathcal{V}} \\
 g \searrow & & \swarrow h \\
 & M &
 \end{array}$$

where ι is injection and g and h are homotopy equivalences. Let g^{-1} and h^{-1} denote the homotopy inverses of g and h , respectively. Since $g = h \circ \iota$, we have $h^{-1} \circ g = h^{-1} \circ h \circ \iota \simeq 1 \circ \iota = \iota$. But $h^{-1} \circ g$ is a homotopy equivalence, so that ι is also a homotopy equivalence. Clearly, $[h(\vec{g})]$ generates $H_{n-2}(N_{\mathcal{V}})$. Thus $h(\vec{g}) = \iota(h(\vec{g}))$ represents a generator of $H_{n-2}(N_{\mathcal{V}}) = H_{n-2}(N_{\mathcal{D}})$. \square

The graph-based approach to characterizing generators can also be used to find useful representations of basis elements of $H_{n-2}(N_{\mathcal{D}^k})$. In particular, let $\vec{g} = (\vec{g}_1, \dots, \vec{g}_k)$ be a tuple of oriented graphs on $[n]$ where each \vec{g}_i has the same unoriented loop (of length n) as its unoriented support (say $1-2-\dots-n-1$).

From \vec{g} , we can extract n vertices $U_{12}^{\vec{\sigma}^1}, \dots, U_{n-1,n}^{\vec{\sigma}^{n-1}}, U_{n1}^{\vec{\sigma}^n}$ as follows: for $U_{i,i+1}^{\vec{\sigma}^i}$, define $\vec{\sigma}^i = (\sigma_1^i, \dots, \sigma_k^i)$ by

$$\sigma_\ell^i = \begin{cases} + & \text{if } \vec{g}_\ell \text{ has arrow } i \rightarrow i+1 \\ - & \text{if } \vec{g}_\ell \text{ has arrow } i \leftarrow i+1 \end{cases} .$$

So, we can think of the vertex $U_{i,i+1}^{\vec{\sigma}^i}$ as keeping track of how the k graphs rank i and $i+1$. Notice that any subset of $n-1$ of these vertices gives a simplex in $N_{\mathcal{P}^k}$: fixing some coordinate ℓ , the intersection of $n-1$ sets (for example, $U_{12}^{\sigma_\ell^1}, \dots, U_{n-1,n}^{\sigma_\ell^{n-1}}$) will result in a transitive partial order on $[n]$, because the unoriented support $1-2-\dots-n-1$ implies that an intransitivity can arise only if n arrows (hence n vertices) are present. Thus each component will correspond to a partial order on $[n]$, so that the $(n-2)$ -dimensional simplex spanned by the $n-1$ vertices will belong to $N_{\mathcal{P}^k}$. Summing (with appropriate coefficients) over all possible subsets of size $n-1$ therefore gives the boundary $h(\vec{g})$ of some $(n-1)$ -dimensional simplex; since each simplex in the sum belongs to $N_{\mathcal{P}^k}$, this proves that $[h(\vec{g})] \in H_{n-2}(N_{\mathcal{P}^k})$.

Next, we define projection maps $p_\ell : N_{\mathcal{P}^k} \rightarrow N_{\mathcal{P}}$, for $1 \leq \ell \leq k$. Obviously, for a profile of preferences we must have $p_\ell(\succ_1, \dots, \succ_k) = \succ_\ell$. So, for a vertex $U_{ij}^{\vec{\sigma}}$ of $N_{\mathcal{P}^k}$, we see that the ℓ^{th} projection is $U_{ij}^{\sigma_\ell} \in N_{\mathcal{P}}$. Extending this map linearly gives our projection map $N_{\mathcal{P}^k} \xrightarrow{p_\ell} N_{\mathcal{P}}$, which is clearly simplicial and therefore continuous. To see how p_ℓ maps the cycle $h(\vec{g})$, consider $S = [U_{12}^{\vec{\sigma}^1}, \dots, U_{n-1,n}^{\vec{\sigma}^{n-1}}]$ (that is, S is one of the $(n-2)$ -dimensional simplices in the sum $h(\vec{g})$). Then $p_\ell(S) = [U_{12}^{\sigma_\ell^1}, \dots, U_{n-1,n}^{\sigma_\ell^{n-1}}]$, an $(n-2)$ -dimensional simplex in $N_{\mathcal{P}}$. The projections of other simplices in $h(\vec{g})$ are computed in a similar fashion, so that $p_\ell(h(\vec{g}))$ is a sum of $(n-2)$ -dimensional simplices in $N_{\mathcal{P}}$. By construction, $p_\ell(h(\vec{g}))$ is a cycle, and therefore belongs to some class of $H_{n-2}(N_{\mathcal{P}})$. Note that $p_\ell(h(\vec{g})) = h(\vec{g}_\ell)$.

