SOCIAL CHOICE *and the* CONTINUITY AXIOM

Heine Rasmussen

Thesis submitted for the Dr. oecon. degree

NHH

SOCIAL CHOICE AND THE CONTINUITY AXIOM

HEINE RASMUSSEN

Norwegian School of Economics and Business Administration
Bergen, 1994

Preface

This thesis has benefited from discussions with several people. I am very grateful to my supervisor Professor Terje Lensberg, in this and many other regards. While working on the thesis, I spent nearly one year at Columbia University, and their hospitality is gratefully acknowledged. In particular, Professors Graciela Chichilnisky and Geoffrey Heal were always very accommodating towards me, and many of the ideas presented here arose through discussions with them. Finally I would like to thank Professor Kurt Jörnsten, who is a member of the thesis committee.

I am thankful for financial support from the Norwegian School of Economics and Business Administration, and the Foundation for Research in Economics and Business Administration.

Contents

1 | Introduction

This treatise is more of an inquiry into the structure of economic theory than the structure of economic reality. Its main theme is the investigation of mutually inconsistent properties of collective decision procedures. It is thus in the tradition initiated by Arrow's impossibility paradox [1]. As such, most of the results we present are of a kind that says that decision or allocation procedures with certain properties do not exist.

They do not exist, not because they are exceedingly difficult to implement, or because they contradict some laws of nature, but because they cannot *possibly* exist, by purely logical considerations. All of the results in the following pages are derived by mathematical methods only. The arguments employed are not only abstract and formal; in a certain sense, they say nothing at all about the real world. So, how is it possible for these results to be relevant or interesting in a field like economics, which aspires to be an empirical science with at least some of the predictive power found in natural sciences like physics or chemistry?

This objection is certainly not unique to social choice theory. It can be raised to almost all of mathematical economics — for instance, it is frequently pointed out that general equilibrium theory by itself implies very little else than the existence of an equilibrium within the model. Even Wassily Leontief maintained that "not having been subjected from the outset to the harsh discipline of systematic fact-finding, traditionally imposed on and accepted by their colleagues in the natural and historical sciences, economists developed a nearly irresistible predilection for deductive reasoning." He goes on to criticize this deductive approach for not "being able to advance, in any perceptible way, a systematic understanding of the structure and the operations of a real economic system."

True, the proportion between deductive reasoning on one hand, and empirical research on the other, is of a quite different magnitude in economics than in the natural sciences. But to claim that this is in some sense only due to the historical background of economic science is an entirely different matter — it should be evident that the explanation must be sought among more substantial causes. It appears to be two main reasons to account for this difference between economics on one hand, and the natural sciences on the other.

First, since the economic systems under investigation are so extremely complicated, we tend to make simplified and highly abstract models that are better suited to provide understanding than prediction. Also, since the models for this reason are farther "removed" from the real systems then what they are in most natural sciences, there is ample scope for the existence of several different models of one and the same system, each emphasizing separate aspects of the system. In physics, an alternative model of a natural phenomenon is a sensation; in economics, every issue of any journal abound with alternative models. And the only way to gain understanding from a formal model is by deduction.

Social choice theory is primarily an attempt to model collective decision processes using tools from economic theory, and is thus a part of economics. All the so-called impossibility results in this theory essentially demonstrate that certain combinations of formulae and statements cannot work as models in the theory without introducing inconsistencies. While this says nothing about the world we try to model, it may say something very interesting about the theory in which we create the models. This is particularly relevant in a social science like economics, where we not only use the theory to create models that describe existing systems, but to an equal degree try to implement real systems that conform to certain models (by deregulating markets, invoking anti-trust laws, etc.). In such cases it is of course extremely important to know what may or may not constitute a model, and what properties a model has. This is the second reason for the prevalence of mathematical deduction in economics. While inquiries of this kind are not empirical, they are essential prerequisites for any empirical study: They establish which concepts can be meaningfully subjected to empirical investigation.

With these methodological remarks behind us, the remaining part of this introduction summarizes the contents of the chapters that follow. The reference to the "continuity axiom" in the title of this work signifies one of the key assumptions made: We require that social welfare functions shall be continuous. This obviously presupposes that a topology is defined on the class of preferences. In Chapter 2 we survey and discuss a few of the most important contributions to the field. Sections 2.2 and 2.3 describes some often used topologies for spaces of preferences. In Section 2.4 we look at their connection to social choice problems, and in Section 2.5 we comment on a debate on the relevance of topological methods in social choice theory.

In Chapter 3, we introduce a new class of topologies for preference spaces. Let X be the choice space, e.g., a space of allocations. We assume that a measure is defined on $X \times X$. We then define the distance between two preferences to be the measure of the symmetric set difference of the graphs of the preferences. The possible domain of preferences is thus very large; the only requirement for a preference to be a member of the domain is that it should have a measurable graph. In particular, typical properties like nonsaturation, continuity, transitivity, or completeness, that are required in many of the other topologies, are not prerequisites with this definition of distance.

It is easily verified that the distance function above has all the properties of a pseudometric. On the other hand, the distance between two different preferences may actually be zero, if the preferences differ by a set of zero mass. But if we define a binary relation between preferences that holds if and only if the distance between two preferences is zero, then this relation is an equivalence relation, and will thus partition the preference space into equivalence classes, where the distance between any two preferences in a particular class is zero. If we now define the distance between two equivalence classes to be the distance between two arbitrarily chosen preferences, one from each of the two classes, the distance becomes a proper metric on the space of equivalence classes.

In general, different measures will generate different metrics. However, if two measures are absolutely continuous with respect to each other, they generate the same topology, and even when the measures are not equivalent and different topologies are generated, these topologies still have many properties in common. We will call them *measure-based topologies.*

As mentioned, the measure-based topologies do not distinguish between preferences that are equal almost everywhere. On the other hand, they do distinguish between all continuous, transitive and complete preferences, as long as the choice space is connected and the generating measure assigns positive mass to open sets, since the symmetric set difference of any two such preferences has a nonempty interior.

One criterion one may use to evaluate topologies for preference spaces is to which extent subsets of preferences with well defined choice theoretic properties are closed in the topology. We show that the space of all preferences with measurable graphs is a complete space. The subset consisting of equivalence classes containing at least one transitive preference is closed, as is the subset of equivalence classes containing at least one complete preference. Even the subset consisting of equivalence classes containing at least one transitive and complete preference is closed.

In Section 3.4 we give an axiomatic characterization of the measure-based pseudometrics. Three conditions are listed, and a pseudometric on preference spaces is a measure-based pseudometric if and only if it satisfies these three conditions. The first condition says that the pseudometric should satisfy a weak convergence criterion, continuity from above. The remaining two conditions concern global properties of the pseudometric. The second condition says that if two preferences agree on the ranking of a pair of alternatives, it should not matter how they rank the alternatives as far as the distance is concerned. The rationale behind this condition is that we want the distance to measure the extent of *disagreement* between preferences, and it should not be influenced by irrelevant information about the ranking of alternatives they agree upon.

The third condition involves the concept of Pareto efficient preferences. We say that a preference P_3 is Pareto efficient relative to two preferences

P_1 and P_2 if whenever P_1 and P_2 agree on the ranking of a pair of alternatives, P_3 ranks this pair in the same way. The condition says that in this case, the distance between P_1 and P_2 should equal the sum of the distance between P_1 and P_3, and the distance between P_3 and P_2. This will ensure that when P_3 is Pareto efficient relative to P_1 and P_2, it is not possible to find a fourth preference that is closer to both P_1 and P_2 than P_3 is.

In the last section of Chapter 3 we investigate continuous aggregation of preferences under the measure-based topologies. We show that the space of transitive and complete preferences without "thick" indifference surfaces, i.e., preferences where all indifference surfaces has zero mass, admits continuous aggregation rules that respect unanimity and is anonymous. This is also true for the space of all continuous, complete and transitive preferences without thick indifference surfaces.

The three remaining chapters all investigate strategy-proofness of social welfare functions, where the agents are allowed to take strategic considerations when revealing their preferences; in other words, an agent is assumed to report the preference that gives him the the best possible social outcome, and not necessarily his "true" preference.

In Chapter 4, we analyze this idea on its most general level. It is known (Chichilnisky & Heal [15]) that if the preference space is topologized in a manner that makes it homeomorphic to an n-sphere, and there is m agents, then, for any continuous social welfare function that satisfy a degree condition, there is always an agent that can achieve any outcome he wants, no matter what preferences the other $m-1$ agents disclose (in general, he will have to misrepresent his preference in order to do this, of course).

This result depends on properties of the function that are only meaningful when the space of preferences is homeomorphic to an n-sphere. The contribution of Chapter 4 is to reformulate the framework in a way that makes it applicable to any space. Assume there are two agents, both with preferences (with unique maxima) over a space Y of social outcomes. An aggregation map f from $Y \times Y$ to Y now gives rise to a two-person noncooperative game, where the possible moves for both players are the points in Y, and with outcome $f(y_1, y_2)$ if the players' moves are y_1 and y_2. An aggregation map is called strategy-proof for a given pair of preferences over Y if it is a Nash equilibrium in this game that both players report their most preferred point in Y.

We define an exhaustive class of preferences over Y as a collection of preferences so that every point in Y is the maximum of some preference in the collection. Given an exhaustive class of preferences, we say that f is strategy-proof for this class if it is strategy-proof for any pair of preferences, both members of the class. We then show that for a certain kind of spaces, retracted H'-spaces, if an aggregation map respects unanimity and is strategy-proof for an arbitrary exhaustive class of preferences, it must be dictatorial. Retracted H'-spaces can be regarded as generalized n-spheres, and this establishes the connection with the result of Chichilnisky & Heal.

In Chapters 5 and 6, we consider the same kind of questions as in Chapter 4, but the assumptions we make about the nature of the preferences are more specific. The concept of a metapreference is introduced in Chapter 5. We now assume that the space Y that is being aggregated upon is a space of preferences. If we want to analyze strategic disclosure of these preference, we need to make assumptions about preferences at a higher level — preferences that have Y as domain, and are used by individuals to rank social preferences (points in Y). We use the term metapreferences for these latter preferences whenever Y is assumed to be a space of preferences.

The main result of Chapter 5 is a theorem that shows how a metapreference over a space of preferences on a choice space X can be represented by a measure on $X \times X$. If we choose an arbitrary finite measure on $X \times X$, and assume that a preference P is singled out, we can define a utility function on the space of preferences by letting the utility of a preference Q be equal to the negative mass of the symmetric set difference between P and Q. This utility function generates a metapreference in the obvious way. The surprising result is that *any* metapreference that satisfy some weak conditions can be generated by this method. This means that we, when analyzing strategic behavior, can work with measures on $X \times X$ instead of the more abstract concept of a metapreference.

The four conditions that characterize the class of metapreferences that can be represented by a measure in such a manner, can be summarized as follows. The first condition says that that only the differences between preferences should determine how they are ranked by the metapreference. So if two preferences P and Q agree on the ranking of two alternatives x and y, it should not matter *how* they rank x and y (i.e., whether it is $x \succsim y$ or $x \not\succsim y$) as far as the metapreferences are concerned.

The second condition says that the metapreferences should have a maximal element, i.e., that there is some preference (to be interpreted as the agents "own" preference) that is considered at least as good as all other preferences. The two remaining conditions are of technical importance only.

In Chapter 6, we use the metapreferences that can be generated by measures to analyze strategy-proofness of social welfare functions. The main result of this chapter appears in Section 6.3, and it says that if all preferences are assumed to be linear orderings, and a social welfare function is onto and strategy-proof for the class of all metapreferences that can be generated by measures on $X \times X$, then the social welfare function must be dictatorial. In Section 6.4, we discuss the possibilities of extending this result to preferences that are not linear orderings.

Finally, observe that the definition of some mathematical concepts can be found in a separate appendix, whenever they are not defined within the text itself.

2 | Topology and preferences

2.1 INTRODUCTION

Since its beginning, with the works of J. C. de Borda and M. de Condorcet in 1781 and -85, social choice theory has mainly been concerned with choices among a discrete set of alternatives, and the methods employed are usually some variants of combinatorics. This is in sharp contrast to almost all other areas of economics, where one deals with infinite sets endowed with a topology, and where the methods rely on real and functional analysis.

This separation has several consequences. For one thing, it means that it is difficult to integrate results in social choice theory with results from other fields in economics. When the choice set is infinite, and thus admits infinitely many different rankings of the feasible choices, it also means that social choice theory misses an important aspect of decision processes, whether these are due to markets or public decisions: continuity. If we want a theory to have any predictive or explanatory capabilities, we should ensure that small errors of observation do not lead to large errors in the predicted outcome or set of outcomes. In other words, the transformation from observations to outcomes should be a continuous mapping, and this presupposes a suitably defined topology. This has been the standard approach in the theory of competitive markets. An economy in its most general form can be described as a set of agents, each with an initial endowment of commodities. To such an economy we associate a set of price vectors (the possible equilibria of the economy). These equilibria should result from continuous transformations on the set of economies. In the same manner, we would like voting rules (or other procedures that aggregate preferences) to be continuous, so that small errors in the observation of individual preferences do not lead to significant changes in the outcome of the aggregation.

In the following sections, we survey and discuss a few of the most important contributions regarding topological preference spaces. The survey is not meant to be exhaustive; in fact, it is more or less limited to those results the reader should be aware of in order to follow the rest of this text.

The choice space, or space of alternatives, will in general be called X. A preference is any preorder over X, but most authors only consider preferences that satisfy some additional restrictions. The space of preferences is called $\mathcal{P}$, but observe that the nature of this space may vary from author to author. Since some writers identify a preference with the preorder itself, and others identify it with the graph of the corresponding preorder, a preference will be denoted in a generic fashion by an italic letter like P, and the associated relation by $\succsim_P$, with $\sim_P$ and $\succ_P$ being the symmetric and antisymmetric parts, respectively.

A preference is said to be *continuous* if its graph is closed in $X \times X$. In the case where X is a subset of $\mathbb{R}^n$, a preference is *monotone* if $x \gg y$ implies $x \succ y$.

2.2 PREFERENCE SPACES

In general, the "reasonableness" of a topology depends on how well it interacts with other structural properties of the space. As an example, take the usual topology on $\mathbb{R}^n$; this topology can be derived from the algebraic structure on $\mathbb{R}^n$ in a straightforward manner.

For spaces without a strong natural nontopological structure, there is usually no single topology that stands out as the "obvious" one. E.g, the family of all subsets of an arbitrary space has (as a starting-point) little other natural structure than the lattice induced by set inclusion ($\subset$). On such families, there are several topologies with equal status.

When constructing a topology for spaces of preferences, one natural method is to relate it to the topology of the choice space X. This is the approach taken by Kannai [28], who was the first to consider sets of preferences as topological spaces. The topology he proposes is the smallest topology that makes the set $\{(x, y, P) : x \succ_P y\}$ open in $X \times X \times \mathcal{P}$. This is the same as requiring that if $x \succ_P y$, $x_n \to x$, $y_n \to y$, and $P_n \to P$, then there is an m such that $x_n \succ_{P_n} y_n$ for all $n \geq m$.

The Kannai topology has a pleasant property: A natural and important subset of $\mathcal{P}$ is metrizable. Let $\mathcal{Q} \subset \mathcal{P}$ be the set of continuous and monotone preferences. The choice space X is now assumed to be the positive orthant of $\mathbb{R}^n$. The subspace topology on $\mathcal{Q}$ induced by the Kannai topology can also be induced by a metric on $\mathcal{Q}$. Every preference P in $\mathcal{Q}$ can be identified with a retraction f_P from X to the diagonal of X in the following way: For any $x \in X$, let $f_P(x)$ be the unique y in the diagonal of X that satisfies $x \sim_P y$.

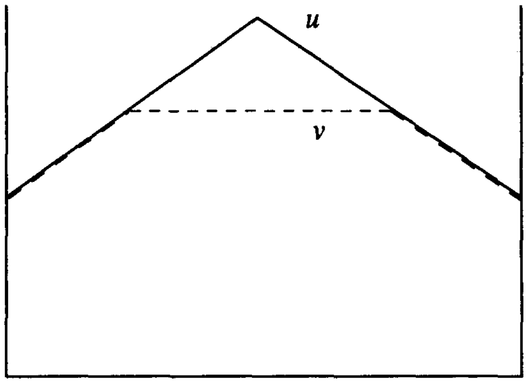

Figure 2.1: The Kannai topology does not distinguish between the two preferences generated by the utility functions u and v.

For every $P \in \Omega$, we can now construct a unique utility function $u_P: X \to \mathbb{R}^1$ as $u_P(x) = \|f_P(x)\|$. Finally, if we let

$$d(P_1, P_2) = \sup_{x \in X} \frac{\|u_{P_1}(x) - u_{P_2}(x)\|}{1 + \|x\|^2},$$

then d is readily seen to be a metric on Ω. The equivalence of the topology induced by d and the Kannai topology is shown in Kannai [**28**, Theorem 3.2].

However, when a space of preferences includes locally saturated preferences, the space is not Hausdorff when endowed with the Kannai topology, as Le Breton [**30**] illustrates. An example is given in Figure 2.1. The choice space is here a closed interval on the real line, and the Kannai topology does not separate the two preferences generated by the utility functions u and v. All neighborhoods of the preference generated by v contains the preference generated by u.

