Topology and Social Choice Theory*

Rafe H. Kinsey
rkinsey@stanford.edu

Sunday, June 8, 2008

This thesis provides an exposition of results that use topological techniques in social choice theory. We first introduce traditional social choice theory, including a proof of Arrow's Impossibility Theorem, and outline the basics of topology that we will need. Then we present Chichilnisky's fundamental results in continuous social choice theory. Finally, we present Baryshnikov's topological proof of Arrow's Theorem.

1 Traditional Social Choice Theory

Social choice theory is the mathematical study of formal mechanisms that aggregate individuals' preferences into one group preference. The most obvious example is democratic elections, where majority voting is used to aggregate voters' preferences over candidates into one selected candidate, the winner. Other examples include matching mechanisms such as those used to assign doctors to residencies at hospitals around the country. In this section, we outline the basic results of the theory.[1]

*Senior Honors Thesis, Stanford University Department of Mathematics. Thanks to Soren Galatius for advising this thesis, to Marc Pauly for introducing me to social choice theory, and to Chris Douglas for introducing me to algebraic topology. Thanks also to the Stanford Mathematics Department and Stanford Undergraduate Research Programs for funding to support this research during summer 2007.

[1] See [1] and [5] for more detailed introductions to social choice theory.

1.1 Preliminaries

First, we introduce the basic language of social choice theory. We call the individuals who have preferences the *voters*. We denote by V the set of voters, which for now we assume to be finite, and which often we will assume for notational simplicity to be $\{1, 2, \ldots, k\}$. The voters have preferences over different *alternatives*. We denote by X the set of alternatives, which again we will assume for now is finite.

We represent preferences by a binary relation R over X. If $(x, y) \in R$, we normally write xRy. We require this relation $R \subset X \times X$ to satisfy three properties:

1. *reflexivity*: xRx for all $x \in X$
2. *transitivity*: $xRy, yRz \Rightarrow xRz$
3. *completeness*: for all $x, y \in X$, either xRy or yRx (or both).

Any such relation is called a *weak preference relation*. The interpretation of xRy is that x is preferred at least as much as y. These requirements correspond to what makes sense for preferences to be "rational."[2]

From this weak preference relation, we can define two more relations. We define a *strict preference relation* $P \subset X \times X$ by xPy iff xRy but not yRx. If xPy, then the interpretation is that x is more preferred than y. We define an *equal preference relation* $E \subset X \times X$ by xEy iff xRy and yRx.[3] If xEy, then the interpretation is that x is preferred equally to y. We observe that E is symmetric. Note that R can be determined from E and P, and vice versa. We will interchange between these notations, as appropriate.

Let $\mathcal{R}$ denote the set of all (weak) preference relations over a given set X.

Definition 1. *A* social welfare function *is a function* $f : \mathcal{R}^k \to \mathcal{R}$.

We define a *profile* of preferences to be a k-tuple $(R_1, \ldots, R_k)$ of preferences, where R_i is voter i's preferences. A social welfare function thus takes a profile of preferences, one for each voter, and aggregates them into one group preference. Social choice theory thus is the study of social welfare functions.

[2] See [1], §1.1, for more on why these are natural requirements.

[3] This is often called *indifference* in the literature, and denoted by I, but "equal preference" I think more accurately describes what it reflects. See [8] for a (rough and preliminary) discussion of this issue.

Notation: When the social welfare function f is clear, we will usually write R for the group preference $f((R_1, \ldots, R_k))$ (and likewise for E and P). We sometimes also write R_f for $f((R_1, \ldots, R_k))$ when the input preference profile is clear, but we want to emphasize the function f.

Observe that a social welfare function takes a profile of preferences and outputs not merely one "winner" but a complete ranked list. A function that instead outputs merely one winner (or a set of winners) is called a *social choice function.* Of course, any social welfare function defines a social choice function: if $f : \mathcal{R}^k \to \mathcal{R}$ is a social welfare function, then we define $F : \mathcal{R}^k \to \mathcal{P}(X)$ to be $F((R_1, \ldots, R_k)) = \{x \in X \mid \forall y \in X, xR_f y\}$. This chooses the top-ranked alternative. Often, all that matters is the social choice function defined; it's not relevant who comes in second or third. But it's standard to develop the more general framework of welfare functions.[4]

1.2 Examples of Social Welfare Functions

We now present some standard examples.

Example 2. *The simplest social welfare function is the* constant *social welfare function that assigns some fixed preference K to all preference profiles $(R_1, \ldots, R_k)$. For example, a rigged presidential election between three candidates a, b, and c where the outcome is going to be $aPbPc$ no matter what the votes are is a constant social welfare function.*

Example 3. *A* dictatorial *social welfare function defines the group preference R to always be the same as some (fixed) voter i's preference R_i.*

Example 4. *Restricting the domain to only strict preferences, we can define the* Borda count *social welfare function. For each alternative x, let $p_i(x) = 1 + |\{y \in X \mid xP_i y\}|$; this assigns $|X|$ points to the highest-ranked alternative, $|X| - 1$ points for the second highest-ranked alternative, all the way down to 1 point for the lowest-ranked alternative. Let $p(x) = \sum_{i \in V} p_i(x)$. Then we define xRy iff $p(x) \geqslant p(y)$.*[5] *This is the mechanism used in sports MVP (most valuable player) voting and in sports ranking polls.*

[4]See [1] for much more on the technical details of the relationship between choice and welfare functions. In particular, one justification for transitivity is that it ensures the existence of a "choice set" of most desired alternatives.

[5]With appropriate modifications, this can be defined on weak preferences as well.

1.3 Properties of Social Welfare Functions

We now present some basic properties a social welfare function might have.

Definition 5. *A social welfare function satisfies the* (weak) Pareto axiom *if, whenever xR_iy for all $i \in V$, then xRy.*

Definition 6. *A social welfare function is* non-dictatorial *if there is no voter $i \in V$ such that xRy if and only if xR_iy, for all $x, y \in X$.*

Definition 7. *A social welfare function satisfies the axiom of* Independence of Irrelevant Alternatives *(IIA) if, given $x, y \in X$, whenever two preference profiles $(R_1, \ldots, R_k)$ and $(R'_1, \ldots, R'_k)$ are such that xR_iy if and only if xR'_iy for all $i \in V$, then xRy if and only if $xR'y$.*

All of these properties are quite reasonable expectations to have for a social welfare function. If everyone agrees that x is better than y, it would be foolish for the outcome to rank x lower than y. If a social welfare function were dictatorial, then it wouldn't in any sense be aggregating the group's preferences. IIA is the most difficult to justify: the intuition is that it shouldn't be relevant to the group's decision between x and y what the group thinks about other alternatives.[6]

REMARK: Social choice theory relies on *ordinal*, not cardinal preferences. There is no sense of "how much" one alternative is preferred to another.

1.4 Arrow's Theorem

The seminal result in social choice theory is Arrow's Impossibility Theorem, which shows that we cannot in general hope to have all of these desireable properties listed above.

Theorem 8 (Arrow's Impossibility Theorem)**.** *When $|X| \geqslant 3$, there is no non-dictatorial social welfare function $f : \mathcal{R}^k \to \mathcal{R}$ that satisfies both the Pareto principle and IIA.*

(Note that Pareto and non-dictatorship imply that the number of voters $k > 1$.)

A constant social welfare function violates the Pareto axiom. A dictatorial social welfare violates, of course, the non-dictatorial axiom. Borda count does not satisfy IIA.

[6]See [1], §2.1, for more discussion on IIA.