Theorem 3.2.7 (Baryshnikov Proposition 5). *Suppose $\vec{g} = (\vec{g}_1, \dots, \vec{g}_k)$ is a tuple of oriented graphs where \vec{g}_ℓ is the oriented cycle $1 \rightarrow 2 \rightarrow \dots \rightarrow n \rightarrow 1$, and all other \vec{g}_i are acyclic with the same unoriented support $1-2-\dots-n-1$. Define $h_\ell = h(\vec{g}_\ell)$. Then the set $\{h_1, \dots, h_k\}$ represents a basis for $H_{n-2}(N_{\mathcal{P}^k}) \cong \mathbb{Z}^k$.*

Proof. Consider h_ℓ . Since each graph \vec{g}_i ($i \neq \ell$) is acyclic, each such $h(\vec{g}_i)$ is trivial in $H_{n-2}(N_{\mathcal{P}})$, and since \vec{g}_ℓ is an oriented cycle, $h(\vec{g}_\ell)$ represents a generator of

$H_{n-2}(N_{\mathcal{D}}) \cong \mathbb{Z}$. This means h_ℓ projects to 0 under p_i ($i \neq \ell$) and to a generator under p_ℓ . So, the collection $\{h_1, \dots, h_k\}$ represents a basis for $H_{n-2}(N_{\mathcal{D}^k})$. \square

With these characterizations in place, we are now prepared to establish the relationship between dictators and generators. We begin with the following lemma:

Lemma 3.2.8. *Suppose $\vec{g} = (\vec{g}_1, \dots, \vec{g}_k)$ has an associated cycle $h(\vec{g})$ representing a generator h_ℓ in $H_{n-2}(N_{\mathcal{D}^k})$ (so \vec{g}_ℓ is an oriented cycle with unoriented support $1 - 2 - \dots - n - 1$). Then $f_*(h(\vec{g})) \in H_{n-2}(N_{\mathcal{D}})$ and corresponds to a graph \vec{g}_f with the same unoriented support.*

Proof. By IIA, the map $f : N_{\mathcal{D}^k} \rightarrow N_{\mathcal{D}}$ sends a vertex $U_{ij}^{\vec{\sigma}}$ to U_{ij}^σ for some $\sigma \in \{+, -\}$. We also know that $h(\vec{g}) \in H_{n-2}(N_{\mathcal{D}^k})$ is a sum of $(n-2)$ -dimensional simplices of the form $[U_{12}^{\vec{\sigma}^1} \dots U_{n1}^{\vec{\sigma}^n}]$ (where one of the vertices $U_{ij}^{\vec{\sigma}^i}$ is dropped). Therefore $f_*(h(\vec{g}))$ is a sum of simplices of the form $[U_{12}^{\sigma_1} \dots U_{n1}^{\sigma_n}]$ where each $\sigma_i \in \{+, -\}$ and, again, some vertex $U_{ij}^{\sigma_i}$ is dropped. Each of these simplices is an $(n-2)$ -dimensional simplex belonging to $H_{n-2}(N_{\mathcal{D}})$ (because such a simplex will represent a partial order on $[n]$, and therefore belong to one of the maximal simplices), and so $f_*(h(\vec{g})) \in H_{n-2}(N_{\mathcal{D}})$. Notice also that $f_*(h(\vec{g}))$ will have vertices $[U_{12}^{\sigma_1} \dots U_{n1}^{\sigma_n}]$, so that the corresponding graph \vec{g}_f will have unoriented support $1 - 2 - \dots - n - 1$. \square

An immediate consequence of this lemma is that $f_*(h(\vec{g}))$ is either a generator or is trivial in $H_{n-2}(N_{\mathcal{D}})$, according to whether \vec{g}_f is an oriented cycle or not. In fact, these two cases correspond to whether or not agent ℓ is a dictator:

Theorem 3.2.9 (Baryshnikov Proposition 7). *Agent ℓ is a dictator if and only if $f_*(h_\ell)$ is a generator.*

Proof. Suppose first that agent ℓ is a dictator. By the previous lemma, we may assume that h_ℓ is represented by a cycle $h(\vec{g})$, where $\vec{g} = (\vec{g}_1, \dots, \vec{g}_k)$ is a tuple of graphs with unoriented support $1 - 2 - \dots - n - 1$, \vec{g}_ℓ is an oriented cycle, and all other \vec{g}_i are acyclic. It will suffice to show that the graph \vec{g}_f corresponding to $f_*(h(\vec{g}))$ is an oriented cycle. Consider alternatives i and $i+1$ in $[n]$. Corresponding to \vec{g} is the vertex $U_{i,i+1}^{\vec{\sigma}^i}$ in $N_{\mathcal{D}^k}$; in particular, we must have $\sigma_\ell^i = +$ since \vec{g}_ℓ has the arrow $i \rightarrow i+1$. Since ℓ is a dictator, this implies f maps $U_{i,i+1}^{\vec{\sigma}^i}$ to $U_{i,i+1}^+$, which means

\vec{g}_f must have the arrow $i \rightarrow i + 1$. Since i was chosen arbitrarily, this means \vec{g}_f is actually equal to the graph \vec{g}_ℓ , so that \vec{g}_f is an oriented cycle. Hence $f_*(h_\ell)$ is a generator.