A larger topology is introduced in Debreu [**17**]. Here, the space of continuous preferences is topologized by identifying preferences with their graphs, and endowing the space of graphs with the Hausdorff distance (Hausdorff [**25**, p. 166]). If δ is a metric on $X \times X$, let ρ be defined as

$$\rho(A, B) = \sup_{b \in B} \delta(A, b),$$

where $A, B \subset X \times X$. In general, whenever X is unbounded, the image of ρ is the nonnegative part of the extended reals. The Hausdorff distance d can now be written as

$$d(A, B) = \max\{\rho(A, B), \rho(B, A)\}.$$

It is easily verified that d has all the properties of a metric except that of always being finite.

The space of preferences endowed with the topology of the Hausdorff distance is in general not a separable space whenever it contains preferences with noncompact graphs. A topology that renders the space separable also in this case is introduced in Hildenbrand [**27**], the *closed convergence topology.* Like the Hausdorff distance, the closed convergence topology is defined on the space of preference graphs. For a sequence of sets $\{A_n\}$ in $X \times X$, define the *superior closed limit* (lim sup) to be the set of all $x \in X \times X$ such that every neighborhood of x has points in common with infinitely many A_n, and the *inferior closed limit* (lim inf) to be the set of all $x \in X \times X$ such that every neighborhood of x has points in common with all but a finite number of the A_n. In the closed convergence topology, a sequence $\{A_n\}$ converges to A if

$$\limsup A_n = \liminf A_n = A.$$

It is easily verified that this topology agrees with the Kannai topology on the space of continuous and monotone preferences, where the choice space is the positive orthant of $\mathbb{R}^n$. Furthermore, the set

$$\{(x, y, P) : x \succ_P y\}$$

is open in $X \times X \times \mathcal{P}$, so the closed convergence topology is larger then the Kannai topology. It is smaller than, but closely related to, the topology induced by the Hausdorff distance; the two topologies coincide when $X \times X$ is a compact metric space, and when $X \times X$ is only locally compact, the space of preferences with the closed convergence topology can be embedded in the space of preferences with graphs in the one-point compactification of $X \times X$ endowed with the topology of the Hausdorff distance (see [**27**]).

Unfortunately, when applied to spaces of preferences with saturated points, the Hausdorff distance (and the closed convergence topology) have some questionable convergence properties. Consider the following example. Let a sequence of preferences $\{P_n\}$ be generated by a sequence of utility functions $\{u_n\}$ on $[0, 1]$ defined as follows: For each n, partition the interval $\langle 0, 1]$ into n consecutive half open intervals E_i^n, $i = 1, \ldots, n$, i.e.,

$$E_i^n = \langle (i-1)/n, i/n].$$

For all n, if $x \neq 0$ let the integer function $\beta_n(x)$ be equal to the index of the set E_i^n that has x as a member, that is, the integer that makes $x \in E_{\beta_n(x)}^n$ hold. Define u_n by $u_n(0) = 0$, and, for $0 < x \leq 1$,

$$u_n(x) = \begin{cases} nx - \beta_n(x) + 1 & \text{for } \beta_n(x) \text{ odd,} \\ \beta_n(x) - nx & \text{for } \beta_n(x) \text{ even.} \end{cases}$$

Figure 2.2 shows the first elements of the sequence, the three functions u_1, u_2, and u_3.

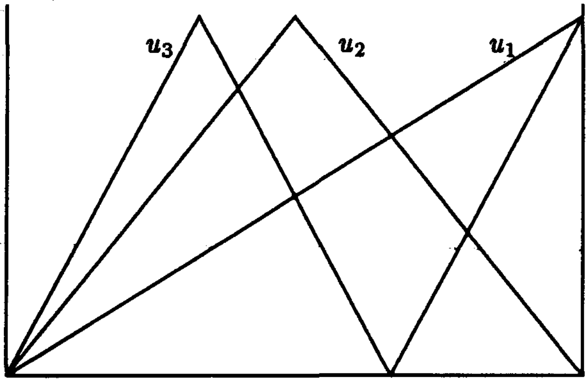

Figure 2.2.

In the topology of the Hausdorff distance, $\{P_n\}$ converges to the trivial preference, i.e., the preference that is indifferent between all alternatives. One can certainly raise legitimate objections to a topology where $\{P_n\}$ has this limit. The preferences in the sequence indicate very complex behavior, while the limit preference indicates very simple behavior (or perhaps one should say no behavior at all). For two arbitrary points x and y, it will be increasingly difficult to predict the ranking of these two points as n goes to infinity if we do not have exact knowledge of the location of the points. In contrast, with the limit preference, the location of the points does not matter; all alternatives are tied. It seems that the best solution would be a topology where $\{P_n\}$ did not converge at all. This is the case with the measure-based topologies presented in Chapter 3.

2.3 SMOOTH PREFERENCES

Even if smoothness is a property usually associated with functions or manifolds, one can imagine various intuitive notions of smoothness applied to preferences as well. However, we know that preferences can be mathematically represented in several ways, and constructing a formal definition of differentiability that agrees on all of these representations seems at the outset to be a nontrivial task.

In Debreu [18], three ways of approaching the question are considered. Preferences are here assumed to be monotone and continuous.

First, a smooth preference can be regarded as a C^1 (i.e., continuous and at least once differentiable) vector field g on the choice space X (which is assumed to be the positive cone of $\mathbb{R}^n$), normalized in such a way that $\|g(x)\| = 1$ everywhere. Such a vector field is obviously equivalent to a map of class C^1 from X to the unit sphere S^{n-1}. Informally, the vector at a point $x \in X$ is orthogonal to the indifference hypersurface through x, and is pointing in the direction of preference. The space of smooth preferences can then be topologized by for instance the C^1 sup norm on vector fields.

It also natural to investigate the possible existence of a utility function u from X to R that is C^2 and has a gradient that is everywhere a positive multiple of $g(x)$. It can be shown that such functions exist whenever $g(x)$ has strictly positive coordinates for all x in X, and in addition satisfies a local integrability condition closely related to a theorem of Frobenius (see Debreu [18] and Chipman *et al.* [16, ch. 9]). This means that there is a second way to represent certain smooth monotone preferences — as C^2 utility functions. This representation will in general not be unique, but since the relation between functions of representing the same preference obviously is an equivalence relation, we can make the representation unique by considering equivalence classes, or families, of functions.

A third approach to smooth preferences consists of making assumptions on their graphs. Following [18], if the boundary of the graph of a continuous, monotone, and complete preorder is a C^2-hypersurface in $\mathbb{R}^{2n}$, the preorder is said to be a preference relation of class C^2.

We have now seen three different approaches to smooth and monotone preferences: A locally integrable and normalized C^1 vector field with strictly positive coordinates, a family of monotone C^2 utility functions, and a preference relation of class C^2. Debreu [18] proves that all three notions are equivalent, in the sense that we can postulate an object of one kind and then derive objects of the two other kinds. This represents a definite solution to the ambiguity of smoothness for preferences.

2.4 CONTINUITY AND SOCIAL CHOICE

Graciela Chichilnisky's article from 1980 [9] is one of the first where continuity is considered as a requirement for social welfare functions. She defines a social welfare function with n agents to be a continuous map ϕ from an n-fold cartesian product $\mathcal{P}^n$ of a preference space $\mathcal{P}$, and into $\mathcal{P}$. She then investigates the existence of social welfare functions that have the following two properties:

(i) ϕ respects unanimity, i.e., $\phi(P, \ldots, P) = P$.

(ii) ϕ is anonymous, i.e., $\phi(P_1, \ldots, P_n) = \phi(P_{\sigma(1)}, \ldots, P_{\sigma(n)})$, where σ is any permutation on $\{1, \ldots, n\}$.

A similar approach has also received some attention in the mathematical literature: A map that satisfies (i) is typically called *idempotent*, a map

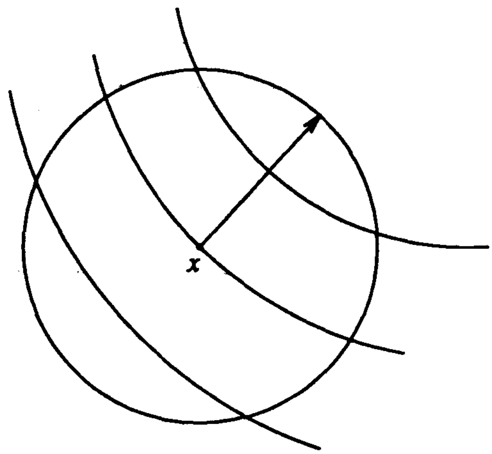

Figure 2.3: This preference maps x to a point in S^1

that satisfies (ii) is called *symmetric*, and one that satisfies (i) and (ii) is called an n-mean. Following Eckmann [19], a space that admits an n-mean is called an M_n-space. Chichilnisky's investigation is thus a special case of the more general problem of characterizing M_n-spaces.

In Chichilnisky's work, preferences are represented by a C^r ($r \geq 1$) locally integrable vector field over the choice space, with vector lengths normalized to unity. As discussed in the previous section, this representation was also considered by Debreu [18]. Chichilnisky, however, replaces the requirement of monotone preferences with one of locally nonsaturated preferences. Informally, the vector at a point x in the choice space is defined to be perpendicular to the indifference surface through x, i.e, it has the same direction as the gradient at x of any utility function locally representing the preference. A preference can then be regarded as a map from the n-dimensional choice space to S^{n-1} (see Figure 2.3). If the choice space is compact, the space of preferences can be topologized with the C^r sup norm on C^r vector fields. If the choice space is not compact (e.g. the positive orthant of $\mathbb{R}^n$), a different topology is needed; in Chichilnisky [7], a Sobolev-Hilbert manifold structure on noncompact spaces is employed to give results analogous to the compact case.

This topology excludes preferences with saturation points, i.e. with local maxima, minima, or saddlepoints, since at these points the vector field would vanish.

It is easy to verify that the property of being an M_n-space is preserved under retractions; in other words, if $\mathcal{P}$ is an M_n-space and $\mathcal{R}$ is a retract of $\mathcal{P}$, then $\mathcal{R}$ is an M_n-space, since an n-mean ϕ on $\mathcal{P}$ induces an n-mean $r \circ \phi$ on $\mathcal{R}$, where r is a retraction. It is furthermore clear that the subspace of linear preferences[1] is a retract of the space of preferences $\mathcal{P}$ (choose an arbitrary point x in the choice space, and let the retraction be the map that takes a preference P to the linear preference that maps x to the same point in S^{n-1} as P does). But in Chichilnisky's topology, the subspace of linear preferences is clearly homeomorphic to S^{n-1}, so if S^{n-1} is not an M_n-space, neither is $\mathcal{P}$.

In [9], Chichilnisky shows that S^m is not an M_n-space for all $m \geq 1$ (see also the 1943 paper by Aumann [2]). In Chichilnisky & Heal [14], the authors show that if the space of preferences $\mathcal{P}$ is a CW-complex with a convex hull that is also a CW-complex, then contractibility of each component of $\mathcal{P}$ is sufficient for $\mathcal{P}$ to be an M_n-space. If $\mathcal{P}$ is a para-finite CW-complex, this condition is also necessary. Their theorem generalizes a result in Eckmann [19, p. 336], where he states that contractibility is necessary and sufficient for a finite polyhedron to be an M_n-space[2].

[1] That is, preferences where the vector field takes on a constant value at all points in the choice space.

[2] See also the article by Eckmann *et al.* [20], where the same problem is analyzed in a category theoretic framework.

2.5 CRITICISM AND DEBATE

Some objections have been raised both to the relevance of the noncontractibility result of Chichilnisky & Heal [14], and of the topological approach to social choice in general. Objections of the first kind occur in two articles by Le Breton & Uriarte [31, 32].

Of course, being a theorem, the correctness of the argument of Chichilnisky & Heal is beyond debate. However, Le Breton & Uriarte question the relevance of this result, as they feel that the noncontractibility of the preference space in the framework of Chichilnisky & Heal is due to their choice of topology and the restricted domain of preferences. Le Breton & Uriarte seem to maintain that noncontractibility is not a typical property of preference spaces, and is possessed only by certain subspaces or for certain topologies.

To support their argument, they topologize the space $\mathcal{P}$ of all continuous and complete preorders over a choice space X with the closed convergence topology. They then show that the subspace of strictly convex preferences (with possibly one saturation point) is an M_n-space. This is not surprising, since the subspace is obviously contractible. They also give a technical definition of another subspace that is dense in $\mathcal{P}$, and show that this subspace is an M_n-space as well.

They then claim that since the subspace is dense in the space of all continuous and complete preorders, this gives an approximate solution of the Chichilnisky problem. However, as they point out, the space $\mathcal{P}$ is not complete, so the theorem on extension of uniformly continuous functions does not apply.

The theorem of Chichilnisky & Heal has an interesting consequence: Even though the conditions are necessary and sufficient, it is in general easier to show that a preference space is not an M_n-space, than to demonstrate that a space *is* an M_n-space. This is partly due to the fact that the theorem requires that the space should be a para-finite CW-complex. For many preference spaces, this can be difficult to show. But if the space is not an M_n-space, one can usually find a retract that is a much simpler space, and easily seen to be a noncontractible, para-finite CW-complex. Since the property of being an M_n-space is preserved under retractions, the original space can then not be M_n.

If the space is an M_n-space, retractions will not be of any help, of course. We are then left with two strategies: (i) Show that the space is a CW-complex (the complexity of this is very dependent upon the topology), and then construct a homotopy that demonstrates contractibility, or (ii) construct an n-mean directly. Preference spaces are in general so complex that both of these strategies may be difficult to follow.

Even though there are some shortcomings in their arguments, Le Breton & Uriarte address an important question: Is it the case that unrestricted preference space with natural topologies are *in general* rendered non-contractible by the topologies? Even though this is more a matter of opinion than of mathematical deduction (since the concept of a "natural" topology is not a mathematical one), in view of the results in the following chapter I believe the question may well be answered with a "no". This, of course, does not in any way invalidate or reduce the relevance of Chichilnisky & Heal's result.

Objections of the second kind can be found in Baigent & Huang [4]. Among other things, they claim that a topological framework is not the right approach for the analysis of issues involving the proposition that large changes in the social preference should not result from small changes in individual preferences. Their main argument to support this claim seems to be that they think it is unlikely that it is possible to find one particular topology that best formalizes the intuitive notion of "closeness" of preferences. They write,

> "The greatest merit of topological analysis is that it permits very general and undemanding ways of expressing continuity. However, for spaces such as preferences, this same generality makes it very difficult to know whether any particular topology does accord with our basic intuitions concerning closeness. If this were not the case, then presumably it would be possible to formulate axioms for a topology on preferences and even state a characterization theorem. That this has not been done, in an area in which axioms are ubiquitous, strongly suggests to us that a topological framework is not the most appropriate for expressing our intuitions concerning closeness of preferences."

We present a class of topologies together with such a characterization theorem in Chapter 3.

It is important to distinguish between the problem of choosing a topology for a particular class of preferences on one hand, and on the other, showing that if a topological space has certain properties, then certain maps do not exist. The appeal of such impossibility results is already subjective in nature. They say that certain combinations of desirable properties are inconsistent. But to what extent the properties are desirable, and thus how troubled one should be by this inconsistency, is something that is an attribute of the individual reader and is, in a sense, beyond mathematical deduction or empirical research. In this perspective, there is no need for an agreement upon what constitutes the "correct" topology for preference spaces. The fact that many different topologies have been proposed does not imply that it it meaningless to topologize preference spaces — it merely means that the reader is free to interpret the impossibility results in different ways, by using the preferred topology of his choice. There appears to be no obvious topology for such spaces (except for certain subspaces) because the intuitive notion of "convergence" is itself an ambiguous and subjective concept when it is applied to preferences.

However, it is in fact reasonable to maintain that continuity of a map is too general a concept to ensure that the effect of observational errors will be negligible. This is especially relevant for incomplete metric spaces; as an example, take the space $\mathbb{R} - \{0\}$, and define a function f on this space by $f(x) = 0$ for $x < 0$ and $f(x) = 1$ for $x > 0$. This is a continuous function, but it seems difficult to argue that observational errors should have less effect on this function than on a discontinuous one. To avoid these "quasi-discontinuities" in continuous maps it is necessary to require that the maps are *uniformly continuous*. For the same reason, one should be extremely cautious with spaces that are not metrizable, or at least does not admit a uniform topology, as uniform continuity is not defined in these contexts. Of course, the impossibility result of Chichilnisky is still equally relevant, since if continuous functions with certain properties do not exist, neither do uniformly continuous ones.

In her reply to Le Breton & Uriarte [**31**] and Baigent & Huang [**4**], Chichilnisky [**12**, p. 310] criticizes the approach of Le Breton & Uriarte on the grounds that (among other things) the space they have chosen is not topologically complete, and she maintains that "the whole meaning of continuity in such a space is questionable." As the simple example in the previous paragraph shows, there is definitely something to be said for this opinion. On the other hand, the requirement that the maps should be uniformly continuous will resolve this.

3 | Measure-based topologies

3.1 INTRODUCTION

This chapter introduces a class of measure-based metric topologies on spaces of preferences.