One might wonder why majority voting doesn't work. We define majority voting in a pairwise manner: xPy if $|\{i \in N \mid xP_iy\}| \geqslant |\{i \in N \mid yP_ix\}|$. This does indeed satisfy Pareto, IIA, and non-dictatorship, but it sometimes produces a non-transitive relation. Consider the following set of preferences:

P_1	P_2	P_2
a	b	c
b	c	a
c	a	b

where the preferences are strict and x is above y in a column when x is strictly preferred to y. (We often use such diagrammatic representations, which provide more intuitive depictions of preferences.) Pairwise majority voting gives aPb, bPa, cPa, contradicting transitivity. This is known as Condorcet's paradox and dates to the 18$^{\text{th}}$ century.

(When there are only 2 alternatives, however, majority voting works just fine. Indeed, it can be shown that it is the *unique* procedure that satisfies a few basic axioms. This is May's Theorem; see [5], §3.1.)

We present a proof of Arrow's Theorem from [1]. First, we need a few definitions. Fix some social welfare function f.

Definition 9. *A subset $L \subset V$ of voters is* semi-decisive *for x over y (with respect to f), denoted $x\tilde{D}_Ly$, if $[xP_iy, \forall i \in L; yP_jx, \forall j \notin L]$ implies that xPy, for all profiles $\rho = (R_1, \ldots, R_k) \in \mathcal{R}^k$.*

Definition 10. *A subset $L \subset V$ of voters is* decisive *for x over y (with respect to f), denoted xDy, if $[xP_iy, \forall i \in L]$ implies that xPy, for all profiles $\rho = (R_1, \ldots, R_k) \in \mathcal{R}^k$.*

The proof requires one lemma.

Lemma 11. *Let f be a transitive social welfare function that satisfies the IIA and weak Pareto axioms, and let $L \subset V$. Suppose there is some $x, y \in X$ such that L is semi-decisive for x over y. Then for all pairs $(u, v) \in X \times X$, L is decisive for u over v.*

Proof. Consider some $z \in X \setminus \{x, y\}$. We will first show that xD_Lz. Let ρ be an arbitrary preference profile such that xP_iz for all $i \in L$. Consider now a profile ρ' where, for $i \in L, xP_i'yP_i'z$ and for $i \notin L, yP_i'x, yP_i'z$. Diagramatically, this looks like:

$P_i, i \in L$	$P_i', i \in L$	$P_i', i \notin L$
x	x	y
	y	
z	z	$x?z$

where $x?z$ means that the relationship betwen x and z is not specified. Because L is semi-decisive for x over y, we have $xP'y$. By Pareto, since $yP_i'z$ for all $i \in V$, we have $yP'z$. Transitivity then gives us $xP'z$. The profiles ρ and ρ' are identical when restricted to the alternatives x and z over L, and we can make the preferences of ρ and ρ' match for $i \notin L$ (since we hadn't specified it yet for ρ'), so we get xPz by IIA. Therefore, L is decisive for x over z. Thus we have

$$\forall z \notin \{x, y\}, x\tilde{D}_L y \Rightarrow xD_L z \tag{1}$$

Since xD_Lz, we have *a fortiori* $x\tilde{D}_Lz$. The same argument, interchanging y and z, gives xD_Ly.

Now, let ρ^* be an arbitrary preference profile such that yP_i^*z for all $i \in L$. Consider a profile ρ^+ where, for all $i \in L, yP_i^+xP_i^+z$ and for all $i \notin L, zP_i^+x, yP_i^+x$. Diagramatically, we have:

$P_i^*, i \in L$	$P_i^+, i \in L$	$P_i^+, i \notin L$
y	y	$z?y$
	x	
z	z	x

Because xD_Lz, we have xP^+z. By Pareto, we have yP^+x. Transitivity then gives yP^+z. Since ρ^* and ρ^+ are identical when restricted to y and z over L, and since we can make the preferences of ρ^* and ρ^+ match for $i \notin L$ (because we hadn't specified it yet for ρ^+), we get yP^*z. Therefore, we have yD_Lz. Thus, we have shown

$$\forall z \notin \{x, y\}\, x\tilde{D}_L y \Rightarrow yD_L z. \tag{2}$$

Since yD_Lz, we have *a fortiori* $y\tilde{D}_Lz$. By the same argument that gave us equation (1), we get yD_Lx. Now consider any $\{u, v\} \in X \setminus \{x, y\}$. By (1) we get xD_lv, and by (2), replacing y by v, we get vD_Lw. This concludes the proof, since we have already shown the lemma in the cases where one of $u, v \in \{x, y\}$. □

With this lemma, we can finish the proof of Arrow's theorem.

Proof. It suffices to show that there is an individual i and alternatives $x, y \in X$ such that $x\tilde{D}_{\{i\}}y$. Note that the set of all voters V is a decisive and thus semi-decisive set for any pair, by Pareto. Therefore, there is a semi-decisive set. Let L be a semi-decisive set of minimal size; without loss of generality, we can assume that L is semi-decisive for x over y and that $i \in L$. Let $\lambda = |L|$.

Assume for the sake of contradiction that $\lambda > 1$. Consider the following preference profile ρ:

P_i	$P_j, j \in L\setminus\{i\}$	$P_k, k \notin L$
x	z	y
y	x	z
z	y	x

Because $x\tilde{D}_L y$, we have xPy. Suppose zPy: then we would have $L\setminus\{i\}$ is semi-decisive for z over y, by IIA, contradicting L's minimality. Therefore, we have $\neg zPy$, that is, yRz. By transitivity, we have xPz. Therefore, by IIA, $xD_{\{i\}}z$ and so $x\tilde{D}_{\{i\}}z$, contradicting $\lambda > 1$. □

We will present an interesting topological proof of this theorem in §4.

1.5 Escaping Arrow's Theorem

Arrow's Theorem is a pretty dismal result. It shows that there is no reasonable way of aggregating preferences in the general case (though majority voting works well if there are just two candidates). Much of social choice theory has consisted of attempts to salvage as many of these as properties as possible.

One escape from this impossibility is to restrict the domain of the social welfare function. The most common domain restriction is known as single-peakedness: if all of the voter's preferences meet a certain pattern and there is an odd number of voters, then simple majority voting does produce a transitive relation.

2 An Outline of Topology and Algebraic Topology

Traditional social choice theory deals with *discrete*, usually finite sets of alternatives, and its formal results rely on straightforward, combinatorial arguments. In the following, we will be considering social choice theory with *continuous* sets of alternatives—subsets of Euclidean space, for example—and continuous social welfare functions. To prove formal results in this continuous framework, we will need more powerful mathematical machinery: the tools of algebraic topology. In this section, we review the basics of topology and algebraic topology.

2.1 Point Set Topology

There are several different ways of thinking of topology. While one simplistic perspective is to see topology as being "rubber-sheet" geometry, a more useful approach for us is to see topology as being fundamentally the study of continuity. In metric spaces, continuity can be defined using open sets, which in turn are defined by ε-balls using the given metric. In topology, the notion of distance is done away with, and the open sets are taken as the fundamental building blocks. Thus, though there is no notion of distance in topological spaces, continuity still makes sense—and, indeed, this is often the appropriate level of abstraction to study continuity at.

For the details of the basics of topology, see, for example, [11]. What is relevant for our purposes is that a topological space is a pair $(X,\mathcal{U})$, where $\mathcal{U} \subset \mathcal{P}(X)$ is the collection of all *open sets* in X. (For this to be a topological space, $\mathcal{U}$ needs to be nicely behaved under various set-theoretic operations.) We often speak of simply X as being a topological space, without explicitly including its collection of open sets $\mathcal{U}$. A function $f : (X,\mathcal{U}) \to (Y,\mathcal{V})$ is *continuous* if, for all open sets V in Y (i.e., all $V \in \mathcal{V}$), the inverse image $f^{-1}(V)$ is open in X (i.e., it is in $\mathcal{U}$). Two topological spaces X and Y are considered "the same" if there is a bijective function $f : X \to Y$ that is continuous and whose inverse is also continuous; such a map is called a *homeomorphism*, two such spaces are called *homeomorphic*, and we write $X \cong Y$.