Next, suppose $f_*(h_\ell)$ is a generator. Take any distinct $i, j \in [n]$; it will suffice to show that f maps $U_{ij}^{\vec{\sigma}}$ to $U_{ij}^{\sigma_\ell}$ for any vector $\vec{\sigma}$ of signs. For concreteness, suppose $\sigma_\ell = +$. Obviously there is some oriented cycle (call it \vec{g}_ℓ) of length n on $[n]$ which has the arrow $i \rightarrow j$. This allows us to represent h_ℓ by a tuple of graphs $\vec{g} = (\vec{g}_1, \dots, \vec{g}_\ell, \dots, \vec{g}_k)$, where each \vec{g}_i ($i \neq \ell$) is acyclic and has the same unoriented support as \vec{g}_ℓ . Suppose we modify one arrow in one of the \vec{g}_i 's ($i \neq \ell$) in such a way that \vec{g}_i remains acyclic (since $n \geq 3$, this can obviously be done). This means one agent has reversed his preference between i and j . By the IIA axiom, the resulting social preference will either be unchanged or it will reverse the order of i and j ; thus at most one arrow of \vec{g}_f will be reversed, and all others will remain the same. By hypothesis, however, \vec{g}_f must be an oriented cycle (the modified tuple \vec{g} still represents h_ℓ , which means \vec{g}_f must still represent a generator). If one arrow of \vec{g}_f were reversed, then \vec{g}_f would no longer be a cycle; thus \vec{g}_f is actually unaffected by changing one arrow in one \vec{g}_i ($i \neq \ell$). Because of this, we may (without affecting \vec{g}_f) modify the tuple \vec{g} , one arrow at a time, until each \vec{g}_i has the arrow $i \rightarrow j$ (the obvious procedure is to just reverse any $i \leftarrow j$ arrows; if doing so causes an oriented cycle, some other arrow in the same graph will need to be reversed beforehand). Corresponding to this modified \vec{g} is the vertex $U_{ij}^{(+, \dots, +)}$ which, by the Pareto property, must be mapped to U_{ij}^+ by f . This means \vec{g}_f has the arrow $i \rightarrow j$. As the above argument shows, any vector $\vec{\sigma}$ can be achieved (with $\sigma_\ell = +$) without changing any arrows of \vec{g}_f ; in particular, the arrow $i \rightarrow j$ of \vec{g}_f is fixed, so that for any vector $\vec{\sigma}$, f will map $U_{ij}^{\vec{\sigma}}$ to U_{ij}^+ . Thus ℓ is a dictator. \square

In view of theorem 3.2.9, Arrow's theorem will be immediate if we can prove the existence of some h_ℓ which f_* maps to a generator. We will give a simple algebraic proof that this is, indeed, the case.

Consider a map $N_{\mathcal{D}} \rightarrow N_{\mathcal{D}^k}$ which sends a vertex U_{ij}^σ to $U_{ij}^{(\sigma, \dots, \sigma)}$. We can extend this map linearly to get a simplicial map $D : N_{\mathcal{D}} \rightarrow N_{\mathcal{D}^k}$ (the diagonal map). Clearly, $D_*(h) = (h, \dots, h)$ for any generator $h \in H_{n-2}(N_{\mathcal{D}})$. Consider the following diagram:

$$\begin{array}{ccccc}
\overbrace{H_{n-2}(N_{\mathcal{P}})}^{\mathbb{Z}} & \xrightarrow{D_*} & \overbrace{H_{n-2}(N_{\mathcal{P}^k})}^{\mathbb{Z}^k} & \xrightarrow{f_*} & \overbrace{H_{n-2}(N_{\mathcal{P}})}^{\mathbb{Z}} \\
& & \nearrow^{i_1} & & \nwarrow_{i_k} \\
\overbrace{H_{n-2}(N_{\mathcal{P}})}^{\mathbb{Z}} & & \dots & & \overbrace{H_{n-2}(N_{\mathcal{P}})}^{\mathbb{Z}}
\end{array}$$

Here, the maps i_ℓ are canonical injections $1 \mapsto (0, \dots, 1, \dots, 0)$, where the 1 is in the ℓ^{th} spot. We call such a setup *purely separated* (in dimension $n-2$). By lemma 3.2.8, $f_* \circ i_\ell$ sends a generator of $H_{n-2}(N_{\mathcal{P}})$ to either 0 or 1. Notice also that the Pareto property implies that $f \circ D$ is the identity map, so that $f_* \circ D_* = (f \circ D)_*$ is the identity map as well. So, we have

$$\begin{aligned}
f_* \circ D_*(1) &= f_*(1, \dots, 1) \\
&= f_*(1, 0, \dots, 0) + \dots + f_*(0, \dots, 0, 1) \\
&= \sum_{j=1}^k f_*(i_j(1)) \\
&= \sum_{j=1}^k d_j \cdot 1,
\end{aligned}$$

where each d_j is either 0 or 1. But since $f_* \circ D_*$ is the identity map, we must have $\sum_{j=1}^k d_j \cdot 1 = 1$; this implies that exactly one d_j is 1, which means $f_*(h_j)$ is a generator. By theorem 3.2.9, then, this agent is a dictator, and the proof of Arrow's theorem is complete.