In Le Breton and Uriarte [**32**] the authors are calling for an extension of the Kemeny distance[1] between preferences on a finite commodity space to the infinite case. The topologies we present here can be regarded as such an extension, although along a different line than that proposed by Le Breton and Uriarte.

Topologies for spaces of preferences have been studied by several authors. They have partly been motivated by problems in the general theory of economic equilibria (e.g. Kannai [**28**], Debreu [**17, 18**], Hildenbrand [**27**]), as in the study of the continuity of the core, and partly by normative problems involving social decision rules and aggregation of preferences (Chichilnisky [**8, 9, 10, 11, 14**] and several other papers by the same author; Uriarte [**40**], Le Breton & Uriarte [**31**]).

The spaces under consideration in this chapter will all be subsets of the class of measurable preferences on a topological space X (i.e. the class

[1] See Kemeny & Snell [**29**].

of preferences having a graph that is a measurable set in $X \times X$). We further assume that $X \times X$ is endowed with a finite Borel-measure[2] μ^*. Let μ be the completion of μ^*. A pseudometric d_μ is then defined on the measurable subsets of $X \times X$ as

$$d_\mu(P, Q) = \mu[(P - Q) \cup (Q - P)],$$

for measurable sets P and Q. By identifying graphs differing by a set of zero measure, we get a metric space Ω_μ of equivalence classes, where each such class consists of preferences that are equal almost everywhere. Topologies on classes of sets defined in this way are sometimes called the *fine* topologies; such spaces are of course homeomorphic to the subset of $L^1(X \times X, \mu)$ consisting of the characteristic functions on $X \times X$. These topologies, defined on general spaces of measurable sets, form an important part of measure theory. The contribution of this chapter lies in their application to spaces of preferences.

In general, different measures may generate different topologies, but it is easily seen that two topologies are equivalent if the generating measures are equivalent (i.e., absolutely continuous with respect to each other). In what follows, topologies defined according to the procedure described above will be referred to as *measure-based topologies.*

The measure-based topologies do not distinguish between preferences that are, in a sense, "equal almost everywhere". However, they *do* distinguish between preferences that are continuous, complete and transitive, whenever the generating measure assigns positive measure to open sets (see Section 3.3).

[2]The Borel-σ-algebra is chosen because this will ensure that the topologies are well-defined on the interesting class of continuous preferences.

The measure-based topologies can be given a very intuitive interpretation that makes them an attractive choice in problems concerning aggregation of preferences: Let ν be a probability measure on X, and endow $X \times X$ with the product measure $\nu \times \nu$ generated by ν. The distance $d_{\nu \times \nu}$ between two preferences is then equal to the probability that the preferences will differ in the ranking of two alternatives x and y drawn independently from the distribution over X generated by ν, i.e., the probability that $x \succsim y$ for one of the preferences and $x \prec y$ for the other.

This also suggests a general principle for selecting a topology generating measure; it can reflect the likelihood of how often a particular point in choice space is expected to be among the feasible alternatives to be decided upon. The exact value of such a measure is of less importance, since, as we have already noted, equivalent measures generate the same topology.

Section 3.2 establishes a convenient notation, and also gives a more formal presentation of the definitions introduced above.

In Section 3.3, we demonstrate some general properties of the measure-based topologies. It is shown that Ω_μ is a complete space, and that some important subsets of Ω_μ are closed. Furthermore, if X is connected and μ assigns positive measure to open sets, then d_μ is a proper metric when restricted to the class of continuous, complete and transitive preferences (i.e. any two of these preferences differ by a set of positive measure).

Section 3.4 gives an axiomatic characterization of the measure-based topologies, and Section 3.5 discusses the existence of social choice rules that are continuous, anonymous, and respect unanimity — a problem originally posed by G. Chichilnisky (see [9]). We show that some natural domains of preferences endowed with a measure-based topology allow social choice rules with these properties.

3.2 NOTATION AND DEFINITIONS

We start with some set-theoretical notation: For any set A, CA is the complement of A. $\overline{A}$ and A° is the closure and interior of A respectively. The boundary of A is indicated by ∂A.

In the following, X is a connected topological space of commodities or resources. Let $X \times X$ have the product topology induced by the topology of X. Let μ be a finite and complete measure on $X \times X$, satisfying the condition that every Borel set is measurable.

A *preference on* X is any measurable subset of $X \times X$ (i.e., preferences are identified with their graphs). Be careful to note that this is a very general use of the word. A preference P is said to be *complete* if for all $(x, y) \in X \times X$, $(x, y) \notin P$ implies $(y, x) \in P$ (this should be understood to imply that $(x, x) \in P$ for all x). We say that P is *transitive* if for all $(x, y), (y, z) \in X \times X$, we have that $(x, y) \in P$ and $(y, z) \in P$ implies $(x, z) \in P$. Complete and transitive preferences are also called *complete preorders*. Finally, P is *continuous* if it is closed in $X \times X$.

For the sake of readability, we sometimes use the operators $\succ_P$, $\succsim_P$ and $\sim_P$, defined as follows: $x \succsim_P y$ means $(x, y) \in P$, $x \sim_P y$ means $x \succsim_P y$ and $y \succsim_P x$, and $x \succ_P y$ means $x \succsim_P y$ and not $y \succsim_P x$.

The operator of *symmetric set-difference*, $\triangle$, is defined as

$$A \triangle B = (A - B) \cup (B - A).$$

Let $\mathcal{A}$ be the class of all μ-measurable subsets of $X \times X$. Define a pseudometric d_μ on $\mathcal{A}$ by $d_\mu(P,Q) = \mu(P \triangle Q)$, and let R_μ be the relation satisfying $P \; R_\mu \; Q$ if and only if $d_\mu(P,Q) = 0$. Finally, let $\mathcal{Q}_\mu = \mathcal{A}/R_\mu$ and endow $\mathcal{Q}_\mu$ with the quotient topology.

The subspace of $\mathcal{Q}_\mu$ consisting of equivalence classes that contain at least one complete preorder is called $\mathcal{P}_\mu$, while $\mathcal{P}_\mu^C$ denotes the subspace of equivalence classes that contain at least one continuous and complete preorder.

An *ultrafilter* $\mathfrak{U}$ on a set Y is a collection of subsets of Y satisfying

(i) $\varnothing \notin \mathfrak{U}$, $Y \in \mathfrak{U}$

(ii) if $A, B \in \mathfrak{U}$, then $A \cap B \in \mathfrak{U}$

(iii) if $A \in \mathfrak{U}$ and $A \subset B \subset Y$, then $B \in \mathfrak{U}$

(iv) for all $A \subset Y$, either $A \in \mathfrak{U}$ or $(Y - A) \in \mathfrak{U}$

An ultrafilter on Y is *nonprincipal* if it contains all the cofinite subsets of Y (that is, subsets with a finite complement).

3.3 SOME PROPERTIES OF THE MEASURE-BASED TOPOLOGIES

It is well known that $\mathcal{Q}_\mu$ is separable if X has a countable base for its open sets (in particular, if X is metrizable and separable), see Halmos [**24**, p. 168]. This section will deal with topological completeness of important subsets of $\mathcal{Q}_\mu$. In Le Breton & Uriarte [**31**] the authors seem to argue that completeness and some other mathematical requirements are irrelevant to the evaluation of a topology's economic appeal. It should be apparent, however, that if subsets of $\mathcal{P}_\mu$ with well-defined economic properties also have well-defined topological properties, we have a clear indication that the topology captures some important economic structure of the space.

Theorem 3.1 below shows that $\mathcal{Q}_\mu$ is a topologically complete space, and is well known from measure theory. We still include a proof, partly because it employs a more direct method than those usually seen, and partly because elements of this proof will be used in the proofs of Theorems 3.2, 3.3, and 3.4.

Theorem 3.1 $\mathcal{Q}_\mu$ *is topologically complete.*

PROOF: Let $\{P_n\}$ be a Cauchy-sequence of preferences. Consider the inferior limit of $\{P_n\}$, written P_* and defined as the set of all points that are members of all but a finite number of the sets in P_n. It can also be expressed as

$$P_* = \bigcup_{k=1}^{\infty} \bigcap_{n=k}^{\infty} P_n. \tag{3.1}$$

We show that $\{P_n\}$ converges to P_*.

We may assume that for any positive integer k there is an integer n_k such that

$$d(P_n, P_m) < \frac{1}{2^k} \text{ for } n, m \geq n_k.$$

Let

$$E_k = P_{n_k} \bigtriangleup P_{n_{k+1}}, \tag{3.2}$$

$$F_m = \bigcup_{k=m}^{\infty} E_k. \tag{3.3}$$

Since $\mu(E_k) = d(P_{n_k}, P_{n_{k+1}})$ it follows that

$$\lim_{m\to\infty} \mu(F_m) = 0. \tag{3.4}$$

From (3.2) and (3.3) we get

$$F_m = \bigcup_{n=m}^{\infty} P_n - \bigcap_{n=m}^{\infty} P_n. \tag{3.5}$$

Consider the expression

$$\left(P_m \cup \bigcup_{k=m}^{\infty} \bigcap_{n=k}^{\infty} P_n\right) - \left(P_m \cap \bigcup_{k=m}^{\infty} \bigcap_{n=k}^{\infty} P_n\right). \tag{3.6}$$

This expression is equivalent to $P_m \bigtriangleup P_*$, since we can replace the lower index of the union operator in (3.1) with any integer without changing the limit.

A comparison of (3.5) and (3.6) reveals that $P_m \bigtriangleup P_* \subset F_m$, and (3.4) then implies $\lim_{m\to\infty} \mu(P_m \bigtriangleup P_*) = 0$. ∎

The two next theorems show that the space of transitive preferences and the space of complete preferences are both topologically complete.

Theorem 3.2 *The set of equivalence classes containing transitive preferences is closed in* Ω_μ.

PROOF: The inferior limit of a set of transitive preferences is obviously transitive. ∎

Theorem 3.3 *The set of equivalence classes containing complete preferences is closed in Ω_μ.*

PROOF: Let $\{P_n\}$ be a Cauchy-sequence of preferences. The superior limit of $\{P_n\}$, written P^*, is the set of all points that are members of an infinite number of the sets in $\{P_n\}$. The superior limit can be expressed as

$$P^* = \bigcap_{k=1}^{\infty} \bigcup_{n=k}^{\infty} P_n. \tag{3.7}$$

We first show that $d_\mu(P^*, P_*) = 0$. Since $P_* \subseteq P^*$, we have

$$P^* \bigtriangleup P_* = \bigcap_{k=m}^{\infty} \bigcup_{n=k}^{\infty} P_n - \bigcup_{k=m}^{\infty} \bigcap_{n=k}^{\infty} P_n \quad \text{for any } m. \tag{3.8}$$

It is easy to see that the right hand side of (3.8) is a subset of F_m (see equation (3.5)) for all m. Equation (3.4) then implies $d_\mu(P^*, P_*) = 0$.

We have thus shown that $\{P_n\}$ also converges to its superior limit, and from the observation that the superior limit of a sequence of complete preferences is complete, the theorem immediately follows. ∎

The last two results imply that the set of equivalence classes containing at least one transitive and at least one complete preference is closed, but they say nothing about the set of equivalence classes containing preferences that are both transitive and complete. But as the next theorem shows, this set turns out to be closed as well. To prove this, we apply a construction called an *ultralimit* of a sequence of sets, defined as follows: Let $\mathfrak{U}$ be a nonprincipal ultrafilter over the natural numbers. P^U is an ultralimit of the sequence $\{P_n\}$ whenever we have $x \in P^U$ if and only if $\{ n : x \in P_n \} \in \mathfrak{U}$.

Theorem 3.4 *$\mathcal{P}_\mu$ is topologically complete.*

PROOF: We will show that a Cauchy-sequence of preferences $\{P_n\}$ converges to an ultralimit P^U of the sequence. It is immediately seen that an ultralimit of a sequence of transitive preferences is transitive, and from the fact that for any ultrafilter $\mathfrak{U}$ on $\mathbb{N}$, $A \cup B = \mathbb{N}$ implies $A \in \mathfrak{U}$ or $B \in \mathfrak{U}$ (see e.g. Eklof [**21**]), we can deduce that an ultralimit of a sequence of complete preferences is complete.

Since the ultrafilter is nonprincipal, every cofinite set of integers is in $\mathfrak{U}$, and every set in $\mathfrak{U}$ is infinite. Hence, for an ultralimit P^U, we have $P_* \subset P^U \subset P^*$. In the proof of Theorem 3.3 we showed that $d_\mu(P^*, P_*) = 0$, and the completeness of μ then ensures that P^U is measurable. ∎

The measure-based topologies identify preferences that differ by a set of zero measure. This is a natural consequence of the interpretation of these topologies: If we attempt to estimate the difference between two preferences by a sampling procedure on differences in the ranking of randomly drawn pairs in $X \times X$, (where the sampling distribution is consistent with the probability measure μ), preferences differing by a set of zero measure are empirically indistinguishable.

Our next result, however, shows that under some general conditions, the measure-based topologies *do* separate all points of an important subset of the space of all preferences. It says that if X is a connected space and the generating measure μ assigns positive measure to open sets, then d_μ is a proper metric when restricted to the class of continuous, complete and transitive preferences.

To prove this result, we use the following lemma:

Lemma 3.5 *Let P be a continuous, complete and transitive preference. If $a \sim_P b$, but $(a,b) \notin \overline{P^o}$, there exist (not necessarily distinct) points z and v so that $(a,z) \in \overline{P^o}$, $(v,b) \in \overline{P^o}$, $(z,v) \in \overline{P^o}$ and $(v,z) \in \overline{P^o}$.*

PROOF: Since there is only one preference involved in this proof, we omit the subscript P on the relations. Let $A = \{\, x \in X : x \prec a \,\}$, $B = \{\, x \in X : x \succ b \,\}$. A and B are open and disjoint, and cannot both be empty as this would imply $P = X \times X$ with every point an interior point, contrary to the assumption of the lemma. Assume first that A and B are both nonempty. Then $\partial \overline{A} \neq \varnothing$. If it was empty, $\overline{A}$ would be a both open and closed nonempty set with a nonempty complement, which is impossible since $X \times X$ is a connected space. By the same argument, $\partial B \neq \varnothing$.

Pick a point $z \in \partial A$, then $z \notin A$; hence by completeness $z \succsim a$. But, by continuity we cannot have $z \succ a$, so $z \sim a$. Now every neighborhood of z (in X) contains a point $u \prec a$, hence $(a, u) \in P^o$, which implies that $(a, z) \in \overline{P^o}$. By a symmetric argument there is a point $v \in \partial B$ such that $(v, b) \in \overline{P^o}$, $v \sim b$. Observe that by transitivity $z \sim v$.

Every neighborhood of z has common points with both A and $C\overline{A}$, and every neighborhood of v has common points with both B and $C\overline{B}$. From this it easily follows that $(v, z) \in \overline{P^o}$, and the following argument shows that $(z, v) \in \overline{P^o}$: By completeness every point r in $C\overline{A}$ has a neighborhood N_r so that $x \succsim z$ for all $x \in N_r$, and every point s in $C\overline{B}$ has a neighborhood N_s so that $v \succsim y$ for all $y \in N_s$. By transitivity $x \succsim y$ for all $(x, y) \in N_r \times N_s$. But then every neighborhood of (z, v) contains an interior point (r, s) of P.

We are left with the case where one of A and B is empty. Without loss of generality, assume $A = \varnothing$. Then a must have a neighborhood N_a where $x \sim a$ for all $x \in N_a$ (if this is not the case, i.e. every neighborhood of a has points strictly preferred to a, then $(a, b) \in \overline{P^o}$, contrary to the assumption of the lemma). Let $z = a$; then certainly $(a, z) \in \overline{P^o}$. By assumption $a \in C\overline{B}$, hence $C\overline{B} \neq \varnothing$, and we can apply the same argument as before to show that $\partial B \neq \varnothing$. Choose a point $v \in \partial B$; then, as before, $(v, b) \in \overline{P^o}$. Observing that $x \sim y$ if $x \in C\overline{B}$ and $y \in C\overline{B}$, it is a trivial exercise to show that $(z, v) \in \overline{P^o}$ and $(v, z) \in \overline{P^o}$. ∎

Theorem 3.6 *If X is a connected space, and μ assigns positive measure to open sets, then d_μ is a proper metric when restricted to continuous, complete and transitive preferences.*

PROOF: We show that for any two continuous, complete and transitive preferences P and Q, if $P \neq Q$, then $P \bigtriangleup Q$ has a nonempty interior. From this, the theorem immediately follows.

By assumption, $P - Q$ and $Q - P$ cannot both be empty. Suppose without loss of generality that $P - Q \neq \varnothing$. This means that there exists a point $(a, b) \in X \times X$ so that $a \succsim_P b$ and $a \prec_Q b$. The proof is in two parts. We first suppose that

(i) $(a,b) \in \overline{P^o}$. Every neighborhood $N_{(a,b)}$ of (a,b) contains a point with a neighborhood $O \subseteq P$. By continuity of Q we can choose $N_{(a,b)}$ so that $x \prec_Q y$ for all $(x,y) \in N_{(a,b)}$. Hence, the open and nonempty set $(N_{(a,b)} \cap O)$ is a subset of $(P - Q)$, which again is a subset of $(P \bigtriangleup Q)$.