2.2 Types of Topological Spaces: Manifolds and CW Complexes

One of the most important types of topological spaces are manifolds. Manifolds are spaces that essentially "look like" Euclidean space up close. Formally, an n-dimensional manifold M is a topological space where, for each point $x \in M$ there is an open set U containing x that is homeomorphic to $\mathbb{R}^n$. For example, the standard sphere S^2 is a 2-dimensional manifold, since locally it looks like a plane. Also of relevance to us will be higher dimensional spheres, defined by

$$S^k = \left\{ x \in \mathbb{R}^{k+1} \mid \|x\| = 1 \right\}.$$

Another mathematical structure that will be useful for us is that of a *CW complex*, which allows us to view certain topological spaces as being "built" from combining together cells of different dimensions. Let us define an n-cell as being either a point for $n = 0$ or else homeomorphic to S^n. Now, we built an n-skeleton X^n inductively from an $n-1$-skeleton X^{n-1} and a collection of n-cells e^n_α by gluing the n-cells onto X^{n-1} via maps $\varphi_\alpha : S^{n-1} \to X^{n-1}$ that glue the boundaries of the n-cells onto the complex X^{n-1}. Then, we either continue this process infinitely, or stop if there are only a finite number of cells. Most useful topological spaces can be given a CW complex structure, and several of our results will be stated for CW complexes. See, e.g., [6], Chapter 0 and Appendix, for more details.

2.3 Algebraic Topology

One of the fundamental problems of topology is to classify different topological spaces. While it's sometimes hard to show directly that there is a homeomorphism between two homeomorphic spaces, since this would require the effective construction of a homeomorphism, at the very least, we would want to be able to show that two non-homeomorphic spaces are, indeed, not homeomorphic.

One technique for doing so is through the use of so-called *algebraic invariants*. For some class of topological spaces $\mathcal{X}$ and some class of algebraic objects $\mathcal{G}$, we define a function $f : \mathcal{X} \to \mathcal{G}$. For this function to be well-defined, it must assign the same algebraic object to homeomorphic topological spaces. (Hence the name "algebraic invariant"; it is an algebraic object that is invariant under homeomorphism.) With this notion, we have a sufficient (though usually not necessary) condition for distinguishing between two spaces: if two spaces X and Y have different assigned algebraic invariants, they cannot be homeomorphic. These invariants will furthermore be *functors*, which means that they not only define a map from $\mathcal{X}$ to $\mathcal{G}$, but also they define in a special way maps between basic functions between spaces in $\mathcal{X}$ (usually continuous functions) and basic functions between objects in $\mathcal{G}$ (usually group homomorphisms); these maps will be called *induced maps*.

Algebraic topology is the study of such algebraic invariants of topological spaces. The three most important algebraic invariants are the homotopy, homology, and cohomology groups, all of which we will use in our results on continuous social choice theory. To explain any of these adequately is far beyond the scope of this exposition; see [6] for such a treatment. We will present a basic outline of fundamental groups; the other two fundamental invariants of algebraic topology—homology and cohomology—are much less intuitive, though often easier to calculate. Finally, we will briefly list some basic results that will be used in the following.

2.3.1 Homotopy Groups

The simplest of the homotopy groups, called the fundamental group, is at least somewhat intuitive. Let X be a topological space, and let x be a fixed point in X. Consider all of the possible loops through X starting at x, that is, the set of all functions $\varphi : [0,1] \to X$ such that $\varphi(0) = \varphi(1) = x$. Now, we define an equivalence relation on such loops: we say that φ and γ are *homotopic* if there is a family of continuous functions $f_t : [0,1] \to X$, for $t \in [0,1]$ such that $f_0 = \varphi, f_1 = \psi$ that are also continuous with respect to t. Intuitively, this means that we can continuously deform the path φ to the path ψ. If φ and ψ are homotopic, we write $\varphi \simeq \psi$. We now define $\pi_1(X, x)$ to be the set of all equivalence classes of such loops based at x. We can add a group structure to this set, by defining multiplication in an appropriate way as the concatenation of different loops. This group is known as the *fundamental group* or *first homotopy group* of X. (It turns out that the basepoint is usually not important.)

A solid sphere, for example, has a trivial fundamental group, since any loop can be continuously deformed into a constant loop at the center point. The circle S^1 has a fundamental group isomorphic to $\mathbb{Z}$; the loop going once around (in the predetermined "positive" direction) corresponds to 1, the loop going twice around corresponds to 2, the loop going once around in the opposite direction corresponds to -1, etc. Since the solid sphere and the circle have different fundamental groups, we know that they must be non-homeomorphic.

There are also higher-dimensional generalizations of the fundamental group, called the *higher homotopy groups*. We could have defined our loops not as functions $\varphi : [0,1] \to X$ with the condition that $\varphi(0) = \varphi(1)$, but rather as functions $\varphi : S^1 \to X$. The higher-dimensional homotopy groups consist of equivalence classes now of functions $\varphi : S^k \to X$ for $k > 1$. These are, however, much more difficult to calculate than the fundamental group. (One useful property of the homotopy groups is their behavior under cartesian products: for connected spaces X_α, $\pi_n(\prod_\alpha X_\alpha) \cong \prod_\alpha \pi_n(X_\alpha)$.)

2.3.2 A Few Results in Algebraic Topology

We present the statement of two important theorems that we will be using.

Theorem 12 (Whitehead's Theorem)**.** *Let X and Y be connected CW complexes, and assume that a map $f : X \to Y$ induces an isomorphism $f_* : \pi_n(X) \to \pi_n(Y)$ for all n. Then f is a homotopy equivalence.*

In particular, if all the homotopy groups of a CW complex are trivial, then it must be homotopy equivalent to a point, i.e., contractible. For details, see [6], §4.1.

The Hurewicz Isomorphism Theorem says that there's a nice relationship between the first non-zero homotopy and homology groups.

Theorem 13 (Hurewicz Isomorphism Theorem)**.** *Let X be a space and $n \geqslant 2$. If $\pi_i(X) = 0$ for $i \leqslant n-1$, then $H_i(X) = 0$ for $i < n$ and $\pi_n(X) \cong H_n(X)$.*

See [6], §4.2 for details.

3 Continuous Social Choice

Traditional social choice theory deals with finite sets of alternatives—a set of candidates in an election, for example. Some scholars did study the cases when these alternatives were infinite, but still the alternatives were treated as distinct, discrete objects.

There are many natural cases, however, where a set of alternatives should be continuous. Much of economic theory is based on the idea of commodity spaces, which are represented as the positive quadrants of $\mathbb{R}^n$ (each dimension represents the amount of a given commodity; for example, one might have $\mathbb{R}^3$ as the commodity space for pounds of rice, gold, and oil). Or voters might choose between different geographic locations: Gaertner [5] gives the example of a group of vacationers choosing which point around a circular lake to camp out at. In the early 1980s, the mathematical economist Graciela Chichilnisky developed a framework to develop social choice theory for such continuous sets of alternatives, and used tools from algebraic topology to prove impossibility results analagous to Arrow's theorem.

3.1 The Basic Framework

Let $V = \{1, \ldots, k\}$ be a finite set of voters. We define our set of alternatives X to be some subset of $\mathbb{R}^n$; we often call this the *choice space*.