This is obviously not the easiest way to prove Arrow's theorem (indeed, the short proof we gave in section 2 is much simpler). The advantage of this topological approach, however, is that it gives at least some intuition for why the theorem is true. Namely, the domain of preferences, together with the IIA axiom, yields a nontrivial topological space, and this nontriviality seems to be responsible for the dictatorship outcome. If one wishes to identify restricted domains where non-dictatorial maps are possible, then the topological model can be used to generate hypotheses about which domains might be suitable. Consider again the diagram of $N_{\mathcal{P}}$ for the case $n=3$. Since dictators map to generators (more or less), one can try to remove the possibility of dictatorship by removing enough orders from \mathcal{P}

to make the nerve contractible. Looking at the diagram, we see that the space will be contractible if we remove the simplices corresponding to the orders $3 \succ 1 \succ 2$ and $1 \succ 3 \succ 2$. The remaining preferences coincide with the single-peaked preferences described in section 2. Remarkably, any contractible subcomplex of $N_{\mathcal{D}}$ will correspond to a (sub)domain of single-peaked preferences! This makes the single-peaked condition almost obvious, whereas the classical framework does not provide any obvious intuition whatsoever as to why single-peaked domains may escape the dictatorship result. Note that our topological framework does not *prove* that single-peaked domains permit non-dictatorial aggregation maps; the relationship “ ℓ is a dictator $\Leftrightarrow f_*(h_\ell)$ is a generator” is built on the premise that the domain is unrestricted, and it is not at all clear that this relationship could hold without that assumption. Still, the topology offers enough intuition that it can be a useful aid in formulating possible candidates for restricted domains. Moreover, the topological construction is, on its own, a very interesting approach to this branch of theory.

Chapter 4

Judgment Aggregation

Judgment aggregation is a new branch of economic theory which studies the way individual *choices* (not necessarily preferences) may be aggregated into a single collective choice. One usually begins with some finite set A containing n alternatives and some subset $X \subseteq P(A)$ of “acceptable choices”. There are k agents, and each selects some member of X (one usually represents members of X as vectors $x = (x_1, \dots, x_n)$ with entries in $\{0, 1\}$, indicating which members of A have been selected).

An aggregation map $f : X^k \rightarrow X$ assigns to each k -tuple $\vec{x} = (x^1, \dots, x^k)$ of X^k a vector $\vec{y} = f(\vec{x})$ of X . If, for every ℓ , we have that $y_\ell = 1$ (resp. 0) whenever $x_\ell^i = 1$ (resp. 0) for each $i = 1, \dots, k$, then f satisfies the *Pareto property*. The IIA axiom can also be adapted to this framework: suppose that for each $\ell = 1, \dots, n$ and any two profiles $\vec{y}, \vec{z} \in X^k$ where $y_\ell^i = z_\ell^i$ for each i , we have that $f(\vec{y})$ and $f(\vec{z})$ agree on the ℓ^{th} coordinate (that is, the collective decision regarding the ℓ^{th} element of A depends only on the individual judgments of the ℓ^{th} element). Then f is *Independent of Irrelevant Alternatives*. Finally, f is *Dictatorial* if there is some agent d such that for every $\vec{x} \in X^k$, $f(\vec{x}) = x^d$ (that is, the aggregated choice always coincides with the choice of individual d , regardless of what others choose).

This framework is clearly very similar to that of the social choice setup. In fact, social choice models are a special case of the judgment aggregation model. To see this, suppose \mathcal{R} is the set of rational preference relations over $[m]$. Since a preference relation \succsim is determined by comparisons between pairs of elements

Table 4.1: The doctrinal paradox

	<i>P</i>	<i>Q</i>	<i>R</i>
Judge 1	T	T	T
Judge 2	T	F	F
Judge 3	F	T	F
Σ	T	T	F

of $[m]$, we can take the set A to be the set of all possible pairwise comparisons; that is, $A = \{i \succ j : i, j \in [m], i \neq j\}$. Then X will be the set of vectors for which the corresponding chosen pairwise relations give a preference relation on $[m]$. It is easy to see that the Pareto and IIA definitions of the judgment aggregation model agree with those of the social choice model; in the case of strict preferences, the Dictatorship axioms agree as well (recall that the dictatorship axiom of social choice only requires that the *strict* part of the aggregated preference agrees with that of the dictator; in the judgment aggregation setup, the aggregated preference must *equal* the dictator’s preference).

The social choice case is a good example of why only certain subsets of A may be admissible. Another possibility is that elements of A may represent logical propositions, and formal relationships between the propositions may force some subsets to be excluded on the grounds of logical consistency. The canonical example is the “doctrinal paradox”. Suppose three judges (J1, J2, and J3) must decide whether a defendant is guilty of a series of crimes (P , Q , and R). Due to legal definitions, one is guilty of R if and only if one is guilty of both P and Q . In this scenario, then, we have $A = \{P, Q, R\}$ and $X = \{(0, 0, 0), (1, 0, 0), (0, 1, 0), (1, 1, 1)\}$, since these are the only logically consistent assignments. To see the difficulty in aggregating the choices of J1, J2, and J3, suppose the aggregation map f is simply a propositionwise majority vote. Then two out of three judges will find the defendant guilty of P (and similarly for Q), but only one of the three will find the defendant guilty of R . But then the aggregated judgment is $(1, 1, 0)$, which is absurd. Situations such as this make the problem of judgment aggregation considerably more delicate than that of preference aggregation (incidentally, this example is also what motivated the name “judgment aggregation” for this branch of theory).