We are left with the case where

(ii) $(a,b) \notin \overline{P^o}$. Since $(a,b) \notin P^o$, we must have $a \sim b$. By Lemma 3.5 there exist points z and v such that $(a,z) \in \overline{P^o}$, $(z,v) \in \overline{P^o}$ and $(v,b) \in \overline{P^o}$. Since $a \prec_Q b$, by transitivity one of the following must hold: $a \prec_Q z$, $z \prec_Q v$, $v \prec_Q b$. Applying part (i) of this proof on the pair where the condition holds gives the desired result. ∎

As mentioned in the introduction, the need for a topology is primarily motivated by the possibility of observational errors. In [28], Kannai defines a topology on a space $\mathcal{P}$ of continuous and complete preorders as the smallest topology in which the set

$$\{(x,y,P) : x \succ_P y\}$$

is open in the product topology of $X \times X \times \mathcal{P}$. The appeal of Kannai's criterion is that whenever $x \succ_P y$, this relation is also true for all x', y', and P' that are sufficiently close to x, y, and P; in other words, whenever one alternative is strictly preferred to another, this will hold even under small observational errors.

In general, the spaces $\mathcal{P}^C_\mu$ do not have the Kannai property. An example will illustrate this. Let a sequence of preferences $\{P_n\}$ be generated by

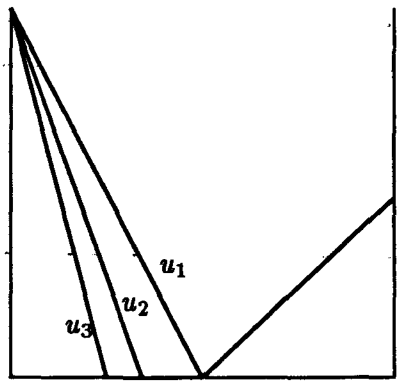

Figure 3.1.

a sequence of utility functions $\{u_n\}$ on $[0,1]$ defined as follows:

$$u_n(x) = \begin{cases} 1-(n+1)x & \text{for } 0 \leq x < 1/(n+1), \\ 0 & \text{for } 1/(n+1) \leq x < 1/2, \\ x - 1/2 & \text{for } 1/2 \leq x \leq 1. \end{cases}$$

Figure 3.1 shows the first three elements of the sequence. If the measure-based topology is generated by for instance the Lebesgue measure, then this sequence converges to the preference P generated by the utility function u defined as

$$u(x) = \begin{cases} 0 & \text{for } 0 \leq x < 1/2, \\ x - 1/2 & \text{for } 1/2 \leq x \leq 1. \end{cases}$$

Clearly, for all P_n, we have $0 \succ_{P_n} 1$, thus, all neighborhoods of P contain preferences where $0 \succ 1$. But for P, we have $1 \succ_P 0$. This violates the Kannai property.

On the other hand, $\mu(\{(x,y) : x \precsim_{P_n} y \ \&\ x \succ_P y\})$ tends to zero as P_n tends to P. If μ is interpreted as a probability measure, this means that we can assume with a high degree of certainty that small observational errors will not affect the outcome. Even if this differs from the Kannai property, it still captures the same idea.

3.4 AN AXIOMATIC CHARACTERIZATION

In the literature we find many alternative topologies for spaces of preferences, all more or less plausible, as is evident from Chapter 2. In most areas of economics, it is common in such cases to resort to what may be labeled "the axiomatic approach". We shall try to put down a few well founded criteria that describe properties we would like a topology to have, and then determine which topologies satisfy these criteria.

For reasons explained in Chapter 2, p. 27, we will aim for a topology that can be generated by a pseudometric d, since we want a metrizable topology if possible. Of course, one of the reasons that so many different topologies have been proposed, is the diffuse intuition one generally has when it comes to "obvious" properties of convergence of preferences. We shall try to circumvent the problem of this lack of intuition by letting as many as possible of the conditions we put down describe properties that do not refer to convergence at all.

Specifically, we list below three conditions[3] we would like a pseudometric to satisfy. Only the first condition makes any mention of convergence, and the convergence concept involved is so weak that, hopefully,

[3] There is a close relationship between Conditions 3.2 and 3.3 in this chapter, and two of the axioms used by Kemeny & Snell [**29**].

it is immediately acceptable to the reader. We say that a sequence of preferences $\{P_n\}$ is *monotonic decreasing* if $P_n \subset P_m$ for $n > m$.

Condition 3.1 *If $\{P_n\}$ is a monotonic decreasing sequence of preferences, then $\lim_{n\to\infty} d(P_n, P) = 0$, where $P = \cap_{n=1}^{\infty} P_n$.*

This is usually called *convergence from above*, and is a standard concept in set and measure theory.

The remaining two conditions does not mention convergence at all, but pertain to "global" properties of the pseudometric. First, we want the distance between two preferences to be a measure of the extent to which the preferences disagree. We also want such a measure to be as specific as possible, in the sense that it should not be influenced by anything else than this disagreement; in particular, if two preferences *agree* on the ranking of a pair of alternatives, it should not matter what ranking they agree on as far as the distance is concerned. The plausibility of this argument can further be seen by taking it to the extreme: If two preferences agree on everything, they are identical, and we want the distance between them to be zero no matter how they rank the alternatives.

We will formalize this by saying that two preferences P and Q *disagree on an ordered pair* (x, y) whenever $x \succsim_P y$ and $x \not\succsim_Q y$, or $x \not\succsim_P y$ and $x \succsim_Q y$.

Condition 3.2 *If two preferences P and Q disagree on exactly the same ordered pairs as two preferences R and S do, and furthermore, when restricted to these pairs, P is equal to R and Q is equal to S, then $d(P, Q) = d(R, S)$.*

The next condition relates to the notion of Pareto optimality. We say that a preference P_3 is Pareto efficient relative to two preferences P_1 and P_2 whenever $x \succsim_3 y$ if $x \succsim_1 y$ and $x \succsim_2 y$, and $x \not\succsim_3 y$ if $x \not\succsim_1 y$ and $x \not\succsim_2 y$, for all $x, y \in X$. We want a condition that ensures that if a preference P_3 is Pareto efficient relative to P_1 and P_2, it should not be possible to find another preference that is closer to *both* P_1 and P_2 than P_3 is. This seems to capture an important intuitive property of "closeness." The condition below implies that the pseudometric will have this property.

Condition 3.3 *If a preference P_3 is Pareto efficient relative to two preferences P_1 and P_2 then*

$$d(P_1, P_2) = d(P_1, P_3) + d(P_3, P_2).$$

In other words, under Pareto efficiency the triangle inequality should hold as an equality.

Theorem 3.7 *A pseudometric on a σ-algebra of preferences satisfies Conditions 3.1–3.3 if and only if it is a measure-based pseudometric.*

PROOF: We first show that if a pseudometric satisfies the conditions, it must be measure-based. The argument will be easier to follow if we first translate Conditions 3.2 and 3.3 into statements involving set operations. Condition 3.2 says that for any four sets P, Q, R, and S, if $P - Q = R - S$ and $Q - P = S - R$, then $d(P, Q) = d(R, S)$. Condition 3.3 says that if $P_1 \cap P_2 \subset P_3 \subset P_1 \cup P_2$, then $d(P_1, P_2) = d(P_1, P_3) + d(P_3, P_2)$.

Define a set function μ by

$$\mu(P) = d(P, \varnothing).$$

We prove that μ is a measure. First, we show that μ is finitely additive: Let $\{A_i\}$, $i = 1, \ldots, n$, be any collection of n mutually disjoint sets. Define E_m by

$$E_m = \bigcup_{i=1}^{m} A_i.$$

By Condition 3.3 we must have

$$d(E_m, \varnothing) = d(E_m, E_{m-1}) + d(E_{m-1}, \varnothing) \quad \text{for } 1 < m \leq n. \tag{3.9}$$

But by Condition 3.2 we have $d(E_m, E_{m-1}) = d(A_m, \varnothing)$, so (3.9) can be written

$$d(E_m, \varnothing) = d(A_m, \varnothing) + d(E_{m-1}, \varnothing) \quad \text{for } 1 < m \leq n. \tag{3.10}$$

Repeated applications of (3.10) then give

$$d(E_n, \varnothing) = \sum_{i=1}^{n} d(A_i, \varnothing),$$

and this shows that μ is finitely additive.

Clearly, since d has the properties of a pseudometric, μ is finite and nonnegative. Condition 3.1 implies that μ is continuous from above at $\varnothing$. Since μ also is finitely additive, we may conclude that μ is a finite measure (see Halmos [**24**, p. 39]). It is left to show that this measure actually generates the metric, i.e., that $d(P, Q) = \mu(P \triangle Q)$.

By Condition 3.3 we must have, for any two sets P and Q,

$$d(P,Q) = d(P, P \cap Q) + d(P \cap Q, Q). \tag{3.11}$$

But by Condition 3.2,

$$d(P, P \cap Q) = d(P - Q, \varnothing)$$

and

$$d(P \cap Q, Q) = d(\varnothing, Q - P).$$

We can then write (3.11) as

$$\begin{aligned} d(P,Q) &= d(P - Q, \varnothing) + d(Q - P, \varnothing) \\ &= \mu(P - Q) + \mu(Q - P) \\ &= \mu(P \triangle Q), \end{aligned}$$

and the first part of the proof is completed.

It remains to show that any measure-based pseudometric satisfies all three conditions. Condition 3.1 is satisfied because any measure is continuous from above. Condition 3.2 is satisfied since $P - Q = R - S$ and $Q - P = S - R$ implies $P \triangle Q = R \triangle S$. For any tree sets P_1, P_2, and P_3, if $P_1 \cap P_2 \subset P_3 \subset P_1 \cup P_2$, then $P_1 \triangle P_3$ and $P_3 \triangle P_2$ are disjoint sets and $P_1 \triangle P_2 = (P_1 \triangle P_3) \cup (P_3 \triangle P_2)$. The additivity of the measure then implies $d(P_1, P_2) = d(P_1, P_3) + d(P_3, P_2)$, so Condition 3.3 is satisfied. This completes the proof. ∎

3.5 AGGREGATION OF PREFERENCES

The so-called "social choice paradoxes" show that seemingly week properties of social aggregation rules may be mutually inconsistent. The best known paradox where continuity of the aggregation rule is involved is due to Chichilnisky [9]. In her paper, the following three properties of an aggregation rule $\phi: \mathcal{P}^n \rightarrow \mathcal{P}$ (where $\mathcal{P}^n$ is the n-fold cartesian product of the preference space $\mathcal{P}$, $n \geq 2$) are considered:

(i) ϕ is continuous

(ii) ϕ respects unanimity, i.e. $\phi(P, \ldots, P) = P$

(iii) ϕ is anonymous, i.e. $\phi(P_1, \ldots, P_n) = \phi(P_{\sigma(1)}, \ldots, P_{\sigma(n)})$ where σ is any permutation on $\{1, \ldots, n\}$

In the mathematical literature, a map that satisfies (i)–(iii) is called an n-mean.

In Chichilnisky & Heal [14] it is shown that for CW-complexes with a convex hull that is also a CW-complex, a sufficient condition for the consistency of (i)–(iii) is that all connected components of the space $\mathcal{P}$ are contractible. For parafinite CW-complexes this condition is also necessary.

The requirement that $\mathcal{P}$ should be a CW-complex is crucial. Consider a space Q constructed in the following way: For each positive integer n, let C_n be the circle in $\mathbb{R}^2$ with center at $(1/n, 0)$ and radius $1/n$. Let $A = \bigcup_{n \geq 1} C_n$ and let $Q = (A \times [0,1])/(A \times \{1\})$, i.e., Q is the cone on A. It is clear that Q is compact and contractible, but in [3], P. Bacon shows that there is no map $\phi: Q^n \rightarrow Q$ that satisfy (i)–(iii).

Since there is no natural way to identify CW-complexes with spaces of preferences endowed with a measure-based topology, the results of Chichilnisky & Heal [14] cannot be transformed to our framework. Even though it is easy to show that some preference spaces (for instance, the space $\mathcal{P}_\mu$ with μ absolutely continuous with respect to the Lebesgue measure) are contractible, this does not immediately imply the existence of an aggregation rule that is an n-mean.

However, if we only consider the subset of $\mathcal{P}_\mu$ consisting of preferences without "thick" indifference surfaces, i.e., the preferences P where $\mu(\{ (x,y) \in P : x \sim_P y \}) = 0$, it is actually possible to construct an aggregation rule with all the necessary properties. We shall also see that this rule aggregates continuous preferences to a continuous preference, so the subset of $\mathcal{P}_\mu$ consisting of continuous preferences without thick indifference surfaces will also admit an aggregation that satisfy (i)–(iii). At present, these are among the widest possibility results known.

We will now assume that the measure μ on $X \times X$ is a product measure generated by a finite measure ν on X. It is natural to investigate the connection between $\mathcal{P}_{\nu\times\nu}$ and the space of real-valued ν-integrable functions on X endowed with the pseudometric d'_ν defined as

$$d'_\nu(f,g) = \int_X |f-g|\, d\nu.$$

As usual, functions with zero distance between them are identified, and the resulting quotient space is called $L^1(X,\nu)$.

For any preference P in $\mathcal{P}_{\nu\times\nu}$, the function

$$f_P(x) = \nu(\{ y : y \precsim_P x \}) \tag{3.12}$$

exists[4], and is easily seen to almost represent P in the sense that $x \succsim_P y$ is almost everywhere equivalent to $f_P(x) \geq f_P(y)$. We shall see that this transformation induces a map U from $\mathcal{P}_{\nu\times\nu}$ to $L^1(X,\nu)$, i.e., the transformation generated by (3.12) takes equivalent preferences to equivalent functions. Furthermore, U is a continuous map.

Lemma 3.8 *The map $U\colon \mathcal{P}_{\nu\times\nu} \to L^1(X,\nu)$ as described above is well defined and continuous.*

PROOF: We show that $d_{\nu\times\nu}(P,Q) \geq d'_\nu(f_P, f_Q)$, so in particular, $d_{\nu\times\nu}(P,Q) = 0$ implies $d'_\nu(f_P, f_Q) = 0$, and it follows that U is well defined in the sense that it maps equivalence classes to equivalence classes.

First consider the inequality

$$\nu(A \triangle B) \geq |\nu(A) - \nu(B)| \tag{3.13}$$

for any measure ν and measurable sets A,B. To see that this inequality must hold, consider that $\nu(A \triangle B)$ can be written as

$$\nu(A - B) + \nu(B - A).$$

It is also trivial that

$$\nu(A - B) \geq \nu(A) - \nu(B),$$

[4]For a somewhat related construction of utility functions, see Mount & Reiter [34].

and

$$\nu(B - A) \geq \nu(B) - \nu(A).$$

The inequality (3.13) follows immediately.

We now consider the expression $(\nu \times \nu)(P \bigtriangleup Q)$. Let a subscript x on a subset of $X \times X$ indicate a section of the subset (with the first coordinate fixed at x), i.e., for a set $A \subset X \times X$, the expression A_x is the set $\{ y : (x, y) \in A \}$. By Fubini's theorem we have

$$(\nu \times \nu)(P \bigtriangleup Q) = \int_X \nu((P \bigtriangleup Q)_x)\, d\nu(x). \tag{3.14}$$

But $(P \bigtriangleup Q)_x$ can obviously be written as $P_x \bigtriangleup Q_x$, so (3.14), together with (3.13), gives us

$$(\nu \times \nu)(P \bigtriangleup Q) \geq \int_X |\nu(P_x) - \nu(Q_x)|\, d\nu(x). \tag{3.15}$$

Since the right hand side is just $\int_X |f_P - f_Q|\, d\nu$, we have shown that $d_{\nu \times \nu}(P, Q) \geq d'_\nu(f_P, f_Q)$, which completes the proof. ∎

Consider now the map Γ that maps functions in $L^1(X, \nu)$ to preferences according to the rule

$$(x, y) \in \Gamma(f) \text{ if and only if } f(x) \geq f(y).$$

It is clear that Γ is well defined, since equivalent functions differ on a set of zero measure, and thus (by Fubini's theorem) generate equivalent preferences.

In general, Γ is not continuous. However, it becomes continuous when restricted to the space $K^1(X,\nu)$ of all functions f in $L^1(X,\nu)$ that satisfy

$$(\nu \times \nu)(\{\,(x,y) : f(x) = f(y)\,\}) = 0,$$

as the following lemma shows.

Lemma 3.9 *The map $\Gamma\colon K^1(X,\nu) \to \mathcal{P}_{\nu\times\nu}$ is continuous.*

Proof: Assume that f_n converges to f in $K^1(X,\nu)$. Then, for all $\epsilon > 0$,

$$\begin{aligned}\Gamma(f) \mathbin{\triangle} \Gamma(f_n) \subset \{(x,y) :\ & |f(x) - f_n(x)| \geq \epsilon \text{ or} \\ & |f(y) - f_n(y)| \geq \epsilon \text{ or} \\ & |f(x) - f(y)| < 2\epsilon\}. \end{aligned} \tag{3.16}$$

To see that (3.16) must hold, consider a pair (x', y') that does not satisfy any of the terms in the disjunction. For this pair, we obviously have $f(x') \geq f(y')$ if and only if $f_n(x') \geq f_n(y')$, so the pair cannot be in $\Gamma(f) \mathbin{\triangle} \Gamma(f_n)$.