How do we define preferences on X? In general economic theory, the standard approach is to consider some utility function $u : X \to [0, \infty)$ that assigns to each point its "utility." Then we define xRy iff $u(x) \geqslant u(y)$. In social choice theory, preferences are ordinal: there is no notion of "how much" one alternative is preferred to another. Thus, two utility functions $u, v : X \to \mathbb{R}$ are equivalent in the ordinal framework if there is a monotonic function $f : \mathbb{R} \to \mathbb{R}$ such that $u = f \circ v$.

It is meaningless for our choice space to be continuous unless in some sense the preferences are continuous: for sufficiently close x and y, the preferences should be similar. Thus, it seems reasonable to assume that, at the very least, such underlying utility functions are continuous.

The approach taken by Chichilnisky and others is to consider the gradient of these utility functions. Let p be a vector field on X; the idea is that $p(x)$ represents the direction of most increasing preference. Furthermore, since the preferences are ordinal, we normalize all of these gradient vectors to length 1. Thus, we require that the vectors all be non-0, which means that our utility functions have no critical points.

Now we can move to the classic results of Chichilnisky.

3.2 Chichilnisky's Impossibility Theorem

Let the choice space X be homeomorphic to the unit ball in $\mathbb{R}^n$. We define preferences on X as a continuous vector field, normalized to length 1. We denote the set of all such vector fields by P and an individual vector field by

p, or p_i for a given voter.[7] In this context, we are curious about the existence of functions $f : P^k \to P$. We require that these functions be continuous, a natural assumption considering that our spaces are continuous.[8] For all theorems, we assume that $k \geqslant 2$.

Of course, to define "continuity," we need to define a topology on the preference space. We use the C^1 topology induced by the sup norm: $\|p - q\| := \sup_{x \in X} \|p(x) - q(x)\|_X$. Various papers discuss the technicalities of this; see [9] for details of which topologies work and for references to the various discussions on the relative merits of different topologies.

Now, we list some relevant properties these functions might have:

Definition 14. *A social welfare function f is* anonymous *if*

$$f(p_1, p_2, \dots, p_k) = f(p_{\sigma(1)}, p_{\sigma(2)}, \dots, p_{\sigma(k)})$$

for any permutation $\sigma : V \to V$.

Definition 15. *A social welfare function f respects* unanimity *if $f(p, \dots, p) = p$.*

Anonymity is a stronger form of non-dictatorship: not only does no individual have complete control over the preferences, but every individual is treated equally. Unanimity is a weaker form of the Pareto criterion.

We can now state the original impossibility theorem of Chichilnisky [4]. Recall that $X \subset \mathbb{R}^n$.

Theorem 16. *There is no continuous social welfare function $\varphi : P^k \to P$ satisfying anonymity and unanimity.*

Proof. Suppose for the sake of contradiction that there were such a function φ. Fix some $x \in X$. At this given x, we note that $\{p(x) \mid p \in P\} = S^{n-1}$.

[7] This seems to be the standard notation for continuous social choice, as opposed to the $\mathcal{P}, P$ notation for discrete social choice.

[8] This is a non-trivial requirement, and the whole continuous social choice paradigm—which is completely dependent on the requirement—is susceptible to the criticism that it is too strong. After all, in discrete social choice, many social welfare functions don't seem "continuous," in the sense that they are not always stable with respect to small perturbations: in majority voting, for example, it takes a switch by only one voter to drastically change the outcome. But, then again, this should only really be troubling and/or surprising if there is a "large" number of voters, in which case the probability of such a perturbation having an effect is very low.

The idea of the proof will be to turn the question about existence of a continuous function $\varphi : P^k \to P$ satisfying the anonymity and unanimity into a question about the existence of a continuous function $\psi : (S^{n-1})^k \to S^{n-1}$. We will then show that ψ cannot exist.

Define $\Gamma_x : P \to S^{n-1}$ to be $\Gamma_x(p) = p(x)$. This continuously maps a preference over the whole space into just its direction at x. Define $\lambda : S^{n-1} \to P$ to take a vector $v \in S^{n-1}$ and send it to the constant vector field $p(y) = v$ for all $y \in X$. This gives a map $\lambda^k : (S^{n-1})^k \to P^k$ that maps a k-tuple of unit vectors $(v_1, \dots, v_n)$ to the k-tuple $(p_1, \dots, p_k)$ of constant vector fields. We note that this is also continuous.

We can now define $\psi : (S^{n-1})^k \to S^{n-1}$ by the following commutative diagram:

$$\begin{array}{ccc} P^k & \xrightarrow{\varphi} & P \\ {\scriptstyle \lambda^k}\big\uparrow & & \big\downarrow{\scriptstyle \Gamma_x} \\ (S^{n-1})^k & \xrightarrow{\psi} & S^{n-1} \end{array}$$

Our new function ψ is continuous since λ^k, φ, and Γ_x are. We observe also that it respects unanimity and anonymity, since φ does.

It suffices now to show that ψ is impossible. We have thus reduced the problem to that of constant (or what Chichilnisky calls "linear") preferences.

We demonstrate the proof for the case $k = 2$, but it can easily be generalized for arbitrary k. Consider the following diagram

$$\begin{array}{ccc} S^{n-1} \times S^{n-1} & \xrightarrow{\psi} & S^{n-1} \\ {\scriptstyle \Delta}\big\uparrow & \nearrow_{Id_{S^{n-1}}} & \\ S^{n-1} & & \end{array}$$

where Δ sends $p \mapsto (p, p)$. We know that this diagram commutes by the condition of unanimity. This induces the following commutative diagram on cohomology with coefficients in $\mathbb{Z}_2$:

$$\begin{array}{ccc} H^{n-1}(S^{n-1} \times S^{n-1}, \mathbb{Z}_2) & \xleftarrow{\psi^*} & H^{n-1}(S^{n-1}, \mathbb{Z}_2) \\ {\scriptstyle \Delta^*}\big\downarrow & \swarrow_{Id^*_{S^{n-1}}} & \\ H^{n-1}(S^{n-1}, \mathbb{Z}_2) & & \end{array}$$

Let A^* and B^* be generators of $H^{n-1}(S^{n-1} \times S^{n-1}, \mathbb{Z}_2) \approx \mathbb{Z}_2 \times \mathbb{Z}_2$, with support by A and B, which, in the case of $n = 1$, we can see as sides of the torus:

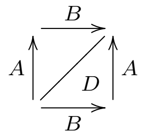

Let G^* be a generator of $H^{n-1}(S^{n-1}, \mathbb{Z}_2)$. We note that we cannot have

$$\psi^*(G^*) = A^*$$

or

$$\psi^*(G^*) = B^*$$

by anonymity. Thus we must have either

$$\psi^*(G^*) = A^* + B^*$$

or

$$\psi^*(G^*) = 0 \text{ (mod 2)}.$$

We cannot have $\psi^*(G^*) = 0$, since if we did,

$$\Delta^*(G^*) \circ \psi^*(G^*) = \Delta^*(0) = 0 \neq G^* = Id^*_{S^n}(G^*),$$

violating the commutativity of the diagram. Thus we must have

$$\psi^*(G^*) = A^* + B^*.$$

By symmetry, $\Delta^*(A^*) = \Delta^*(B^*) = k$ for some $k \in \mathbb{Z}_2$. Thus $\Delta^*(A^* + B^*) = k + k = 2k = 0$ (mod 2). But this contradicts the commutativity of the diagram:

$$\begin{aligned} \Delta^*(\psi^*(G^*)) &= \Delta^*(A^* + B^*) \\ &= \Delta^*(0) = 0 \\ &\neq 1 = Id^*_{S^{n-1}}(G^*) \end{aligned}$$

Thus, we have a contradiction. A similar proof applies for arbitrary $k \geqslant 2$ by using coefficients in $\mathbb{Z}_k$. □

(Note that, for what we have proved, all we need is "homotopy anonymity.") There are also more intuitive, geometric proofs in special cases, e.g., when $k = 2$. Chichilnisky also discusses a connection to fixed point theorems. Mehta [10] presents this result in terms of degree theory. We present a second proof below, as part of Theorem 17.