Despite these difficulties, the problem of characterizing all *impossibility domains* (those domains X for which every aggregation map $f : X^k \rightarrow X$ satisfying the IIA and Pareto properties is Dictatorial) has recently been resolved. Specifically, Dokow and Holzman [5] have shown that X is an impossibility domain if and only if X is *totally blocked* and is not an affine subspace (over \mathbb{F}_2). We will not discuss this result (nor the definition of total blockedness), but remark that their result has been helpful for identifying some simple examples for use in this section.

Our objective is to explore the potential for understanding such generalized impossibility results in a topological fashion, just as Baryshnikov has done for Arrow's theorem in social choice. Keeping in mind the judgment aggregation form of preference aggregation, we proceed as follows: given $A = [n]$ (we can make this identification since A is assumed finite) and $X \subseteq \{0, 1\}^n$, we define for each $i \in A$ and $\sigma \in \{+, -\}$ the sets U_i^σ by

$$U_i^+ = \{x = (x_1, \dots, x_n) \in X \mid x_i = 1\} \quad \text{and} \quad U_i^- = \{x = (x_1, \dots, x_n) \in X \mid x_i = 0\} .$$

The analogy with Baryshnikov's approach is clear; indeed, the above sets coincide in a 1-1 fashion with the sets U_{ij}^σ in Baryshnikov's model when A is the set of all strict relations $i \succ j$. Naturally, we are interested in the collection \mathcal{U} of all such sets U_i^σ , and the corresponding nerve N_X .

Notice that a vector $(x_1, \dots, x_n) \in X$ will be the only element in the intersection of sets $U_i^{\sigma_i}$, where $\sigma_i = +$ if and only if $x_i = 1$, and that any intersection of more than n sets must be empty (because then some index i is repeated, so that sets U_i^+ and U_i^- are in the family; but obviously these two sets are disjoint). This means the simplices of maximal dimension in N_X will be in one-to-one correspondence with vectors in X , and that these simplices will have dimension $n - 1$ (being simplices spanned by n vertices). In a similar fashion, we may define sets $U_i^{\vec{\sigma}}$ (where $\vec{\sigma}$ is a vector of k signs in $\{+, -\}$) by

$$U_i^{(\sigma_1, \dots, \sigma_k)} = U_i^{\sigma_1} \times \dots \times U_i^{\sigma_k} .$$

The nerve corresponding to the collection of all such sets $U_i^{\vec{\sigma}}$ is denoted N_{X^k} . The maximal simplices of N_{X^k} correspond to k -tuples of vectors in X , and therefore are

also of dimension $n - 1$. Finally, any aggregation map $f : X^k \rightarrow X$ satisfying the IIA axiom induces a simplicial map $N_{X^k} \rightarrow N_X$ (which we also call f), just as in Baryshnikov's model.¹

At this stage, one encounters the first major difficulty in constructing a general topological theory of judgment aggregation: using this setup, not every impossibility domain X yields a nontrivial nerve N_X . For example, take

$$X_I = \{(0, 0, 0, 0), (0, 1, 0, 1), (1, 0, 0, 1), (1, 1, 1, 1)\}.$$

This is Example 4 in Dokow and Holzman [5], and it is easily shown (using their theorem) that X_I is an impossibility domain. However, the nerve corresponding to this domain is contractible, and therefore of no use in our setting. A similar problem occurs if one tries to use our procedure for the case of preference aggregation when preferences need not be strict (that is, when the entire domain \mathcal{R} of preference relations is permitted). In that case, A consists of all possible pairwise comparisons $i \succsim j$; in particular, the vector $(1, \dots, 1)$ (corresponding to the indifference relation $1 \sim 2 \sim \dots \sim m$) will belong to every set U_i^+ ; in general, the possibility of indifferences add enough vectors to each set U_i^σ that the resulting nerve is trivial (due to the increased number of nonempty intersections). This is why Baryshnikov's model considers only the case of strict preferences. Can a similar domain restriction be found for judgment aggregation models? This is probably a difficult task: even if one removes $(1, 1, 1, 1)$ or $(0, 0, 0, 0)$ (or both) from X_I , the resulting nerve is still trivial. Recently, however, Tanaka [10] has modified Baryshnikov's approach by adding vertices \bar{U}_{ij} containing those orders where the relation $i \sim j$ is *not* satisfied. In the judgment aggregation model, the set \bar{U}_{ij} would consist of those vectors where (for some particular s, t) $s = t = 1$ is not satisfied. This certainly yields a nontrivial nerve (indeed, Tanaka only studies the case where there are three elements to be ordered, but the result is a somewhat complicated space).