Without loss of generality, assume that $\mu(X) = 1$. By (3.16), Fubini's theorem, and some elementary logic we have, for all $\epsilon > 0$,

$$\begin{aligned}(\nu \times \nu)(\Gamma(f) \mathbin{\triangle} \Gamma(f_n)) \leq\ & 2\nu(\{\,x : |f(x) - f_n(x)| \geq \epsilon\,\}) + \\ & (\nu \times \nu)(\{\,(x,y) : |f(x) - f(y)| < 2\epsilon\,\}).\end{aligned} \tag{3.17}$$

We want to show that, for any $\delta > 0$, we can find a k such that $(\nu \times \nu)(\Gamma(f) \mathbin{\triangle} \Gamma(f_n)) < \delta$ for all $n \geq k$. Since

$$(\nu \times \nu)(\{\,(x,y) : f(x) = f(y)\,\}) = 0,$$

by continuity from above we can find an ϵ sufficiently small so that

$$(\nu \times \nu)(\{\,(x, y) : |f(x) - f(y)| < 2\epsilon\,\}) < \delta/2.$$

Convergence in $L^1(X, \nu)$ implies convergence in measure, so for this ϵ, we can find a k sufficiently large to make

$$\nu(\{\,x : |f(x) - f_n(x)| \geq \epsilon\,\}) < \delta/4$$

hold for all $n \geq k$. But then the right hand side of (3.17) is less than δ for all n greater than or equal to this k, and the proof is complete. ∎

Let $\mathcal{P}^K_{\nu\times\nu} \subset \mathcal{P}_{\nu\times\nu}$ be the space of preferences P that satisfy

$$(\nu \times \nu)(\{\,(x, y) : x \sim_P y\,\}) = 0.$$

Theorem 3.10 *The space $\mathcal{P}^K_{\nu\times\nu}$ admits an n-mean.*

PROOF: Define the map $F\colon [K^1(X, \nu)]^n \to K^1(X, \nu)$ by

$$[F(f_1, \ldots, f_n)](x) = \min\{f_1(x), \ldots, f_n(x)\} \text{ for all } x.$$

This map is obviously an n-mean on $K^1(X, \nu)$. Consider the map ϕ defined by the diagram

$$\begin{tikzcd}
(\mathcal{P}^K_{\nu\times\nu})^n \arrow[r, "\phi"] \arrow[d, "U\times\cdots\times U"'] & \mathcal{P}^K_{\nu\times\nu} \\
{[K^1(X,\nu)]^n} \arrow[r, "F"] & K^1(X,\nu) \arrow[u, "\Gamma"']
\end{tikzcd}$$

Since $\Gamma \circ U$ is the identity map on $\mathcal{P}^K_{\nu\times\nu}$, and all involved maps are continuous, the composite map ϕ must be an n-mean. ∎

With some further restrictions on μ, we also get the following result:

Corollary 3.11 *If X is a subset of $\mathbb{R}^n$ and ν is finite and absolutely continuous with respect to the Lebesgue measure, the subspace of $\mathcal{P}^K_{\nu\times\nu}$ that consists of equivalence classes that contain continuous and complete preorders allows an n-mean.*

PROOF: A result in Neuefeind [**35**, p. 174] implies (with some trivial modifications of Neuefeind's proof) that U maps continuous preferences in $\mathcal{P}^K_{\nu\times\nu}$ to continuous functions. Since both F and Γ maps continuous objects to continuous objects, this also holds for the map ϕ constructed above. ∎

4 Strategy-proofness: A homotopy approach

4.1 INTRODUCTION

In this chapter we consider strategy-proofness of continuous social welfare functions and other aggregation procedures. Similar problems have been examined by several authors, starting with the seminal papers by Gibbard [23] and Satterthwaite [38], and surveyed by Pely [37], Pattanaik [36] and Sen [39]. In most of the literature on this subject, the social aggregation procedures under investigation have been social choice functions or social choice correspondences. It is assumed that all voters have preferences over a choice space X. A *social choice correspondence* is then a relation between profiles of preferences (i.e., n-tuples of individual preferences if there are n voters) and points in X. The interpretation is then that a point in X is a possible outcome of the social decision process for the given preference profile whenever the relation holds. When only *one* point in X corresponds to each preference profile, we get the special case of a *social choice function.*

In this setting, it is assumed that the result of the social decision process is a point in X. A *social welfare function*, on the other hand, assumes that society should choose a preference over X, rather than a point in X. It is frequently maintained that while the inherent rationality that is embedded in the concept of a preference may be a reasonable assumption for individual behavior, there are few arguments to support the claim that this rationality should also apply to collective decisions. However, nearly all of economic theory is founded on the assumption that the economic agents, or decision units, display sufficient rationality for their choices to be modeled by preferences. In real world applications, one will find a number of cases where the economic agent in question is not an individual, but a group of individuals, and with a choice behavior resulting from some kind of collective decision process. If these decisions cannot be (more or less closely) modeled by preferences, major parts of economic theory will simply not apply. For this reason alone, the investigation of social welfare functions is an important one.

While the literature on manipulation of social decisions concentrates on social choice functions, the possibility of strategic voting should of course be taken just as seriously when social welfare functions are involved. Loosely speaking, a social welfare function can be manipulated if any of the voters can benefit from insincere disclosure of their preferences.

The results we present here are closely related to two articles by Chichilnisky [**13**] and Chichilnisky & Heal [**15**]. Their work on this and related issues differs from most of the literature by the requirement that the social welfare function is continuous.

This, of course, presupposes that a topology is defined on the space of preferences. If we (for the moment) confine the study to linear preferences over $\mathbb{R}^n$ (i.e., preferences that can be represented by linear utility functions), the orthogonal vectors to the indifference surfaces will have the same direction everywhere. The space of linear preferences can then be given the same topology as S^{n-1} (the $n-1$ dimensional sphere) by

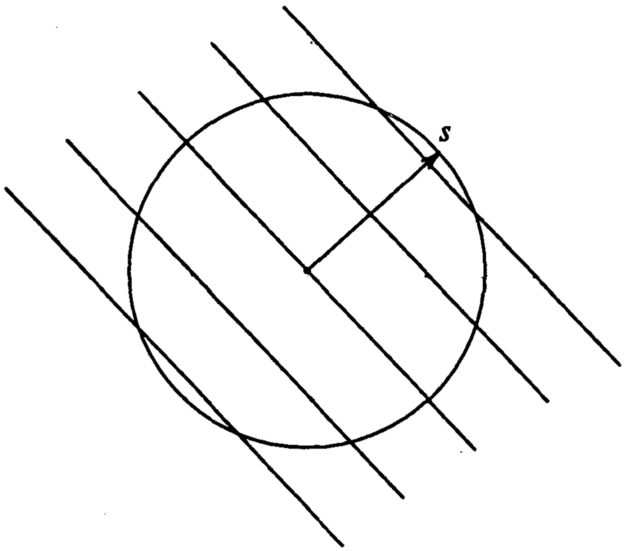

Figure 4.1: This construction maps a linear preference to a point s in S^1

identifying any preference with the point where the orthogonal vector at the origin of $\mathbb{R}^n$ intersects an $n-1$ dimensional sphere centered at the origin of $\mathbb{R}^n$. The two-dimensional case is illustrated in Figure 4.1.

With linear preferences over $\mathbb{R}^n$ thus regarded as points in S^{n-1}, a continuous social welfare function defined for m voters becomes a map from the product of m copies of S^{n-1} to the space of social preferences, which we assume is identical to the space of individual preferences, i.e. S^{n-1}.

For instance, if we assume that the social welfare function is defined for two voters with preferences over $\mathbb{R}^2$, the welfare function can be regarded as a map from the torus to the circle. Call this map $f: S^1 \times S^1 \to S^1$. It can be shown that the degree of the restriction of f to the diagonal of $S^1 \times S^1$ is equal to the sum of the degrees of the restrictions of f to the two subspaces $x \times S^1$ and $S^1 \times x$, respectively, where x is an arbitrary point in S^1. If the restriction of f to the diagonal is of nonzero degree (e.g., if f respects unanimity), then at least one of the last two restrictions of f must be of nonzero degree. From the fact that any map of nonzero degree is onto, it is easily seen that at least one of the two agents can always achieve any social outcome he may desire, given the preference the other agent reports (in general, of course, he will have to misrepresent his preferences to accomplish this). In the terminology of Chichilnisky & Heal [15], such an agent is called a strategic dictator.

This argument is easily generalized to spheres of arbitrary dimension, and an arbitrary number of voters, and the result then becomes:

Theorem 4.1 (Chichilnisky & Heal) *If a continuous social welfare function $f: (S^n)^m \to S^n$ (where $n \geq 1$ and $m \geq 2$) is of nonzero degree when restricted to the diagonal, there is always a strategic dictator.*

The proof, using degree arguments, depends on the spherical structure of the preference space. However, one can easily imagine spaces of preferences that are not homeomorphic to an n-sphere. In several papers, e.g. Chichilnisky & Heal [14] and Heal [26], it has been shown that *noncontractibility* of the preference space gives rise to social choice paradoxes. It is then natural to ask whether there is a connection between these paradoxes and the result above; in other words, can Theorem 4.1 (or some version of it) be extended to noncontractible spaces in general? The rest of this chapter is devoted to that question.

4.2 TWO COUNTEREXAMPLES

So far, we have been referring to spaces of preferences. However, since the mathematical results presented here can be applied to other problems than aggregation of preferences (for some examples, see Chichilnisky & Heal [**15**]), we will henceforth consider just an abstract topological space Y. We furthermore assume there is defined an *aggregation map* $f: Y \times Y \to Y$, to be interpreted in the following way: Two agents each report a point in Y (their "votes"), and the aggregation map then selects a social outcome in Y.

An example from game theory can illustrate the generality of this approach. Consider a game with two agents and a state space Y. The two agents act independently. It is proposed that the agents either remain in the initial state y_0, or move a positive distance in direction x. Let V be the space of all feasible directions. The two agents have different preferences over the states in the system, and will bargain about a common solution. The bargaining process can thus be modeled as a map $f: Y \times Y \to Y$, where Y is the space $V \cup \{y_0\}$. It is a consequence of the results in this chapter that for a fairly general class of spaces Y, even though the game is perfectly symmetric in the sense that the two agents are treated equally, the bargaining outcome will necessarily be of a very asymmetric nature.

In order to keep notation and arguments as simple as possible, we will not consider cases with more than two agents, even though there are some generalizations to an arbitrary number.

Let both agents have preferences[1] over the points of Y. It will be assumed that these preferences have unique global maxima. If we regard the agents' preferences over Y as preferences over social outcomes, the aggregation map gives rise to a two-person noncooperative game, where the possible moves for both players are the points in Y, and with outcome $f(x_1, x_2)$ if the players' moves are x_1 and x_2. We say that the aggregation map is *strategy-proof* for a given pair of preferences over Y if it is a Nash equilibrium in this game that both players report their most preferred point (the unique maximum of their preferences).

In general, an aggregation map can clearly be strategy-proof for some pairs of preferences, but not for other pairs. By an *exhaustive class of preferences over* Y we mean a collection of preferences with unique maxima, such that every point in Y is the maximum of some preference in the collection. Given an exhaustive class of preferences, we say that f is strategy-proof for this class if it is strategy-proof for any pair of preferences where both preferences are members of the class. In the following, when we say that a map is strategy-proof without referring to a particular class, we mean that the map is strategy-proof for at least *some* exhaustive class of preferences.

We say that an aggregation map is *dictatorial* (with agent m as a dictator) if

$$f(x_1, \ldots, x_m, \ldots, x_n) = x_m \quad \text{for all } x_1, \ldots, x_n.$$

[1]Note that the preferences we consider here have no connection with the preferences we described in Section 4.1. There it was assumed that the result of the social decision process (the social outcome) was a social preference; the preferences we introduce here are individuals' preferences over social outcomes, whether these outcomes are social preferences, or something else. If Y is regarded as a space of social preferences, the preferences we describe in this section should then be interpreted as preferences over social preferences.

It is an easy consequence of Theorem 4.1 that given an arbitrary exhaustive class of preferences over the points in S^n, the only aggregation maps that may be strategy-proof for this class are the dictatorial ones (assuming the map is of nonzero degree when restricted to the diagonal).

There is one requirement in the formulation of Theorem 4.1 that cannot immediately be generalized to other spaces than S^n, and that is the nonzero degree condition. We will replace this with a stronger criterion; we require that the aggregation map shall respect unanimity, i.e., $f(x,\ldots,x) = x$ for all x. For the spaces S^n, this implies that the restriction of f to the diagonal is of degree 1. Theorem 4.1 then implies:

Theorem 4.1′ *Let Y be homeomorphic to S^n for some $n \geq 1$. For an arbitrary exhaustive class of preferences over Y, if an aggregation map respects unanimity and is strategy-proof, it is dictatorial.*

We will try to extend this weaker version of Theorem 4.1 to a wider class of spaces than the n-dimensional spheres. However, a simple counterexample shows that any attempt to extend it to all noncontractible spaces will fail. Let Y be homeomorphic to a cylinder, and define an exhaustive class of preferences over Y as follows: Let $d: Y \times Y \to \mathbb{R}$ be the natural metric on Y, i.e., the distance between two points is defined to be the Euclidean length of the shortest path in Y that connects the points. To any point $x \in Y$, we associate a preference over Y that has x as its unique global maximum, and where $y \succsim z$ if $d(x,y) \leq d(x,z)$.

The cylinder can be written as $S^1 \times I$, where I is the unit interval. Denote points in S^1 by σ (with subscripts to distinguish between different

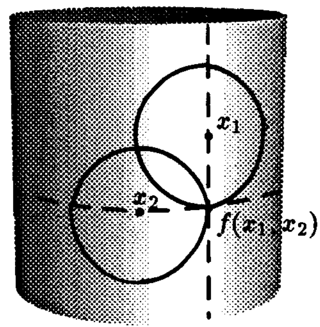

Figure 4.2.

points), and points in I by ι. A point $x_1 \in Y$ can then be written (σ_1, ι_1). Consider the aggregation map $f: Y \times Y \to Y$ defined by

$$f[(\sigma_1, \iota_1), (\sigma_2, \iota_2)] = (\sigma_1, \iota_2).$$

This map is strategy-proof for the exhaustive class of preferences described above, but it is not dictatorial. An example is depicted in Figure 4.2. We assume that agent 1's most preferred point is x_1, and agent 2's most preferred point is x_2. For each of the two corresponding preferences, we have drawn the indifference curve that contains $f(x_1, x_2)$. We see that agent 1 has nothing to gain by reporting anything else than x_1, since he can only achieve outcomes that lie on the horizontal broken line in the figure (assuming that agent 2 reports x_2). The situation is the same for agent 2, since he can only achieve outcomes on the vertical broken line.

We have here a noncontractible space, with a very natural class of preferences and a nondictatorial aggregation map, where neither of the agents have incentives to report anything else than their most preferred points.

Consequently, Theorem 4.1′ does not extend to noncontractible spaces in general. We also observe that the cylinder is homotopy equivalent to the circle, where Theorem 4.1′ is valid, so the validity of the theorem must depend on topological characteristics below the homotopy level. The obvious question is then: Which of the topological properties that are possessed by the circle, and not by the cylinder, are relevant to our problem?

As a first attempt, we note that the proof of Theorem 4.1 relies heavily on the important fact that any map from S^n to S^n of nonzero degree is onto. A suitable generalization of this property to topological spaces in general will be to require that any map from Y to Y that is homotopic to the identity map on Y, is onto. We shall call a space Y with this property *a retracted space*. However, adding the requirement that Y must be a retracted space is still not enough to ensure that only dictatorial maps are strategy-proof for an arbitrary exhaustive class of preferences. We use the torus as a counterexample.

Let Y be homeomorphic to the torus. Y can then be written as $S^1 \times S^1$, and points in Y will be denoted by (σ, σ'). Let $\hat{d}$ be the natural metric on the circle, i.e., the distance between two points is the Euclidean length of the shortest path in S^1 that connects the points. Define a metric d on the torus by

$$d[(\sigma_1, \sigma_1'), (\sigma_2, \sigma_2')] = \sqrt{\hat{d}(\sigma_1, \sigma_2)^2 + \hat{d}(\sigma_1', \sigma_2')^2}.$$

Define an exhaustive class of preferences on Y by exactly the same method as in the previous counterexample. Let the aggregation map be defined by

$$f[(\sigma_1, \sigma_1'), (\sigma_2, \sigma_2')] = (\sigma_1, \sigma_2').$$

The situation is illustrated in Figure 4.3, which is to be interpreted in

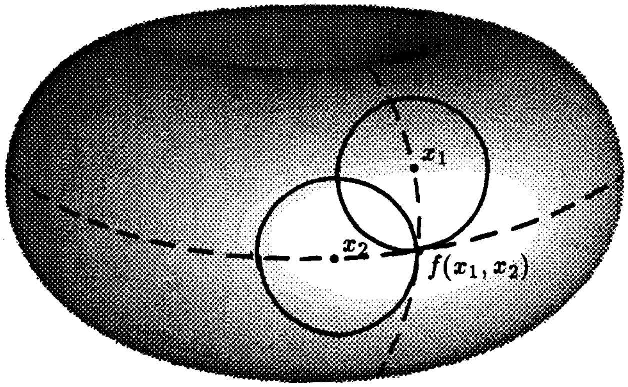

Figure 4.3.

the same way as Figure 4.2.