3.3 Necessary and Sufficient Conditions for Existence

A later theorem provides necessary and sufficient conditions for the existence of a satisfactory social welfare function. This strengthens Chichilnisky's original result (Theorem 16).[9]

Theorem 17. [3] *Let P be a parafinite CW complex*[10] *reflecting preferences; furthermore, assume that it is a simplicial complex. A necessary and sufficient for the existence of a continuous social welfare function $\varphi : P^k \to P$ satisfying unanimity and anonymity, where $k \geqslant 2$, is that the spaces P be contractible.*

Proof. First, suppose the spaces P are contractible. We want to show the existence of a map $P^k \to P$ that is anonymous and satisfies unanimity. The anonymity criterion corresponds to requiring the existence of a map from $P^k/\Sigma_k \to P$, where P^k/Σ_k is the quotient group of the k-fold product of P modulo the action of the symmetric group permuting the factors. Consider the following diagram

$$P \xhookrightarrow{i} P^k/\Sigma_k \xrightarrow{\varphi} P, \qquad \varphi \circ i = id$$

We want to show the existence of this function φ such that the diagram commutes; this will ensure unanimity. Since P is a simplicial complex, P^k/Σ_k is a CW complex (see [6], p. 482). Therefore, we can indeed extend by induction on cells the identity map $id : P \to P$ to the map $\varphi : P^k/\Sigma_k$, since P is contractible and therefore has trivial homotopy groups ([6], Lemma 4.1.7). Thus, such a map φ exists.

Now suppose that there is a continuous social welfare function $\varphi : P^k \to P$ satisfying unanimity and anonymity. We will show that P must be contractible. Since P is a CW complex, by Whitehead's Theorem (Theorem 12), it suffices to show that all the homotopy groups are trivial.

[9] We prove a slightly weaker result than Chichilnisky and Heal present. They present the necessary and sufficient conditions for CW complexes, but their proof of sufficiency seems somewhat flawed. (The gist of their proof is that, if the preference space can be made convex, then convex averaging works; the problem is that it's not clear from their proof how to get the convex space from the preference space.) Our proof requires that the spaces be simplicial complexes for sufficiency; [12] and [7] provide proofs for a much wider range of spaces.

[10] A CW complex is parafinite if it has a finite number of cells in any dimension; this term is used in the continuous social choice literature, though not in the topology literature.

Our map φ induces a homomorphism

$$\varphi_* : \pi_i(P^k) \to \pi_i(P)$$

on homotopy groups. We note that

$$\pi_i(P^k) \cong \prod_k \pi_i(P).$$

By anonymity, since

$$\varphi(x, x, \ldots, x) = x,$$

we have, for any $\gamma \in \varphi_i(P)$,

$$\varphi_*(\gamma, \gamma, \ldots, \gamma) = \gamma. \tag{3}$$

Let e denote the identity element of $\pi_i(P)$. By φ_* being a homomorphism and by anonymity, we have

$$\begin{aligned} \varphi_*(\gamma, \gamma, \ldots, \gamma) &= \underbrace{\varphi_*(\gamma, e, e \ldots, e) + \varphi_*(e, \gamma, e, \ldots, e) + \cdots + \varphi_*(e, e, \ldots, \gamma)}_{k} \\ &= k \cdot \varphi_*(\gamma, e, \ldots, e) \end{aligned}$$

Thus, we have, for any $\gamma \in \pi_i(P)$,

$$\gamma = k \cdot \varphi_*(\gamma, e, e, \ldots, e). \tag{4}$$

We want to show that $\varphi_i(P) = 0$ for all $i \geqslant 0$. First, let us consider $i \geqslant 2$. Let i be the first non-zero homotopy group. By the Hurewicz Isomorphism Theorem (Theorem 13), $\pi_i(P) \cong H_i(P)$. Since P is a parafinite CW complex, $H_i(P)$ is finitely generated. It therefore consists of a free part consisting of copies of $\mathbb{Z}$ and a torsion part consisting of $\mathbb{Z}_k$ for various $k \geqslant 2$. We will show that both of these parts must be trivial.

Let γ be a generator of the free part of $\pi_i(P)$. By (4), $\gamma = k\eta$, for some $\eta = \varphi_*(\gamma, e, \ldots, e) \in \pi_i(P)$. Since $k \geqslant 2$, if γ is a generator, then $\gamma = e$, and so the free part is trivial. Now let γ be a generator of the torsion. Once again, we have $\gamma = k\eta$ for some $\eta \in \pi_i(P)$, which forces $\gamma = e$ and so the torsion is trivial.

All that remains is $\pi_1(P)$. If this is abelian, then $\pi_1(P) \cong H_1(P)$ and the above argument applies. We will show that $\pi_1(P)$ is abelian.[11]

[11] Above we used additive notation. Here, since $\pi_1(P)$ might not be abelian, we use multiplicative notation.

Let $a, b \in \pi_1(P)$. Then, using (3),

$$\begin{aligned}
a \cdot b &= \varphi_*(a, a, \ldots, a)\varphi_*(b, b, \ldots, b) \\
\text{(by (4))} &= \varphi_*(a, e, e, \ldots, e)^k \varphi_*(e, b, e, \ldots, e)^k \\
&= \varphi_*(a^k, e, e, \ldots, e)\varphi_*(e, b^k, e, \ldots, e) \\
&= \varphi_*(a^k, b^k, e, \ldots, e) \\
&= \varphi_*(b^k, e, e, \ldots, e)\varphi_*(e, a^k, e, \ldots, e) \\
&= \varphi_*(b, b, \ldots, b)\varphi_*(a, a, \ldots, a) \\
&= b \cdot a
\end{aligned}$$

This completes the proof. □

Some of these results were earlier discovered in a purely topological context by Auman and Eckman; see [7] for references.

3.4 Other Results

Since Chichilnisky's initial work in the 1980s, there have been a couple of dozen more recent papers on continuous social choice theory. See [9] and [10] for recent surveys. It should be noted that it's not clear what, if any, impact this continuous framework has had, and that there are some critics.

4 Baryshnikov's Topological Proof of Arrow's Theorem

In the 1990s, Baryshnikov [2] gave a topological proof of Arrow's Theorem that suggests a unification with Chichilnisky's results.

4.1 Statement and Simplification

We recall the statement of Arrow's Theorem:

Theorem (Arrow's Impossibility Theorem)**.** *When $|X| \geqslant 3$, there is no non-dictatorial social welfare function $f : \mathcal{R}^k \to \mathcal{R}$ that satisfies both the Pareto principle and IIA.*

First, Baryshnikov shows that it suffices to show non-existence of such a function on strict preferences. (Recall from §1.1 that a strict preference is a total order, whereas a weak preference is a weak order.)

Lemma 18 (Baryshnikov Lemma 1)**.** *The image of the restriction of a social welfare function satisfying IIA and Pareto to strict orders lies in the set of strict orders.*

If there were a non-dictatorial function on weak preferences that satisfied IIA and Pareto, then, by the lemma, its restriction to strict preferences would be a non-dictatorial function on strict preferences satisfying IIA and Pareto.