¹There is a small subtlety here. The IIA axiom tells us that a set U_i^σ is mapped to either U_i^+ or U_i^- . This suggests the following problem: suppose $X = \{(0, 0), (1, 1)\}$ and that $k = 2$. What if $U_1^{(+,-)} \xrightarrow{f} U_1^+$ and $U_2^{(+,-)} \xrightarrow{f} U_2^-$? The problem is that $U_1^+ \cap U_2^- = \emptyset$, but $U_1^{(+,-)} \cap U_2^{(+,-)} = \{(1, 1), (0, 0)\}$ is nonempty, so that f cannot be a simplicial map. But, by hypothesis, f maps every vector in X^k to some element of X . So it cannot be the case that $U_1^{(+,-)} \xrightarrow{f} U_1^+$ and $U_2^{(+,-)} \xrightarrow{f} U_2^-$; f will always send intersecting vertices to intersecting vertices.

It seems plausible that a similar approach could be successful in judgment aggregation models. However, by introducing additional vertices, the overall study of the problem becomes significantly more complicated; in particular, we lose the simple correspondence between maximal simplices and vectors of X , which will make it difficult to characterize homologically trivial chains. Although we do not develop these ideas in this thesis, this nonetheless seems like a (possibly) fruitful avenue for future research.

We now turn our attention to a simple impossibility domain which *does* yield a nontrivial nerve. Our examination of this space serves as a prototype for a (possibly) more general theory. Let A be a finite set with n elements ($n \geq 3$), and take X to be those subsets of A consisting of only one element. This means we can represent members of X by the “basis” vectors $(1, 0, \dots, 0), (0, 1, 0, \dots, 0), \dots, (0, \dots, 0, 1)$. Although we will not give a proof here, the theorem of Dokow and Holzman can easily be applied to verify that X is, in fact, an impossibility domain.

For each i , we have

$$U_i^+ = \{(0, \dots, 0, 1, 0, \dots, 0)\} \text{ and } U_i^- = X \setminus U_i^+,$$

where the 1 is in the i^{th} position of the vector. Notice that $U_i^+ = \bigcap_{j \in [n] \setminus i} U_j^-$. Obviously, the intersection over all sets U_j^- is empty (X does not contain the zero vector), but (as the above expression for U_i^+ shows) any collection of $n - 1$ of the U_j^- intersect. This means the nerve N_X will consist of the boundary of the $(n - 1)$ -dimensional simplex $[U_1^- \dots U_n^-]$, along with a collection of n simplices (corresponding to the vectors in X), each of which is $(n - 1)$ -dimensional and has exactly one face in common with $[U_1^- \dots U_n^-]$. This is illustrated in figure 4.1 for the case $n = 3$.

Clearly, the nerve N_X is homotopy equivalent to the sphere S^{n-2} , and the boundary $\partial[U_1^- \dots U_n^-]$, being a sum of $(n - 2)$ -dimensional simplices, is a cycle representing a generator of $H_{n-2}(N_X) \cong \mathbb{Z}$.

Our goal, of course, is to achieve a purely separated setup in dimension $n - 2$. Rather than attempting to compute $H_{n-2}(N_{X^k})$, however, we will achieve this by focusing on the topology of N_X instead. This is because (as we have seen in Baryshnikov’s model) the homology of N_{X^k} might be very complicated, and it




Figure 4.1: The nerve N_X for $X = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$.

seems that one can only perform the relevant computations by constructing a clever model for N_{X^k} (by this, we mean a collection of sets with the same intersection properties as the $U_i^{\bar{\sigma}}$, but which satisfy the requirements of the nerve theorem and cover a space M for which the homology can be computed). Any theory which aims to work for a wide array of impossibility domains should, ideally, not require the formulation of a clever new covering every time a new impossibility domain presents itself.

In fact, it is quite easy to obtain a purely separated setup without studying the topology of N_{X^k} in depth. First, we define projection maps $p_\ell : N_{X^k} \rightarrow N_X$ ($\ell = 1, \dots, k$) similar to those in section 3.2. Namely, we begin with the projection map $(x^1, \dots, x^k) \xrightarrow{p_\ell} x^\ell$ and extend the map linearly to yield a simplicial map $N_{X^k} \rightarrow N_X$. Consider a chain $S_1 = \partial[U_1^{\bar{\sigma}^1} \dots U_n^{\bar{\sigma}^n}]$ satisfying $p_1(S_1) = \partial[U_1^- \dots U_n^-]$, and where all other p_i send S_1 to trivial chains of $H_{n-2}(N_X)$ (such a chain S_1 can be constructed by taking $\sigma_1^j = -$ for each $j = 1, \dots, n$, and for each $\ell \neq 1$ choosing one value j^* for which $\sigma_\ell^{j^*} = +$ and setting all other $\sigma_\ell^j = -$). Since $\partial[U_1^- \dots U_n^-]$ is a generator of $H_{n-2}(N_X) \cong \mathbb{Z}$, this means that $(p_\ell)_*$ sends the class $[S_1] \in H_{n-2}(N_{X^k})$ to $1 \in \mathbb{Z}$,

and that all other projections send S_1 to 0. Similarly, we can define chains S_ℓ for each $\ell = 2, \dots, k$ satisfying $p_\ell(S_\ell) = 1$ and $p_j(S_\ell) = 0$ for all $j \neq \ell$. Thus, the classes $[S_1], \dots, [S_k]$ form a basis for a subgroup G of $H_{n-2}(N_{X^k})$, with $G \cong \mathbb{Z}^k$.