Since the torus is a retracted space, this example shows that even retracted spaces admit strategy-proof aggregation maps that are not dictatorial. However, if we also require that the space shall be an H'-space (to be defined in the next section), it turns out that we get an interesting generalization of Theorem 4.1′.

4.3 SUFFICIENT CONDITIONS

Before we give the definition of an H'-space, we shall explain a few mathematical terms. A *pointed space* is a topological space where an arbitrary point, called the base point, is singled out. The $\vee$-*union* of two pointed spaces X and Y, written $X \vee Y$, is constructed from the disjoint union of the spaces by identifying the two basepoints. Informally, the two spaces can be thought of as being "glued together" at the base points, and at all other points they retain their original topology. More formally, if X has base point x_0 and Y has base point y_0, $X \vee Y$ may be regarded as the subspace $X \times y_0 \cup x_0 \times Y$ of $X \times Y$. E.g, the figure eight can be regarded as the $\vee$-union of two circles.

A *based map* from a pointed space to a pointed space is a map that takes the base point in the first space to the base point in the second. If X, Y, and Z are pointed spaces, two based maps $f: X \to Z$ and $g: Y \to Z$ can be combined into a map $(f, g): X \vee Y \to Z$ by letting f apply to the $X \times y_0$ part of $X \vee Y$, and g to the $x_0 \times Y$ part.

An H'-space[2] is a pointed space Y together with a map $\mu: Y \to Y \vee Y$ with the property that the composite maps (where c is the constant map that maps everything to the base point, and i is the identity map)

$$Y \xrightarrow{\mu} Y \vee Y \xrightarrow{(c,i)} Y \quad \text{and} \quad Y \xrightarrow{\mu} Y \vee Y \xrightarrow{(i,c)} Y \tag{4.1}$$

are both homotopic to i. Some examples of H'-spaces will be given in Section 4.4.

If Y is an H'-space, we can introduce a binary operator '$*$' on based maps from Y to Y by defining $f * g$ to be the composite map $(f, g) \circ \mu$, or equivalently, with the diagram

$$Y \xrightarrow{\mu} Y \vee Y \xrightarrow{(f,g)} Y.$$

We shall call $f * g$ *the product* of f and g.

[2]See for instance Maunder [33].

Before we state our main result, we shall deduce a general property of strategy-proof aggregation maps. For an aggregation map f and an arbitrary point $x_0 \in Y$, define the two maps $f_1^{x_0}, f_2^{x_0}: Y \to Y$ by

$$f_1^{x_0}(x) = f(x, x_0) \tag{4.2}$$
$$f_2^{x_0}(x) = f(x_0, x). \tag{4.3}$$

If we assume that f is strategy-proof, we must have

$$f(x_1, x_2) = f[f(x_1, x_2), x_2] \quad \text{for all } x_1, x_2. \tag{4.4}$$

This is a consequence of the following argument: Let x_1 and x_2 be two arbitrary points in Y, and let agent 1's and agent 2's most preferred points be $f(x_1, x_2)$ and x_2, respectively. Since f by assumption is strategy-proof, we may suppose that agent 2 reports x_2. By reporting x_1, agent 1 can always achieve $f(x_1, x_2)$ (his most preferred point). But since the map is strategy-proof, the outcome $f(f(x_1, x_2), x_2)$ must be at least as good for him. Considering that $f(x_1, x_2)$ is the unique maximum of his preference, this implies (4.4). By a symmetric argument we have

$$f(x_1, x_2) = f[x_1, f(x_1, x_2)]. \tag{4.5}$$

It is now an immediate consequence of (4.4) and (4.5) that $f_1^{x_0}(x)$ and $f_2^{x_0}(x)$ are both *retractions*; that is, they are equal to the identity map when restricted to their respective images.

We can now prove the main result in this chapter:

Theorem 4.2 *Let Y be a path connected, retracted H'-space. For an arbitrary exhaustive class of preferences over Y, if an aggregation map respects unanimity and is strategy-proof, it is dictatorial.*

PROOF: Let $f: Y \times Y \to Y$ be an aggregation map that respects unanimity. Our aim is to show that if f is strategy-proof, it must be dictatorial.

Pick an arbitrary base point $x_0 \in Y$. Consider the three maps $f_\Delta, f_1^{x_0}, f_2^{x_0}: Y \to Y$, where f_Δ is defined as

$$f_\Delta(x) = f(x, x),$$

and $f_1^{x_0}$ and $f_2^{x_0}$ are defined by (4.2) and (4.3). We first show that f_Δ is homotopic to $f_1^{x_0} * f_2^{x_0}$. Let $u: Y \vee Y \to Y \times Y$ be the inclusion map $[(i, c), (c, i)]$. Call the two composite maps in (4.1) ϕ and ϕ', respectively. By the definition of the product of two functions, $f_1^{x_0} * f_2^{x_0}$ is equivalent to the composite map

$$Y \xrightarrow{\mu} Y \vee Y \xrightarrow{u} Y \times Y \xrightarrow{f} Y, \tag{4.6}$$

that is, $f \circ [(i, c), (c, i)] \circ \mu$. An equivalent way of writing this is

$$f \circ [(i, c) \circ \mu, (c, i) \circ \mu],$$

or shorter, $f \circ (\phi, \phi')$. The following diagram will thus define the same composite map as (4.6):

$$Y \xrightarrow{(\phi, \phi')} Y \times Y \xrightarrow{f} Y$$

This map must be homotopic to f_Δ, since both ϕ and ϕ' are homotopic to i, and if we replace both ϕ and ϕ' in the above diagram with i, the composite is seen to be identical to f_Δ.

It is clear from the definition of the product that for any two based maps $g, h: Y \to Y$, we have $(g * h)(Y) = g(Y) \cup h(Y)$. Since $f_1^{x_0} * f_2^{x_0}$ is homotopic to f_Δ (which is equal to the identity map, since f respects unanimity), it must be onto, because Y is assumed to be a retracted space. We thus have

$$f_1^{x_0}(Y) \cup f_2^{x_0}(Y) = Y.$$

We now show by contradiction that one of $f_1^{x_0}$ and $f_2^{x_0}$ must be a constant map that maps everything into x_0. Assume this is not the case; we can then find a point x_1 different from x_0 so that $x_1 \in f_1^{x_0}(Y)$, which implies $f(x_1, x_0) = x_1$. But then we have $x_0 \notin f_2^{x_1}(Y)$, as the converse would imply $f(x_1, x_0) = x_0$. By the same argument there is a point x_2 different from x_0 so that $x_0 \notin f_1^{x_2}(Y)$. By the properties of Y and the fact that $f_1^{x_0} * f_2^{x_0}$ is homotopic to the identity map, we must have

$$f_1^{x_2}(Y) \cup f_2^{x_1}(Y) = Y.$$

We then arrive at a contradiction, since the above argument implies that x_0 is not a member of this union.

This means that one of $f_1^{x_0}$ and $f_2^{x_0}$ is a constant map. Without loss of generality, assume this is true of $f_2^{x_0}$. The map $f_1^{x_0}$ must then be onto, since $f_1^{x_0}(Y) \cup f_2^{x_0}(Y) = Y$. As $f_1^{x_0}$ is a retraction, and is onto, it is the identity map on Y. For any x, f_1^x must be homotopic to $f_1^{x_0}$, hence onto (since Y is a retracted space), hence f_1^x must be the identity map. But then agent 1 is a dictator. This completes the proof. ∎

Remark: It is not hard to see that a strategy-proof aggregation map that is onto, must also respect unanimity. If the map is onto, then, for any point x, there must exist points x_1 and x_2 such that $f(x_1, x_2) = x$. This means that $f_1^{x_2}(x_1) = x$, and since $f_1^{x_2}$ is a retraction, we must have $f_1^{x_2}(x) = x$, i.e., $f(x, x_2) = x$. Again this means that $f_2^x(x_2) = x$, and $f_2^x(x) = x$, but this is the same as $f(x, x) = x$. The condition of Theorem 4.2 that the aggregation map should respect unanimity can thus be replaced with a condition that only requires the map to be onto.

4.4 CONCLUDING REMARKS

We have seen that a path connected space that is both a retracted space and an H'-space does not admit nondictatorial, strategy-proof aggregation maps that respect unanimity. The counterexamples in Section 4.2 illustrate that neither of the two properties are sufficient by themselves. The space in Figure 4.2 is an H'-space, but it is not retracted, as the identity map is homotopic to a map from the cylinder to a circle. The space in Figure 4.3 is retracted, but it is not an H'-space.

To get an intuitive feeling for the characteristics of H'-spaces, it is instructive to consider an important class of such spaces, the *suspensions*. The suspension of an arbitrary topological space Y is defined to be the quotient space of $Y \times I$ where $Y \times 0$ is identified to one point and $Y \times 1$ is identified to another point. For example, the suspension of a circle is a cylinder with the two ends collapsed into one point each; in other words, a space homeomorphic to a sphere.

The property of being an H'-space is preserved under homotopy equivalence. This means that the figure eight, for instance, is an H'-space, since it is homotopy equivalent to the suspension of a space consisting of three discrete points.

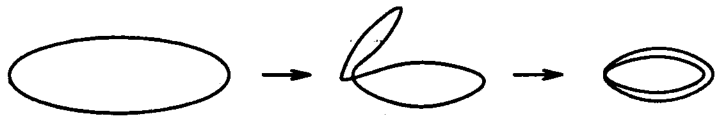

Figure 4.4: Folding a circle around itself.

Informally, retracted H'-spaces can be regarded as generalized spheres, in the sense that they retain two properties possessed by the spherical spaces: (a) They do not contain any proper subspaces homotopy equivalent to themselves (so they are "minimal" representatives of their homotopy type), and (b) loosely speaking, there is always a way to fold such a space "around itself", as illustrated in Figure 4.4.

5 | Representation of metapreferences

5.1 INTRODUCTION

In the last chapter, we considered strategy-proofness of general aggregation procedures. Earlier work on strategic disclosure of preferences has mainly focused on agents' behavior when faced with *social choice functions*, i.e. functions $f: \mathcal{P}^n \to X$, where $\mathcal{P}^n$ is an n-fold product space of individual preferences over the choice space X. The idea is then that individuals will not disclose their true preferences if they by reporting some other preference can ensure a more preferred (according to their own true preferences) social choice in X.

In this context, all preference relations involved have the same domain (the choice space X), although any preference may serve two different purposes: As an argument to the social choice function, and as a criterion for the individual to rank social choices when optimal strategic behavior is to be determined. This is of course possible only because the range of the social choice function is identical to the domain of the individual preferences.

With social welfare functions, where the range of the function is not the choice space X, but a space of social preferences, matters are different. We can not use the same class of preferences both as arguments to the welfare function, and as individuals' ranking criteria for social outcomes. If we want to analyze strategic disclosure of preferences when social welfare functions are involved, this requires two different classes of preferences: The preferences in one class are to be used as inputs to the function (these are the preferences that individuals reveal), and the other class, consisting of preferences with the range of the welfare function as domain, is used by individuals to rank social outcomes.

In the following, we will identify the space of social preferences (the range of the social welfare function) with the space of individual preferences. The two classes, or spaces, of preferences mentioned in the previous paragraph can thus be described as (a) one space of individual preferences over a choice space X, to be used as arguments to the welfare function, and (b) one space consisting of preferences with the space in (a) as domain; in other words, a space of preferences over preferences. Preferences in this latter space will be called *metapreferences.*

In this context, any individual may disclose a preference from the space in (a) above, and for each individual, one of the preferences in this space will be assumed to be his "true" preference. In addition, we assign a metapreference to each individual that (by ranking the possible outcomes of the welfare function) models his strategic disclosure of preferences.

The main result in this chapter is a theorem that shows how a metapreference over a space of preferences on a choice space X can be represented by a measure on $X \times X$. As in Chapter 3, we identify preferences with their graphs, so that any preference can be regarded as a set in $X \times X$. If there is a finite measure μ defined on $X \times X$, and an arbitrary preference P is singled out, we can define a utility function u on the space of preferences by letting $u(Q) = -\mu(P \triangle Q)$ for any preference

Q. This utility function generates a metapreference in the obvious way. It is, however, somewhat surprising that *any* metapreference that can be generated by a utility function satisfying some natural conditions, can be generated by a utility function *of this particular kind*, i.e., one that measures the symmetric set difference between graphs. This is the subject of the next section.

5.2 A REPRESENTATION THEOREM

We use the same notation as we did in Chapter 3, so that for a preference P, the symbol P itself refers to the preference graph, and $\succsim_P$ refers to the corresponding relation ($x \succsim_P y$ if and only if $(x, y) \in P$). To avoid confusion, we shall use a slightly different notation regarding metapreference relations; for a metapreference M, the relation is denoted $\overset{M}{\succsim}$.

As we did in Chapter 3, we still assume that the class of all possible preferences constitutes a σ-algebra $\mathcal{P}$. With preferences represented in this way, a metapreference is thus a special instance of what we will call a *set preference*. We define a set preference to be any preference relation on a σ-algebra of subsets of some space Y. Set preferences and metapreferences will have the same notation, i.e., for a set preference M, the relation is denoted $\overset{M}{\succsim}$.

In order to prove this chapter's main result, we shall first deduce a general characterization of set preferences that can be generated by a utility function that is a finite, nonatomic measure on Y. These will be the set preferences that satisfy the four conditions 5.1 to 5.4 below. Let M be a set preference on a σ-algebra $\mathcal{S}$ of subsets of Y.

The first condition essentially says that when ranking two sets in $\mathcal{S}$, the set preference should only depend upon the difference between the two sets, i.e., the ranking of the sets should be determined only by points that are in one of the two sets, but not in the other. We formulate this as follows:

Condition 5.1 *For any four sets A, B, C, and D in $\mathcal{S}$, if $A - B = C - D$ and $B - A = D - C$, then $A \overset{M}{\succsim} B$ if and only if $C \overset{M}{\succsim} D$.*

The next condition says that anything is as least as good as nothing, and thus introduces a direction to the set preference.

Condition 5.2 *For any set A in $\mathcal{S}$, we have $A \overset{M}{\succsim} \varnothing$.*

Conditions 5.1 and 5.2 together imply that the set preference is *monotonic*, that is, $F \supset E$ implies $F \overset{M}{\succsim} E$. The argument is simple. We obviously have

$$\begin{aligned} F - E &= (F - E) - \varnothing, \\ E - F &= \varnothing - (F - E). \end{aligned}$$

By Condition 5.1, this gives us $F \overset{M}{\succsim} E$ if and only if $(F - E) \overset{M}{\succsim} \varnothing$. By Condition 5.2, the latter expression is true.

We say that a set preference M is *nonatomic* if for any set A and any integer $n > 0$, it is always possible to find a partition of A into n disjoint subsets so that for any two subsets B and C in the partition, we have $B \overset{M}{\sim} C$.

Condition 5.3 *M is nonatomic.*

Finally, we require that the set preference shall be sufficiently well behaved, to the extent that it can be generated by some utility function.

Condition 5.4 *M can be generated by a real valued utility function f defined on the σ-algebra $\mathcal{S}$, in the sense that for any two sets A and B, we have $A \overset{M}{\succsim} B$ if and only if $f(A) \geq f(B)$. The function f is continuous*[1].

This condition clearly ensures that the set preference is transitive, but it also implies that the family of indifference classes of M (with the order induced by M) is order homeomorphic to a subset of $\mathbb{R}$.

It is clear that any set preference that satisfies Condition 5.1 has the following property: For any three sets A, B, and C, if C is disjoint from $A \cup B$, then $A \cup C \overset{M}{\succsim} B \cup C$ if and only if $A \overset{M}{\succsim} B$. This procedure (and the associated equivalence) will be called *disjoint addition.*

Lemma 5.1 *Let M be a transitive set preference that satisfies Condition 5.1. If A and B are disjoint sets, and C and D are disjoint sets, and $A \overset{M}{\succsim} C$ and $B \overset{M}{\succsim} D$, then $A \cup B \overset{M}{\succsim} C \cup D$.*

PROOF: The assumption that $A \overset{M}{\succsim} C$ implies (by disjoint addition) that

$$A \cup (B - C) \overset{M}{\succsim} B \cup C, \tag{5.1}$$

[1] For a formal definition of continuity for set preferences, see p. 112.

and $B \overset{M}{\succsim} D$ in the same way implies

$$B \cup C \overset{M}{\succsim} D \cup (C - B). \tag{5.2}$$

From (5.1) and (5.2) we get (by transitivity)

$$A \cup (B - C) \overset{M}{\succsim} D \cup (C - B).$$

By disjoint addition it follows that

$$A \cup (B - C) \cup (B \cap C) \overset{M}{\succsim} D \cup (C - B) \cup (B \cap C),$$

but this is of course nothing but

$$A \cup B \overset{M}{\succsim} D \cup C,$$

and the proof is complete. ∎

Trivial modifications of the proof give us the following two corollaries:

Corollary 5.2 *Let M be a transitive set preference that satisfies Condition 5.1. If A and B are disjoint sets, and C and D are disjoint sets, and $A \overset{M}{\succ} C$ and $B \overset{M}{\succ} D$, then $A \cup B \overset{M}{\succ} C \cup D$.*

Corollary 5.3 *Let M be a transitive set preference that satisfies Condition 5.1. If A and B are disjoint sets, and C and D are disjoint sets, and $A \overset{M}{\sim} C$ and $B \overset{M}{\sim} D$, then $A \cup B \overset{M}{\sim} C \cup D$.*

In order to prove Theorem 5.11, which says that a set preference satisfying Conditions 5.1–5.4 can be represented by a measure, we shall develop some intermediary results. We will in the following assume that the set preference M satisfies all of the four conditions above. We shall also make another assumption; it will be tacitly assumed that M is not the trivial set preference, i.e., that there is at least one set A in $\mathcal{S}$ for which $A \overset{M}{\succ} \varnothing$. By monotonicity and transitivity, this naturally implies $Y \overset{M}{\succ} \varnothing$.