Proof. Let $\mathcal{P}$ be the set of strict preferences on X and $\mathcal{R}$ the set of weak preferences on X. Let $\varphi : \mathcal{R}^k \to \mathcal{R}$ satisfy IIA and Pareto. We want to show that $\varphi|_{\mathcal{P}^k} \subset \mathcal{P}$.

Suppose for the sake of contradiction that there is some *strict* preference profile $P^k = (P_1, \ldots, P_k)$ such that there are $x \neq y \in X$ with $xE_{\varphi(P^k)}y$. By IIA, any profile P'^k that agrees with P^k on x and y will also produce $xE_{\varphi(P^k)}y$. Since $|X| \geqslant 3$, there is some third alternative z. Let $P^{k(0)} = (P_1^{(0)}, \ldots, P_k^{(0)})$ be a profile of strict preferences that agrees with P^k on x and y, and where z is ranked immediately below y for all voters i. Let $P^{k(1)} = (P_1^{(1)}, \ldots, P_k^{(1)})$ differ from $P^{k(0)}$ only in that z is now ranked directly above y instead of directly below it. By Pareto, $yP^{(0)}z$ and $zP^{(1)}y$. $P^{k(0)}$ and $P^{k(1)}$ offer the same rankings of x relative to y. Therefore by IIA we have $yE^{(0)}x$ and $yE^{(1)}x$, which implies that $xP^{(0)}z$ and $zP^{(1)}x$. But $P^{k(0)}$ and $P^{k(1)}$ agree in their rankings of x and z, so by IIA, $xP^{(0)}z \iff xP^{(1)}z$, giving a contradiction. □

The idea of our proof will be to provide a topological structure to the discrete, combinatorial rankings in Arrow's framework. In §4.2, we introduce *nerves*, which we will use as our tool to create this topology. In §4.3 and §4.4, we will construct the appropriate topological spaces for preferences and profiles of preferences, respectively, and calculate their homologies and cohomologies. Finally, in §4.5, we will provide a general axiomatic framework for these topological spaces, and prove Arrow's theorem using properties of their homology and cohomology.

4.2 Nerves and the Nerve Theorem

Let X be a topological space, and let $\mathcal{U} = \{U_i\}_{i\in I}$ be a cover of X. The *nerve* $\mathcal{N}(\mathcal{U})$ of this cover is the simplicial complex with vertices the $U \in \mathcal{U}$ and with simplices $\{U_1, \ldots, U_k\}$ iff $U_1 \cap \cdots \cap U_k \neq \emptyset$. The concept of a nerve was introduced to provide a combinatorial basis for topological spaces, but we will use it to be provide a non-trivial topological space on a combinatorial (discrete) space.

The key theorem we will need about nerves is the nerve theorem.

Theorem 19 (Nerve Theorem)**.** *Let X be a paracompact topological space and $\mathcal{U} = \{U_i\}_{i\in I}$ an open cover of X. Suppose that the intersection $\cap_{j\in J} U_j = \emptyset$ or is contractible for all $J \subseteq I$. Then $\mathcal{N}(\mathcal{U}) \simeq X$.*

A proof is given in [6], §4.G.

4.3 Topology on Preferences

Let $\mathcal{P}$ be the set of strict preferences on a set X of alternatives of cardinality n. For notational convenience, let $X = \{1, 2, \ldots, n\}$. Let $i < j \in N$. We define

$$\mathcal{U}_{ij}^{+} = \{\, P \in \mathcal{P} \mid iPj \}$$

and

$$\mathcal{U}_{ij}^{-} = \{\, P \in \mathcal{P} \mid jPi \}\,.$$

(Note that $\mathcal{U}_{ij}^{+} = \mathcal{U}_{ji}^{-}$.) Let $\mathcal{U} = \left\{ \mathcal{U}_{ij}^{\sigma} \;\middle|\; \sigma \in \{-,+\}\,, i < j \right\}$. We observe that $\mathcal{U}$ covers $\mathcal{P}$. Let us denote the nerve of this covering of $\mathcal{P}$ by $\mathcal{N}_{\mathcal{P}}$.

Theorem 20 (Baryshnikov Theorem 1)**.** *The simplicial complex $\mathcal{N}_{\mathcal{P}}$ is homotopy equivalent to the sphere S^{n-2}.*

Proof. Let $\Delta = \{\, (x, x, \ldots, x) \in \mathbb{R}^n \mid x \in \mathbb{R}\}$. Let $M = \mathbb{R}^n \backslash \Delta$.

Claim: $M \simeq S^{n-2}$.

Proof of Claim: For simplicity, we assume that Δ instead is the nth axis. Thus $M = \{\, (x_1, \ldots, x_n) \mid \neg(x_1 = \cdots = x_{n-1} = 0)\}$. Consider the orthogonal complement of Δ: $\Delta^{\perp} = \{(x_1, \ldots, x_{n-1}, 0)\}$. We can easily deformation retract M onto this orthogonal complement, giving us $M \simeq \{\, (x_1, \ldots, x_n) \mid x_n \neq 0\}$. This is simply $\mathbb{R}^{n-1}$ minus a point, which is clearly homotopy equivalent to S^{n-2}. □

Our plan will be to show that M with a certain covering has the same nerve as $\mathcal{P}$ does, but also satisfies the conditions of the nerve theorem, and so $\mathcal{N}_{\mathcal{P}} \simeq S^{n-2}$.

Define $U_{ij} = \{(x_1, \dots, x_n) \in \mathbb{R}^n \mid x_i > x_j\}$. Now, we show that any $P \in \mathcal{P}$ can be represented by a vector in M. A simple way to do this is to define $f : \mathcal{P} \to M$ by sending $P \mapsto (x_1, \dots, x_n)$ where if i is most preferred by P, then $x_i = n$, if i is second-most preferred by P, then $x_i = n - 1$, and so on. (Equivalently, x_i is the Borda count value for i (see Example 4).) Now we see of course the link between the U_{ij}'s in M and the $\mathcal{U}_{ij}^{\sigma}$'s in $\mathcal{P}$. We observe that the intersection properties of these are the same, thus the nerve $\mathcal{N}_M$ of M with the covering $U = \{U_{ij}\}_{i \neq j \in X}$ is the same as $\mathcal{N}_{\mathcal{P}}$.

Finally, we observe that these U_{ij}'s in M are convex, and thus any intersection of them will be convex, and thus contractible. Thus, by the nerve theorem, $\mathcal{N}_M \simeq M$. We know $\mathcal{N}_M = \mathcal{N}_{\mathcal{P}}$ and $M \simeq S^{n-2}$, establishing our result. □

Now, we want to figure out what the generator of the $H_{n-2}(\mathcal{N}_{\mathcal{P}})$ is.

Let Δ_{tot} be the simplex on all the vertices $\mathcal{U}_{ij}^{\sigma}$. Note that $\mathcal{N}_{\mathcal{P}}$ is a subcomplex of this. First, we set up a bijective correspondence between oriented graphs on the vertices $X = \{1, \dots, n\}$ and sub-simplices of Δ_{tot}. (We pick some lexigraphic ordering on the $\mathcal{U}_{ij}^{\sigma}$s to give an orientation to these simplicial complexes.) Consider some oriented, simple graph on X. For each directed edge from i to j that this graph has, the corresponding simplex has as a vertex $\mathcal{U}_{ij}^{+}$; likewise, if there's an arrow from j to i, then the simplex has $\mathcal{U}_{ij}^{-}$.