Next, observe that $f_*([S_\ell]) \in \{0, 1\}$ for each ℓ . This is because the IIA axiom implies that for each i , $f(U_i^{\vec{\sigma}}) = U_i^\sigma$ for some $\sigma \in \{0, 1\}$. Since $S_\ell = \partial[U_1^{\vec{\sigma}^1} \dots U_n^{\vec{\sigma}^n}]$, this means either $f(S_\ell) = \partial[U_1^- \dots U_n^-]$ (so $f_*(S_\ell) = 1$) or $f(S_\ell) = \partial[U_1^{\sigma_1} \dots U_n^{\sigma_n}]$, where exactly one σ_i is $+$ and all others are $-$ (so that $f_*([S_\ell]) = 0$). The simplicial map $f : N_{X^k} \rightarrow N_X$ induces a homomorphism $H_{n-2}(N_{X^k}) \rightarrow H_{n-2}(N_X)$, which obviously restricts to a homomorphism $G \rightarrow H_{n-2}(N_X)$ since G is a subgroup of $H_{n-2}(N_{X^k})$ (we will still call this restricted homomorphism f_*). Applying the Pareto property, we see that $f_*(1, \dots, 1) = 1$.² As in Baryshnikov's model, this is enough to imply that $1 = \sum_{\ell=1}^k f_*([S_\ell]) = \sum_{\ell=1}^k d_\ell$, where each $d_\ell \in \{0, 1\}$, so that there is exactly one index, d , for which $f_*([S_\ell]) = 1$.

To see that this agent d is a dictator, we need to show that an agent ℓ is a dictator if $f_*([S_\ell]) = 1$. In fact, the converse also holds:

Theorem 4.0.10. *Agent ℓ is a dictator if and only if $f_*([S_\ell]) = 1$.*

Proof. Suppose first that ℓ is a dictator. By definition, $S_\ell = \partial[U_1^{\vec{\sigma}^1} \dots U_n^{\vec{\sigma}^n}]$, where each vector $\vec{\sigma}^i$ has a $-$ sign in the ℓ^{th} coordinate. Since ℓ is a dictator, f maps $U_1^{\vec{\sigma}^1}$ to $U_1^{\sigma_1} = U_1^-$; similarly, f maps each $U_i^{\vec{\sigma}^i}$ to U_i^- . Thus f maps S_ℓ to $\partial[U_1^- \dots U_n^-]$, which represents a generator of $H_{n-2}(N_X)$. Thus $f_*([S_\ell]) = 1$.

Conversely, suppose f maps S_ℓ to $\partial[U_1^- \dots U_n^-]$. We prove that for each i , f maps $U_i^{\vec{\sigma}^i}$ to U_i^- whenever $\sigma_\ell^i = -$. Given $S_\ell = \partial[U_1^{\vec{\sigma}^1} \dots U_n^{\vec{\sigma}^n}]$, we can modify the " j^{th} coordinate" ($j \neq i$) as follows: given $U_i^{\vec{\sigma}^i}$, we can reverse the sign of σ_j^i ; if σ_j^i was $+$, then the j^{th} coordinate of all other vectors is $-$, so we choose some $m \neq i$ and set $\sigma_j^m = +$ (this way, the j^{th} coordinate still projects to a trivial chain of $H_{n-2}(N_X)$). If σ_j^i was $-$, then some other σ_j^m was $+$; so we change that σ_j^m to $-$ (again, this ensures that the j^{th} coordinate projects to 0). These changes give a modified chain S'_ℓ with $p_\ell(S'_\ell) = \partial[U_1^- \dots U_n^-]$, and all other projections p_j send

²There is a technicality here which needs to be addressed. In order to apply the Pareto property to reach this conclusion, we need to know that in $H_{n-2}(N_{X^k})$, we have $[S_1] + \dots + [S_n] = [\partial[U_1^{\vec{\sigma}} \dots U_n^{\vec{\sigma}}]]$, where $\vec{\sigma} = (-, \dots, -)$. That is, the sum of the basis classes should coincide with the class represented by this *particular* chain which projects to 1 in each coordinate. This does not seem unreasonable. We will discuss this (and a related technical problem) later.

S'_ℓ to a trivial chain of $H_{n-2}(N_X)$. By hypothesis, then, f maps S'_ℓ to the generator $\partial[U_1^- \dots U_n^-]$ as well. Notice that for U_i , the only difference between S_ℓ and S'_ℓ is that the j^{th} coordinate of $\vec{\sigma}^i$ has switched signs; but the image under f remains the same. Repeating this argument, we see that any vector $\vec{\sigma}^i$ can be achieved (with $\sigma_\ell^i = -$) without changing the image of $U_i^{\vec{\sigma}}$ under f . In particular, the vector $(-, \dots, -)$ is attainable, and by the Pareto property, f maps $U_i^{(-, \dots, -)}$ to U_i^- . Thus, for all i , we have that $f(U_i^{\vec{\sigma}^i}) = -$ whenever $\sigma_\ell^i = -$. This means that for any k -tuple $(x^1, \dots, x^k) \in X^k$ with $y = f(x^1, \dots, x^k)$, we have that $x_i^\ell = 0 \Rightarrow y_i = 0$. But by choice of our domain X , this forces y to equal x^ℓ , because each vector in X has exactly one coordinate equal to one, and all others zero. Hence ℓ is a dictator. \square

Combining this theorem with the result above (namely, that there is some agent d for which $f_*([S_d]) = 1$), this gives a topological proof that X is an impossibility domain.