For a given set preference M, we say that a set A has the ***k*-property** (or is a *k-set*) for some integer k whenever there exists a partition (called a k-partition) $\mathcal{E}$ of Y into k sets such that A is in $\mathcal{E}$, and for all B and C in $\mathcal{E}$, we have $B \overset{M}{\sim} C$.

Lemma 5.4 *For any k, if A is a k-set, then $A \overset{M}{\succ} \varnothing$.*

PROOF: Assume the proposition does not hold, i.e., $A \overset{M}{\sim} \varnothing$ (by monotonicity of M, we cannot have $\varnothing \overset{M}{\succ} A$). Let the k-partition that contains A be called $\mathcal{E}$, and enumerate the sets in $\mathcal{E}$ by E_i, $i = 1, \ldots, k$. By disjoint addition we have

$$\bigcup_{i=1}^{n} E_i \overset{M}{\sim} \left(\bigcup_{i=1}^{n-1} E_i \right) \cup \varnothing,$$

since (by transitivity) $E_n \overset{M}{\sim} \varnothing$. But this is of course the same as

$$\bigcup_{i=1}^{n} E_i \overset{M}{\sim} \bigcup_{i=1}^{n-1} E_i,$$

which, by induction, gives $Y \overset{M}{\sim} \varnothing$. This is contrary to the assumption we made at p. 75 that $Y \overset{M}{\succ} \varnothing$. ∎

Lemma 5.5 *For a given k, if A and B are both k-sets, then $A \overset{M}{\sim} B$.*

Proof: Assume this is not true, and that we have (without loss of generality) $A \overset{M}{\succ} B$. Clearly, this must imply that A and B come from two different partitions, which we call $\mathcal{E}$ and $\mathcal{F}$, respectively. By transitivity of M, any set in $\mathcal{E}$ must be strictly preferred to any set in $\mathcal{F}$. Enumerate the sets in $\mathcal{E}$ and $\mathcal{F}$ by E_i and F_i respectively, $i = 1, \ldots, k$. If

$$\bigcup_{i=1}^{n} E_i \overset{M}{\succ} \bigcup_{i=1}^{n} F_i,$$

then (by Corollary 5.2)

$$\bigcup_{i=1}^{n+1} E_i \overset{M}{\succ} \bigcup_{i=1}^{n+1} F_i,$$

for $n < k$. Of course, $E_1 \overset{M}{\succ} F_1$, so an induction argument leads to $Y \overset{M}{\succ} Y$, and this contradiction proves the lemma. ∎

Lemma 5.6 *If $A \overset{M}{\sim} B$ and $A \subset B$, then $(B - A) \overset{M}{\sim} \varnothing$.*

Proof: By contradiction. Assume $(B - A) \overset{M}{\succ} \varnothing$ (by Condition 5.2, we cannot have $\varnothing \overset{M}{\succ} (B - A)$). By disjoint addition we get

$$(B - A) \cup A \overset{M}{\succ} \varnothing \cup A$$

which is equivalent to $B \overset{M}{\succ} A$. This contradicts the assumption of the Lemma, and completes the proof. ∎

We say that A is a *weighted* set if there exists an n and a k such that A can be written as a union of n disjoint sets from a k-partition. In this case, A has *weight* n/k.

Lemma 5.7 *If A and B are weighted sets with weights n/k and m/j respectively, then $A \overset{M}{\succsim} B$ if and only if $n/k \geq m/j$.*

PROOF: We first show that $A \overset{M}{\succsim} B$ implies $n/k \geq m/j$. This is done by contradiction; assume $A \overset{M}{\succsim} B$ and $m/j > n/k$.

Since M is nonatomic, we can subdivide each of the k-sets into j subsets all indifferent to each other, and with an induction argument involving Corollary 5.2, it is not hard to see that the indifference relation will also hold between such subsets from different k-sets. This means that all the kj sets are indifferent to each other, and thus have the kj-property. Call this partition $\mathcal{E}$. In a similar way, divide all the j-sets into k sets each, and call the resulting partition of jk-sets $\mathcal{F}$. The set A is then written as nj disjoint sets from $\mathcal{E}$, and the set B is written as mk disjoint sets from $\mathcal{F}$. We assume that $mk > nj$, and derive a contradiction.

Enumerate the sets in $\mathcal{E}$ by E_i, $i = 1, \ldots, jk$, and the sets in $\mathcal{F}$ by F_i, $i = 1, \ldots, jk$. By Lemma 5.5, any set in $\mathcal{E}$ is indifferent to any set in $\mathcal{F}$. Thus, by Corollary 5.3, we have

$$\bigcup_{i=1}^{l+1} E_i \overset{M}{\sim} \bigcup_{i=1}^{l+1} F_i$$

whenever

$$\bigcup_{i=1}^{l} E_i \overset{M}{\sim} \bigcup_{i=1}^{l} F_i,$$

for $l < nj$. An induction argument then gives

$$A \overset{M}{\sim} \bigcup_{i=1}^{nj} F_i.$$

A is thus indifferent to a subset of B. By monotonicity and transitivity, this implies $B \overset{M}{\succsim} A$. This, together with the assumption of the lemma that $A \overset{M}{\succsim} B$, gives us $A \overset{M}{\sim} B$. Transitivity and Lemma 5.6 then implies

$$\bigcup_{i=nj+1}^{mk} F_i \overset{M}{\sim} \emptyset.$$

However, this, together with monotonicity and transitivity, contradicts Lemma 5.4 whenever $nj < mk$. This concludes the first part of the proof.

We then show that $n/k \geq m/j$ implies $A \overset{M}{\succsim} B$. With the same procedure as used above, write A as nj jk-sets, and B as mk jk-sets. By assumption, $nj \geq mk$, and an induction argument similar to the one above will show that B is indifferent to a subset of A. By transitivity and monotonicity, we get $A \overset{M}{\succsim} B$, and this proves the lemma. ∎

Corollary 5.8 *The weight of any weighted set A is unique.*

Lemma 5.9 *For any weighted set A with weight p, and any rational q where $0 \leq q \leq p$, there exists a weighted set $B \subset A$ with weight q.*

PROOF: Write p as n/k, and q as m/j. Partition the n k-sets into j subsets each, all indifferent to each other. Pick mk of these subsets, and form their union. This union has weight mk/kj. ∎

Lemma 5.7 does in particular imply that if A and B have the same weights, they are indifferent, and thus $f(A) = f(B)$. For any rational number q in $[0, 1]$, by Lemma 5.9 it is clearly possible to construct a set $W \subset Y$ with weight q. Define a map g' by

$$g'(q) = f(W),$$

where W has weight q. By the above remarks it does not matter which set W we choose, as long as its weight is q, since these sets all give the same value for f.

We now have a map g' from the rational numbers in $[0, 1]$ into the range of f. By Lemma 5.7 this map is strictly monotonic, and we shall see that is it uniformly continuous, and hence has a unique extension to a continuous one-to-one map on all of $[0, 1]$.

Lemma 5.10 *The map g' is uniformly continuous.*

PROOF: We must show that for any $\delta > 0$ there is an $\epsilon > 0$ so that $|x - y| < \epsilon$ implies $|g'(x) - g'(y)| < \delta$ for any two rational $x, y \in [0, 1]$.

Let $\mathcal{T}$ be the class of all pairs of weighted sets (A, B) that satisfy $|f(A) - f(B)| \geq \delta$. Let $w(A)$ be the weight of a weighted set A. By contradiction, we prove that $\inf |w(A) - w(B)|$, where the infimum is taken over all $(A, B) \in \mathcal{T}$, must be strictly positive. We can then set ϵ to the value of this infimum.

Assume that $\inf_{\mathcal{T}} |w(A) - w(B)| = 0$. This means that there exists a sequence of pairs $\{(A_i, B_i)\}$ from $\mathcal{T}$ such that

$$\lim_{i\to\infty} |w(A_i) - w(B_i)| = 0. \tag{5.3}$$

Consider now the sequence $\{w(A_i)\}$. From elementary topology we know that any sequence in $[0,1]$ has a convergent subsequence, so there is a subsequence of $\{w(A_i)\}$ with a limit in $[0,1]$. Also, since any convergent sequence in $[0,1]$ has a weakly monotonic (increasing or decreasing) subsequence, there will thus be a monotonic subsequence $\{w(A_{i_j})\}$ of $\{w(A_i)\}$. Without loss of generality, assume $\{w(A_{i_j})\}$ is weakly monotonic decreasing. By Lemma 5.9, we may thus assume $A_{i_k} \subset A_{i_j}$ for $k > j$.

It is a consequence of (5.3) that $\{w(B_{i_k})\}$ is a convergent sequence as well, with the same limit as $\{w(A_{i_k})\}$. Thus, $\{w(B_{i_k})\}$ has a weakly monotonic subsequence $\{w(B_{i_{k_l}})\}$. Without loss of generality, assume it is weakly monotonic decreasing. By Lemma 5.9, we may again assume $B_{i_{k_l}} \subset B_{i_{k_j}}$ for $l > j$.

Let $\underline{A}$ be the limit of $A_{i_{k_l}}$, and $\underline{B}$ the limit of $B_{i_{k_l}}$. Since

$$\lim_{l\to\infty} w(A_{i_{k_l}}) = \lim_{l\to\infty} w(B_{i_{k_l}}),$$

it is not hard to see that $\underline{A} \overset{M}{\sim} \underline{B}$, and thus $f(\underline{A}) = f(\underline{B})$. Since f is continuous from above, we must have

$$\lim_{l\to\infty} f(A_{i_{k_l}}) = \lim_{l\to\infty} f(B_{i_{k_l}}),$$

and thus

$$\lim_{l\to\infty} |f(A_{i_{k_l}}) - f(B_{i_{k_l}})| = 0.$$

This implies $\inf |f(A) - f(B)| = 0$, $(A, B) \in \mathfrak{T}$. But this contradicts the assumption that all $(A, B) \in \mathfrak{T}$ satisfy $|f(A) - f(B)| \geq \delta$, and the proof is complete. ∎

With the uniform continuity of g' established, we can extend it to a unique continuous map on all of $[0, 1]$. Call this extension g. Since g' is strictly increasing, this must also hold for g. Thus g has a strictly increasing inverse, g^{-1}. Consider the composite map $g^{-1} \circ f$. Call this map μ. Since g^{-1} is orderpreserving, the map μ considered as a utility function will generate the same set preference as f. It is left to show that μ is a measure.

Theorem 5.11 *If a set preference M on a σ-algebra $\mathcal{S}$ satisfies Conditions 5.1–5.4, then and only then can it be represented by a finite nonatomic measure μ, in the sense that $A \overset{M}{\succsim} B$ if and only if $\mu(A) \geq \mu(B)$.*

PROOF: If M is the trivial set preference, i.e. the preference that is indifferent between all alternatives, then we let μ be the measure that assigns zero to any set, and the theorem is obviously true. If M is nontrivial, we may (due to monotonicity) assume $Y \overset{M}{\succ} \varnothing$, and may thus employ all the results and definitions in the preceding lemmas and paragraphs.

Let μ be defined as $g^{-1} \circ f$. We must prove that μ is a finite nonatomic measure. We first show that μ is additive. From the definition of μ it is clear that if A is a weighted set, then $\mu(A) = w(A)$ (were as before, $w(A)$ is the weight of A). Furthermore, the reader should by now have no problems verifying that w is an additive set function on the weighted sets.

Let A and B be two disjoint sets in the σ-algebra. Thus, $\mu(A)$ and $\mu(B)$ are two real numbers in $[0, 1]$. Assume that μ is not additive for these two sets, and furthermore, that

$$\mu(A) + \mu(B) < \mu(A \cup B).$$

Pick two rational numbers c and d that satisfy $c > \mu(A)$, $d > \mu(B)$, and $c+d < \mu(A \cup B)$. Let C and D be two weighted sets with weights c and d, respectively. Since $c+d$ is less than 1 (because $\mu(A \cup B)$ must be less than or equal to 1), we may clearly pick D disjoint from C. By the fact that μ considered as a utility function generates M, and that for weighted sets, μ evaluates to the weights, we must have $C \overset{M}{\succ} A$ and $D \overset{M}{\succ} B$. Since w is additive, we must also have $A \cup B \overset{M}{\succ} C \cup D$. But this contradicts Corollary 5.2. By a symmetric argument, we may in the same way derive a contradiction from the assumption that $\mu(A)+\mu(B) > \mu(A \cup B)$. Since A and B were arbitrary disjoint sets, μ must be additive.

By Condition 5.4 and the fact that g^{-1} is continuous, μ is continuous from above. Since $\mu(\varnothing) = 0$, μ is nonnegative, and $\mu(Y) = 1$ implies that μ is finite. We may thus conclude that μ is a finite measure (Halmos [24, p. 39]). Condition 5.3 implies that μ is a nonatomic measure.

The converse statement, that a set preference generated by a finite measure satisfies the four conditions, follows trivially from the properties of measures. This completes the proof. ∎

Theorem 5.11 will for instance have applications in location theory; it gives necessary and sufficient conditions for a preference over different areas of land to be consistent with a per-unit price structure, i.e., a price structure where the land can be divided into arbitrarily small units, and where the price of any piece of land is equal to the sum of the prices of the units it consists of.

Our motivation for deriving the theorem, however, is due to the fact that it provides us with the necessary foundation to prove Theorem 5.12 below. This theorem says that under some reasonably weak conditions, any metapreference can be generated by a finite measure in a particular way.

For the rest of this chapter, it will be assumed that the space of preferences is a σ-algebra $\mathcal{P}$ of subsets of $X \times X$, where X is the choice space. The first condition displays certain similarities with Condition 3.2 in Chapter 3. Intuitively, the condition states that only the differences between preferences should determine how they are ranked by some metapreference. To formalize this condition, we use the same terminology we defined on p. 42.

Condition 5.5 *If two preferences P and Q disagree on exactly the same ordered pairs as two preferences R and S do, and furthermore, when restricted to these pairs, P is equal to R and Q is equal to S, then $P \overset{M}{\succsim} Q$ if and only if $R \overset{M}{\succsim} S$.*

Converted to set notation, this condition looks exactly the same as Condition 5.1.

In order to illustrate Condition 5.5 with an example, assume that the preferences in question are interpreted as the possible preferences of a set of candidates up for election, and the metapreference is a given voter's ranking of the candidates (or more correctly, their preferences). If two candidates agree on a particular issue, and then both change their opinions on the issue, but in such a way that their new opinions are still in agreement with each other, Condition 5.5 dictates that the voter's ranking of the two candidates should not change. In other words, whenever the voter ranks two candidates, he takes into consideration only issues where the candidates disagree.

The next condition says that for an agent with a metapreference M, there should be a preference P (the agents own preference, or his most preferred preference), such that for any preference Q, we have $P \overset{M}{\succsim} Q$.

Condition 5.6 *There is a preference P in $\mathcal{P}$ such that $P \overset{M}{\succsim} Q$ for all Q in $\mathcal{P}$.*

In other words, M should have a maximal element.

The last two conditions are just restatements of Conditions 5.3 and 5.4:

Condition 5.7 *M is nonatomic.*

Condition 5.8 *M can be generated by a real valued utility function u defined on the σ-algebra $\mathcal{P}$, in the sense that for any two sets Q and R, we have $Q \overset{M}{\succsim} R$ if and only if $u(Q) \geq u(R)$. The function u is continuous.*

Theorem 5.12 *If a metapreference M satisfies Conditions 5.5–5.8, then and only then can it be represented by a finite nonatomic measure μ, in the sense that $Q \overset{M}{\succsim} R$ if and only if $\mu(P \triangle Q) \leq \mu(P \triangle R)$.*

PROOF: Let $\mathcal{S}$ be the class of sets defined by

$$P \triangle Q \text{ is in } \mathcal{S} \text{ if and only if } Q \text{ is in } \mathcal{P},$$

where P is the preference referred to in Condition 5.6. First of all, observe that $\mathcal{S}$ is a σ-algebra, since we actually have $\mathcal{S} = \mathcal{P}$: Since $\mathcal{P}$, being a σ-algebra, is closed under symmetric set difference, $\mathcal{S}$ must be a subclass of $\mathcal{P}$, and $\mathcal{P}$ must be a subclass of $\mathcal{S}$ since any set Q in $\mathcal{P}$ can be written as $P \triangle (P \triangle Q)$.