Note that the corresponding simplicial complex is in $\mathcal{N}_{\mathcal{P}}$ iff the graph g has no oriented cycles. First, if g has an oriented cycle, then it will violate transitivity, so there can be no preference relation satisfying it. Now, suppose that g has no oriented cycles. We claim that we can then extend it to a total order. This extension proceeds by induction on the number of vertices. Since there are no oriented cycles, there is some vertex v that has no arrows pointing to it. Add the edge $v \to w$ for all vertices w that don't already have such an edge. Now, inductively repeat this process on the graph omitting the vertex v. This process will terminate with a total order. This order will correspond to some simplex in $\mathcal{N}_{\mathcal{P}}$, and so the original partial order will correspond to a subsimplex.

Consider a simple cycle of length n on the set of alternatives $X = \{1, \dots, n\}$. Let this be an oriented graph; call it g. Denote the corresponding simplex by $S(g)$. Now consider the boundary $\partial S(g)$; it is a cycle.

Proposition 21 (Baryshnikov Proposition 2)**.** *The cycle $\partial S(g)$ is an $n-2$-dimensional cycle of $\mathcal{N}_{\mathcal{P}}$. If g is an oriented cycle, then $\partial S(g)$ is a generator of $H_{n-2}(\mathcal{N}_{\mathcal{P}})$; otherwise, it is* 0.

Proof. Suppose, first, that g is not an oriented cycle. Then $S(G) \in \mathcal{N}_{\mathcal{P}}$. Therefore, $[\partial S(G)] = 0 \in H_*(\mathcal{N}_{\mathcal{P}})$.

Now, suppose that g is an oriented cycle. We define a second cover of $\mathbb{R}^n \backslash \Delta$. Let

$$\mathcal{V}_j = \{ x \in \mathbb{R}^n \mid \exists i \text{ such that } x_i < x_j \} .$$

Note that $\mathcal{U}_{ij} \subset \mathcal{V}_j$; indeed, $\mathcal{V}_j = \bigcup_{i \neq j} \mathcal{U}_{ij}$. (Here we conflate $\mathcal{U}_{ij}$ with U_{ij}.) The $\mathcal{V}_j$ thus also cover $\mathbb{R}^n \backslash \Delta$. Furthermore, this cover still satisfies the conditions of the nerve theorem: all intersections are contractible or non-empty. Call the nerve of this cover $\mathcal{N}_{\mathcal{V}}$. For clarity, for this proof, we will also refer to the nerve $\mathcal{N}_{\mathcal{P}}$ with cover $\mathcal{U}$ as $\mathcal{N}_{\mathcal{U}}$.

We have a map of inclusions $\mathcal{U}_{ij} \to \mathcal{V}_j$. Since both $\mathcal{N}_{\mathcal{U}}$ and $\mathcal{N}_{\mathcal{V}}$ are homotopy equivalent to M, by the nerve theorem, we get induced maps $\mathcal{N}_{\mathcal{U}} \to \mathcal{N}_{\mathcal{V}}$ that give the homotopy equivalence $\mathcal{N}_{\mathcal{U}} \simeq \mathcal{N}_{\mathcal{V}}$.

Now, because the intersection properties of $\{\mathcal{V}_j\}_j$ are simple, it's not hard for us to figure out the geometric realization of $\mathcal{N}_{\mathcal{V}}$: all intersections are non-empty except for the intersection of all the $\mathcal{V}_j$'s. Therefore, the geometric realization is the boundary of the $n-1$-dimensional simplex, $\partial \Delta^{n-1}$.

Now, we see where $\partial S(g)$ gets mapped. When g is an oriented cycle, $\partial S(g)$ is going to be a sum of what corresponds to oriented graphs like $1 \to 2 \to \cdots \to n$, $2 \to 3 \to \cdots \to n \to 1$, etc. Each of these will be mapped to an $n-2$ dimensional face of Δ^{n-1}: $1 \to 2 \to \cdots \to n$, for example, will be mapped to the $n-2$ dimensional face containing every vertex except $\mathcal{V}_n$. Thus, we have that $\partial S(g)$ gets mapped to $\partial \Delta^{n-1}$, which is a generator of the homology. □

4.4 Topology on Profiles

So far, we have looked at the associated topology for just one voter's preferences $\mathcal{P}$. We now consider a profile of preferences $\mathcal{P}^k$. We have that $\mathcal{U}_{ij}^{\sigma}$ cover $\mathcal{P}^k$. Let $\mathcal{N}_{\mathcal{P}^k}$ be the nerve of this covering.

Now, again, we construct a manifold with an associated cover that will allow us to use the nerve theorem. Let

$$U_{ij}^{\overline{\sigma}} = \left\{ (x^{(1)}, \ldots, x^{(k)}) \,\middle|\, x^{(l)} \in \mathbb{R}^n, x_i^{(l)} > x_j^{(l)} \text{ if } \sigma = +, x_i^{(l)} < x_j^{(l)} \text{ if } \sigma = - \right\},$$

where $i < j \in X, \sigma \in \{-,+\}^n$. Let

$$MK = \bigcup_{i<j, \overline{\sigma} \in \{+,-\}^n} U_{ij}^{\overline{\sigma}}.$$

Once again, since each of these U_{ij}^{σ}s is convex, the intersection of any collection of them will be contractible. Thus, the nerve theorem applies. Likewise, we can associate the nerve of MK with the nerve of $\mathcal{P}^k$. Thus, we again are curious about MK. Note that $MK = M^k$.

Now, we calculate the topology of MK. First, we define a projection mapping $p_l : \mathcal{N}_{\mathcal{P}^k} \to NP$ that sends a tuple of comparisons to its lth component. Let c be a generator of $H^{n-2}(\mathcal{N}_{\mathcal{P}})$.

Proposition 22 (Baryshnikov Propositions 4 and 5). *$H^i(MK, \mathbb{Z}) = 0$ for $1 \leqslant i < n-2$, and $H^{n-2}(MK, \mathbb{Z}) \cong \mathbb{Z}^k$, with basis consisting of $p_l^*(c)$, $1 \leqslant l \leqslant k$.*

Proof. We use Künneth's formula; see [6], §3.2. Recall that the cohomology of M has $\mathbb{Z}$ in dimensions 0 and $n-2$, and is otherwise 0. Künneth's formula gives us the isomorphism

$$H^*(M)^{\otimes k} \to H^*(M^k)$$

given by

$$x_1 \otimes \cdots \otimes x_k \mapsto p_1^*(x_1) \smile \cdots \smile p_k^*(x_k).$$

Since the x_i can be non-trivial only in dimensions 0 and $n-2$, this gives us that $H^i(MK) = 0$ for $1 \leqslant i < n-2$. In dimension $n-2$, it's clear that the elements

$$\left\{ 1 \otimes \cdots \otimes 1 \otimes \underbrace{x}_{l^{\text{th}} \text{ spot}} \otimes 1 \otimes \cdots \otimes 1 \,\middle|\, 1 \leqslant l \leqslant k \right\}$$

form a basis, and these elements correspond precisely to the $p_l(c)$'s. □

4.5 The Axiomatic Framework and Conclusion of Proof

4.5.1 The Axiomatic Framework

The proof of Arrow's theorem relies on homological properties of nerves that we derive from coverings of sets of preferences and profiles. We can easily axiomatize the properties needed for the actual proof of the theorem.

Definition 23. *Let A and B be simplicial complexes, with simplicial maps*

$$\begin{array}{ccccc} A & \xrightarrow{D} & B & \xrightarrow{f} & A \\ & \swarrow^{p_1} & & \searrow^{p_k} & \\ A & & \cdots & & A \end{array}$$

so that the p_i's induce the isomorphisms $H^m(B) = \bigoplus H^m(A)$, the p_i induce maps i_l on homology that project onto the appropriate factor, and all composed maps $A \to B \to A$ are the identity. We call such a setup purely separated (in dimension m).