Remark. The proof of this theorem depends on the fact that any two chains α, β of the form $\alpha = \partial[U_1^{\vec{\sigma}^1} \dots U_n^{\vec{\sigma}^n}], \beta = \partial[U_1^{\vec{\tau}^1} \dots U_n^{\vec{\tau}^n}]$ (where $p_\ell(\alpha) = p_\ell(\beta) = \partial[U_1^- \dots U_n^-]$ and all other p_m send α and β to trivial chains of $H_{n-2}(N_X)$) must, in fact, belong to the class $[S_\ell]$ (this is needed to ensure that f maps the modified chain S'_ℓ to a generator). This does not seem unreasonable at all; in fact, it would follow easily if we knew that $H_{n-2}(N_{X^k}) \cong \mathbb{Z}^k$ (because then each projection would send the class $[\alpha] - [\beta]$ to 0, forcing $[\alpha] - [\beta] = 0$). Despite this, we have not been successful in finding a rigorous proof. Note that this is essentially the same difficulty mentioned previously, where we need to be able to represent the class $[S_1] + \dots + [S_k]$ by the chain $\partial[U_1^{\vec{\sigma}} \dots U_n^{\vec{\sigma}}]$, where $\vec{\sigma} = (-, \dots, -)$. There seem to be two main options. First, one might attempt to give a general proof that, for nerves constructed in this manner, one always has $H_m(N_{X^k}) \cong (H_m(N_X))^k$, where $H_m(N_X)$ is the top homology group of N_X . Such a theorem would be quite valuable since, as previously discussed, the homology of N_{X^k} is difficult to compute directly, but only information about the top group seems to be required to establish dictatorship results. Alternatively, one might attempt to construct an algorithm yielding a sequence of chains $\gamma_1, \dots, \gamma_t$ (each trivial in $H_m(N_{X^k})$) with $\beta = \alpha + \gamma_1 + \dots + \gamma_t$. We regard both approaches as avenues for future research (indeed, either result

would be a significant feature of a generalized topological theory of judgment aggregation).

Note that for our simple space X , we do not rely on any graph-based or other approaches to identifying generators. This is because the generator $\partial[U_1^- \dots U_n^-]$ is the only one needed to establish the relationship between dictatorship and mapping properties of f_* , and no extra language or notation needs to be introduced to describe it. In general, however, there may not be a simple way to represent the relevant generators in an arbitrary domain X . Finding a convenient way of describing generators, just as Baryshnikov has done for the case of strict preferences, seems to be another challenge that must be overcome.

At this stage, there is also nothing to rule out the possibility that some domains X may have associated nerves with a more complicated topology. This, in turn, could cause the algebraic portion of the proof to fail. What is the appropriate way to generalize the purely-separated setup in such situations? How will mapping properties of the various generators relate to the dictatorship condition? These questions, and the issues previously mentioned, are quite difficult. Still, the analysis given above of the simple space X is encouraging enough that further research into these problems seems worthwhile, and that (eventually) a more general theory might be found.

Bibliography

- [1] K. Arrow. *Social choice and individual values*. Yale Univ Pr, 1963. → pages 6
- [2] Y. Baryshnikov. Unifying Impossibility Theorems: A Topological Approach. *Advances in Applied Mathematics*, 14(4):404–415, 1993. → pages 11, 18
- [3] G. Chichilnisky. Social Choice and the Topology of Spaces of Preferences. *Advances in Mathematics*, 37(2):165–176, 1980. ISSN 0001-8708. → pages 10
- [4] G. Chichilnisky and G. Heal. Necessary and Sufficient Conditions for a Resolution of the Social Choice Paradox. *Journal of Economic Theory*, 31(1):68–87, 1983. ISSN 0022-0531. → pages 10, 11, 12
- [5] E. Dokow and R. Holzman. Aggregation of binary evaluations. *Journal of Economic Theory*, 145(2):495–511, 2010. → pages 31, 32
- [6] J. Geanakoplos. Three brief proofs of arrows impossibility theorem. *Economic Theory*, 26(1):211–215, 2005. → pages 7
- [7] A. Hatcher. *Algebraic topology*. Cambridge: Cambridge University Press, 2002. → pages 14, 20
- [8] E. Spanier. *Algebraic topology*. Springer Verlag, 1994. → pages 13
- [9] E. Szpilrajn. Sur l'extension de l'ordre partiel. *Fund. Math*, 16:386–389, 1930. → pages 19
- [10] Y. Tanaka. A topological approach to the arrow impossibility theorem when individual preferences are weak orders. *Applied mathematics and computation*, 174(2):961–981, 2006. → pages 32