Definition 24. *Suppose we have purely separated data, as above. Fix some index $l \in \{1, \ldots, k\}$. Consider the composite map*

$$H_m(A) \xrightarrow{i_l} H_m(B) \xrightarrow{f_*} H_m(A) ,$$

where i_l is the inclusion of the lth factor. If this map is an isomorphism, we say that l is a homological dictator*; if this map is trivial, we say l is* homologically irrelevant.

We will construct in §4.5.2 simplicial complexes that give precisely this sort of purely separated data, and where $H^m(A) = \mathbb{Z}$. Then, we will show in §4.5.3 that dictators are homological dictators and non-dictators are homologically irrelevant. With these facts, the following theorem will give us Arrow's theorem.

Theorem 25. *Suppose there is purely separated data in dimension m where the mth homology group of A contains a free part, and suppose every index is either a homological dictator or is homologically irrelevant. Then there is exactly one homological dictator.*

Proof. Consider some free generator $h \in H_m(A)$. By pure separatedness, we have $D_*(h) = (h, \ldots, h) \in H_m(A)$. Therefore, since we know that every factor is either a homological dictator or is homologically irrelevant, $f_* \circ D_*(h) = \sum_{i=1}^{k} d_i h$, where d_i is 1 if i is a homological dictator and 0 if it is homologically irrelevant. Since we know that this mapping is also the identity, $\sum_{i=1}^{k} d_i = 1$, so there is exactly one dictator. $\square$

4.5.2 Pure Separatedness

Now, we can get our purely separated data. We recall that we have a function $f : \mathcal{P}^k \to \mathcal{P}$. IIA is equivalent to saying that $f(\mathcal{U}_{ij}^{\overline{\sigma}}) \subset \mathcal{U}_{ij}^{\overline{\sigma}}$ for some $\overline{\sigma} \in \{-,+\}^n$ and $\sigma \in \{-,+\}$. This therefore gives us a mapping from $\mathcal{N}_{\mathcal{P}^k}$ to $\mathcal{N}_{\mathcal{P}}$, which we will still call f. Define a diagonal mapping $D : \mathcal{N}_{\mathcal{P}} \to \mathcal{N}_{\mathcal{P}^k}$ that sends $\mathcal{U}_{ij}^{\sigma}$ to $\mathcal{U}_{ij}^{(\sigma,\ldots,\sigma)}$. We can now define a diagram of simplicial mappings:

$$\begin{array}{ccccc} \mathcal{N}_{\mathcal{P}} & \xrightarrow{D} & \mathcal{N}_{\mathcal{P}^k} & \xrightarrow{f} & \mathcal{N}_{\mathcal{P}} \\ & \swarrow^{p_1} & & \searrow^{p_k} & \\ \mathcal{N}_{\mathcal{P}} & & \cdots & & \mathcal{N}_{\mathcal{P}} \end{array}$$

Pareto gives us the fact that all composite maps from $\mathcal{N}_{\mathcal{P}}$ to itself are identical.

4.5.3 Homological Dictatorship

Now it suffices to show that our setup meets the criteria about dictatorship. Let c be the generator of $H^{n-2}(\mathcal{N}_{\mathcal{P}})$, and let h_i, $i = 1, \ldots, k$ be the dual basis of homologies corresponding to the $p_i^*(c)$ cohomology basis in $\mathcal{N}_{\mathcal{P}^k}$. We write the pairing of homology and cohomology in bracket form: $(a, b) := b(a)$, where $b \in H^*(M)$ and $a \in H_*(M)$.

Proposition 26 (Baryshnikov Proposition 7)**.** *If a voter l is a dictator, then $(h_l, f^*(c)) = 1$; otherwise, it is* 0.

Proof. Let $\vec{g} = (g_1, \ldots, g_k)$ be a k-typle of oriented graphs, each of which has the unoriented support of the cycle $1 - 2 - \cdots - n - 1$, and which is an oriented cycle for g_l and is acyclic for all other factors. This defines a simplex of n vertices $S(\vec{g})$ in $\mathcal{N}_{\mathcal{P}}$: $\mathcal{U}_{ii+1}^{\overline{\sigma}}$ has the sign σ_j determined by the direction of the edge between i and $i+1$ in graph g_j. Again, we have the boundary

$\partial S(\vec{g})$. The way we chose our $\vec{g}$, we will have that $\partial S(\vec{g})$ represents h_l in $H_{n-2}(\mathcal{N}_{\mathcal{P}^k})$.

Now, consider the image of $\partial S(\vec{g})$ under f_*. This is going to be a cycle in $\mathcal{N}_{\mathcal{P}}$, and thus correspond to an oriented graph g with the same unoriented support. Suppose g is not oriented. Then each simplex in $\partial S(\vec{g})$ is going to agree, under f_*, on some vertex, say $\mathcal{U}_{12}^{\sigma}$. Suppose, first, that they all have $\mathcal{U}_{12}^{+}$. Then f sends a profile with $2 \to \cdots \to n \to 1$ to a result with 1 preferred over 2, and thus l is not a dictator. Now suppose that they all have $\mathcal{U}_{12}^{-}$. But then f_* maps $1 \to 2 \to \cdots \to n$ to something that l disagrees with, and so l is not a dictator.

We note that $(h_l, f^*(c)) = (f_*(h_l), c) = (f*(h(\vec{g})), c)$. Now, $(f_*(h(\vec{g})), c) = 0$ is equivalent to g not being an oriented cycle. This completes one direction.

Now we show that if $(h_l, f^*(c)) = 1$, then l is a dictator. We have that the corresponding g, as above, now must be oriented. We can vary the components g_m, $m \neq l$ for $\vec{g}$ without changing g, since g must remain an oriented cycle and by IIA, each change only affects one arrow at a time. Thus, we must have that, for any $i \to j$, l is a dictator. □

4.5.4 Proof of Arrow's Theorem

Arrow's theorem is now immediate. Suppose for the sake of contradiction that every voter were a non-dictator. Then, by Proposition 26, every voter is homologically irrelevant. In §4.5.2 we finished the construction of purely separated data, which allows to use Theorem 25 to show that there is a homological dictator. This gives a contradiction, completing the proof. □

References

[1] David Austen-Smith and Jeffrey S. Banks, *Positive political theory 1: Collective preference*, University of Michigan Press, Ann Arbor, 1999.

[2] Yuliy M. Baryshnikov, *Unifying impossibility theorems: A topological approach*, Advances in Applied Mathematics **14** (1993), 404–415.

[3] Graciela Chichilnisky and Geoffrey Heal, *Necessary and sufficient conditions for a resolution of the social choice paradox*, Journal of Economic Theory **31** (1982), 68–87.

[4] Graciela Chichilnisky, *Social choice and the topology of spaces of preferences*, Advances in Mathematics **37** (1980), 165–176.

[5] Wulf Gaertner, *A primer in social choice theory*, Oxford University Press, Oxford, 2006.

[6] Allen Hatcher, *Algebraic topology*, Cambridge University Press, New York, 2002.

[7] Charles D. Horvath, *On the topological social choice problem*, Social Choice and Welfare **18** (2001), 227–250.

[8] Rafe H. Kinsey, *Generalized social choice: Partial preferences*, 2008. `http://www.stanford.edu/~rkinsey/partial-prefs.pdf`.

[9] Luc Lauwers, *Topological social choice*, Mathematical Social Sciences **40** (2000), 1–39.

[10] Paras Mehta, *Topological methods in social choice theory: An overview*, Social Choice and Welfare **14** (1997), 233–243.

[11] James R. Munkres, *Topology*, 2nd ed., Prentice Hall, Upper Saddle River, NJ, 2000.

[12] Shmuel Weinberger, *On the topological social choice model*, Journal of Economic Theory **115** (2004), 377–384.