Now define a set preference N on $\mathcal{S}$ by

$$P \triangle R \overset{N}{\succsim} P \triangle Q \text{ if and only if } Q \overset{M}{\succsim} R.$$

Clearly, if we can show that N satisfies Conditions 5.1–5.4, then, by Theorem 5.11, we are almost done.

Let us start with Condition 5.1. Assume that A, B, C, and D are four sets in $\mathcal{P}$. Furthermore, assume that

$$(P \triangle A) - (P \triangle B) = (P \triangle C) - (P \triangle D), \tag{5.4}$$

$$(P \triangle B) - (P \triangle A) = (P \triangle D) - (P \triangle C). \tag{5.5}$$

By elementary set theory we have

$$A - B = [(P \triangle A) - (P \triangle B) - P] \cup (P \cap [(P \triangle B) - (P \triangle A)])$$

for any two sets A and B, so (5.4) and (5.5) implies

$$A - B = C - D,$$
$$B - A = D - C,$$

and by Condition 5.5 we get

$$B \overset{M}{\succsim} A \text{ if and only if } D \overset{M}{\succsim} C. \tag{5.6}$$

By the definition of N this is equivalent to

$$(P \bigtriangleup A) \overset{N}{\succsim} (P \bigtriangleup B) \text{ if and only if } (P \bigtriangleup C) \overset{N}{\succsim} (P \bigtriangleup D).$$

But then we may conclude that N satisfies Condition 5.1.

Condition 5.2 is trivially satisfied by N, since Condition 5.6 implies that $P \bigtriangleup Q \overset{N}{\succsim} P \bigtriangleup P$ for all Q. It is furthermore obvious that N satisfies Condition 5.3. Finally, to see that Condition 5.4 is satisfied, consider the function f on $\mathcal{S}$ defined by $f(P \bigtriangleup Q) = -u(Q)$. This function will generate N, and it has the necessary continuity properties by the continuity properties of u.

By Theorem 5.11 we may conclude that there is a finite nonatomic measure μ such that we have $(P \bigtriangleup Q) \overset{N}{\succsim} (P \bigtriangleup R)$ if and only if it is the case that $\mu(P \bigtriangleup Q) \geq \mu(P \bigtriangleup R)$, and by the definition of N, this is equivalent to $Q \overset{M}{\succsim} R$ if and only if $\mu(P \bigtriangleup Q) \leq \mu(P \bigtriangleup R)$.

That a metapreference M that is generated by a finite nonatomic measure satisfies Conditions 5.5–5.8 follows trivially from the properties of measures. This completes the proof. ∎

5.3 CONCLUDING REMARKS

We have seen that under some plausible and weak conditions, any metapreference can be generated by a measure on $X \times X$. If this measure is a product measure, it can be derived from a measure on X. This is somewhat surprising, since it is conceptually a long way from a measure on X to preferences over preferences over X. The result also gives further support to the relevance of the topologies we examined in Chapter 3.

The measure on $X \times X$ that generates a metapreference, can be interpreted as reflecting the relative importance an agent with this metapreference assigns to disagreement on the ranking of pairs of alternatives. Disagreement on sets of pairs with small mass will thus (in a certain sense) be deemed less important than disagreement on sets with larger mass.

By applying Theorem 5.12, we can work with measures on $X \times X$ instead of the more abstract concept of a metapreference. We will do this in the next chapter, where we investigate strategic disclosure of preferences under social welfare functions.

6 Strategy-proofness and measure-based metrics

6.1 INTRODUCTION

In an innovative article by Bossert & Storcken [6], the authors investigate strategy-proofness of social welfare functions. In their paper, they assume that the choice set is finite, and that preferences are transitive, complete, and antisymmetric (i.e., linear orderings). It is also assumed that an individual with a preference P has a metapreference M generated by the Kemeny distance d_K between preferences, in the sense that $Q \overset{M}{\succsim} R$ if and only if $d_K(P,Q) \leq d_K(P,R)$.

With metapreferences defined in this way, a social welfare function is *manipulable by strategic voting* if a coalition of one or more individuals can achieve a more preferred outcome (according to the metapreferences) by insincere disclosure of preferences. A social welfare function that is not manipulable is called *coalitional strategy-proof.* Bossert & Storcken show that for choice sets with at least four alternatives, there are no coalitional strategy-proof welfare functions that satisfy some additional requirements (these requirements imply, among other things, that the social welfare functions must be nondictatorial).

It can be argued that the concept of coalitional strategy-proofness is a very strong one. It assumes that any group of individuals can form a coalition, agree on a disclosure strategy, and enforce the implementation of such a strategy. There may well be a number of factors that will make such an assumption implausible; if the electorate is large, the cost of organizing some coalitions may be prohibitive, and anyway, informational concealments like secret ballot may make it impossible for the coalitions to enforce their strategies. In that case, coalitional strategies are only plausible if they also are Nash equilibria.

A weaker assumption would be to require only that the social welfare function should be strategy-proof (as defined in Chapter 4); that is, that no single individual can achieve a more preferred outcome by insincere disclosure of his preferences. This means that the strategy where all agents report their true preferences should be a Nash equilibrium.

Our second remark on the framework of Bossert & Storcken concerns their use of the Kemeny distance. For linear orderings of a finite choice set, the Kemeny distance between two orderings P and Q (where P and Q denote the graphs of the orderings) can be defined as the number of elements in the set $P-Q$ (see Bogart [**5**]). This distance has a particular, arbitrary property associated with it; a property that is essentially due to one of the axioms Kemeny & Snell [**29**, p. 10] use to characterize the Kemeny distance. This is their Axiom 2, which says that if a preference Q_1 results from P_1 by a permutation of the alternatives, and Q_2 results from P_2 by the same permutation, then $d_K(P_1, P_2) = d_K(Q_1, Q_2)$.

We shall illustrate this with an example. Assume there are four alternatives in the choice space: Orange juice (o), water (w), skim milk (s), and regular milk (m). Consider the two preferences P_1 and P_2, where P_1 ranks

$$o \succ w \succ m \succ s$$

and P_2 ranks

$$o \succ w \succ s \succ m.$$

The point here is that P_1 and P_2 can be regarded as being rather close, since they only differ in the ranking between two alternatives of similar quality and nature (regular milk and skim milk). They both prefer orange juice to water, and water to milk.

Now consider the two preferences Q_1 and Q_2, where Q_1 ranks

$$s \succ m \succ w \succ o$$

and Q_2 ranks

$$s \succ m \succ o \succ w.$$

One can easily argue that these two preferences are not very close, since they differ in the ranking of two alternatives that possess rather different qualities, water and orange juice. But according to Kemeny & Snell's Axiom 2, the distance between Q_1 and Q_2 should be the same as the distance between P_1 and P_2, since Q_1 and Q_2 can be derived from P_1 and P_2, respectively, by one and the same relabeling of the alternatives (swap o and s, and w and m).

While disguised as an independence of what names we assign to the choices, this condition is of course much more comprising than that. Axiom 2 induces us to disregard *any* information pertaining to the alternatives; we are not at all allowed to take into consideration that a difference in the ranking of one pair of alternatives may be much more significant (and indicate a more substantial disagreement between the preferences) than a difference in the ranking of another pair.

While some applications involving the Kemeny distance may justify the validity of Axiom 2, it is highly unlikely that metapreferences generated by a metric should necessarily conform to this condition. Typically, when ranking social preferences, an individual will put more weight on differences between the ranking of some alternatives than of other alternatives. For instance, if the choice space is a space of possible resource allocations, an individual will likely be more concerned about differences over issues that affect how much he himself will be allotted, than issues that do not affect his share. We should allow for this possibility, and use a larger class of admissible metapreferences than those generated by the Kemeny distance.

The subject of this chapter integrates the material in the previous chapters. Like Bossert & Storcken, we shall analyze strategic disclosure of preferences under social welfare functions, but we will be content with social welfare functions that are strategy-proof (as in Chapter 4), instead of coalitional strategy-proof. On the other hand, we shall admit a wider class of metapreferences, and will only assume that they can be generated by a measure-based pseudometric — in other words, the metapreferences that were characterized in Chapter 5. This also allows us to consider choice spaces that are not necessarily finite, and preferences that are not necessarily linear orderings.

In Section 6.2, we deduce a few general properties of social welfare functions that are strategy-proof for the class of all measure generated metapreferences. These properties are used in Section 6.3 to prove that if all preferences are linear orderings and the social welfare function is onto, then strategy-proofness is equivalent to the existence of a dictator. In Section 6.4, we discuss the possibilities of extending this result to preferences that are not linear orderings.

6.2 STRATEGY-PROOF WELFARE FUNCTIONS

We shall investigate social welfare functions that are strategy-proof for a particular exhaustive[1] class of metapreferences, namely the class of all metapreferences that can be generated by a measure according to the procedure described in the previous chapter. As we did in Chapter 4, for reasons of notational simplicity we will only consider the case where there are two agents. We shall denote by Ω the space of all measurable preferences on a choice space X, i.e., Ω is a σ-algebra on $X \times X$, and whenever we talk about measures, it is assumed that the measures all have Ω as their domain. We shall also take it for granted that X contains three or more alternatives. Initially we analyze welfare functions that are defined on product spaces of Ω; later on, we will restrict the attention to the case where the preferences are assumed to be complete and transitive.

Assume now that the social welfare function $f\colon \Omega \times \Omega \to \Omega$ is strategy-proof for the class of metapreferences $\mathcal{M}$ that consists of any metapreference that can be generated by a measure with Ω as a domain. We shall assume that $\mathcal{M}$ also includes metapreferences generated by measures with atoms.

If f is strategy-proof for $\mathcal{M}$, this means that for any two preferences P and Q, and any measure-based pseudometric d, we must have

$$d[P, f(P,Q)] = \min_{R \in \Omega} d[P, f(R,Q)], \tag{6.1}$$

[1]See p. 58 for a definition.

and symmetrically,

$$d[Q, f(P,Q)] = \min_{R \in \Omega} d[Q, f(P,R)]. \tag{6.2}$$

We will achieve this if f has the property that for any P and Q

$$P \bigtriangleup f(P,Q) \subset P \bigtriangleup f(R,Q) \quad \text{for all } R \in \Omega \tag{6.3}$$

and

$$Q \bigtriangleup f(P,Q) \subset Q \bigtriangleup f(P,R) \quad \text{for all } R \in \Omega, \tag{6.4}$$

since in this case $d[P, f(P,Q)]$ will always be less or equal to $d[P, f(R,Q)]$ for any R in Ω, and the same relation holds between $d[Q, f(P,Q)]$ and $d[Q, f(P,R)]$. So if (6.3) and (6.4) hold, f will be strategy-proof. In fact, the converse is also true; f is strategy-proof for the class $\mathcal{M}$ only if it satisfies (6.3) and (6.4). To see this, assume that there is some R such that $P \bigtriangleup f(P,Q)$ is not a subset of $P \bigtriangleup f(R,Q)$. This means that the set $E = [P \bigtriangleup f(P,Q)] - [P \bigtriangleup f(R,Q)]$ is nonempty. Pick a point x in E, and define a measure μ with domain Ω by

$$\mu(A) = \begin{cases} 1 & \text{if } x \in A, \\ 0 & \text{if } x \notin A. \end{cases}$$

Clearly, since $x \in P \bigtriangleup f(P,Q)$ and $x \notin P \bigtriangleup f(R,Q)$, we must have $d_\mu[P, f(P,Q)] > d_\mu[P, f(R,Q)]$, and f would not be strategy-proof for $\mathcal{M}$. It is in general not necessary to resort to a measure with such extreme properties as μ, of course; the point is that we can always find a measure where E has a mass that is sufficiently large relative to that of sets disjoint with E.

The first result below shows how strong the notion of strategy-proofness is in this context, where we require that the functions should be strategy-proof for all of $\mathcal{M}$. We assess one agent's influence on the social outcome, under the *ceteris paribus* assumption that the other agent reports a preference Q.

Lemma 6.1 *If a social welfare function f is strategy-proof for $\mathcal{M}$, then, for a given preference Q, any ordered pair $(x, y) \in X \times X$ satisfy one of the following:*

(i) $(x, y) \in f(P, Q)$ *for all* P.

(ii) $(x, y) \notin f(P, Q)$ *for all* P.

(iii) $(x, y) \in f(P, Q)$ *if and only if* $(x, y) \in P$ *for all* P.

The situation is similar if we take P as given, and let Q vary.

PROOF: Expression (6.3) implies that for any P and Q, and any pair (x, y), if $(x, y) \notin P$, but $(x, y) \in f(P, Q)$, then $(x, y) \in f(R, Q)$ for all R. So (x, y) satisfy (i) in the theorem. Consider the set

$$\bigcup_{P \in \Omega} [f(P, Q) - P].$$

All ordered pairs (x, y) in this set will clearly satisfy (i). It is also implied by (6.3) that if $(x, y) \in P$, but $(x, y) \notin f(P, Q)$, then $(x, y) \notin f(R, Q)$ for any R. The set

$$\bigcup_{P \in \Omega} [P - f(P, Q)].$$

will thus be a set of pairs that satisfy (ii). The set of pairs that we have not accounted for so far is thus

$$X \times X - \bigcup_{P \in \Omega} [f(P,Q) - P] \cup [P - f(P,Q)].$$

It is easy to see that all pairs in this set will satisfy (iii).

Running through the same argument by using (6.4) instead of (6.3) will naturally prove the symmetric case with P fixed and Q varying. ∎

Lemma 6.1 says that given a preference Q, the space $X \times X$ can be partitioned into three sets, corresponding to the conditions (i)–(iii), such that if (x, y) is in one of the first two sets, the choice of P has no effect on the way $f(P,Q)$ ranks (x, y), and when restricted to the third set, P is equal to $f(P,Q)$. So, given the second agents preference, then, for any (x, y), either the first agent has no influence on the way the social preference ranks (x, y), or the social preference ranks (x, y) in the same way as he does himself. The result is of course similar if we fix the preference of the first agent, and analyze the situation for the second agent. Observe that Lemma 6.1 does not require that the social welfare functions should be continuous; it will hold for any social welfare function that is strategy-proof for $\mathcal{M}$.

This is about as far as we get without making further assumptions. However, many will object to this setting as being too general; we should at least require that whenever all individuals report preferences that are transitive, the social preference should be transitive as well, and whenever all individuals report preferences that are complete, the social preference should be complete. This will also make the conclusions more comparable to other results in the social choice literature, where it is usually assumed that all preferences involved are transitive and complete.

If we require that the social preference should be transitive/complete whenever all individuals report preferences that are transitive/complete, we can develop a stricter classification of strategy-proof welfare functions. Lemma 6.1 is of course still valid. Call the subsets of $X \times X$ that correspond to (i)–(iii) in Lemma 6.1 for I, II, and III, respectively; e.g., $(x, y) \in$ I if and only if (x, y) satisfy (i). Let Δ be the diagonal of $X \times X$, i.e., $\Delta = \{ (x, x) : x \in X \}$.

Lemma 6.2 *Assume that f is strategy-proof for $\mathcal{M}$, and evaluates to a transitive/complete preference whenever both arguments are transitive/complete. Assume also that the given preference Q in Lemma 6.1 is complete and transitive. Then* I, II, *and* III $\cup\, \Delta$ *are transitive sets.*

PROOF: We first show that I is transitive. Assume that (x, y) and (y, z) are in I. Then (x, z) cannot be in II, as $f(P, Q)$ would then never be transitive for any P. Obviously, (x, z) must be in $f(P, Q)$ for all transitive P. But there clearly exists some transitive P such that $(x, z) \notin P$, and it follows that (x, z) cannot be in III. This leaves one possibility only: (x, z) is in I.

Next assume that (x, y) and (y, z) are in II. We must have x, y, and z all different, since otherwise $f(P, Q)$ would never be complete for any P. Since (x, y) and (y, z) are never in $f(P, Q)$ for any P, we have $(x, z) \notin f(P, Q)$ for any complete and transitive P (since $f(P, Q)$ must then be complete and transitive). So (x, z) is not in I. Also, there is clearly some complete and transitive P for which $(x, z) \in P$, and thus, if (x, z) was in III, $(x, z) \in f(P, Q)$. It follows that (x, z) cannot be in III, so it must be in II.

Finally, we show transitivity of $\text{III} \cup \Delta$. Let (x, y) and (y, z) be in III. If either $x = y$, $y = z$, or $x = z$, we immediately have $(x, z) \in \text{III} \cup \Delta$. Assume that x, y, and z are all different. We cannot have $(x, z) \in \text{II}$, as there clearly exists a transitive P for which (x, y) and (y, z) are in P, so $(x, z) \in f(P, Q)$. Suppose that $(x, z) \in \text{I}$. There is clearly some complete and transitive P so that neither (x, y) nor (y, z) are in P, thus both (y, x) and (z, y) must be in $f(P, Q)$ (by completeness of $f(P, Q)$). But if $(x, z) \in \text{I}$, we must then have $(x, y) \in f(P, Q)$, by transitivity of $f(P, Q)$. But we picked a P for which $(x, y) \notin P$, so we have arrived at a contradiction to the assumption that (x, y) is in III. This shows that (x, z) is not in I.

Since (x, z) is neither in I nor II, it must be in III, and the proof is complete. ∎

The reason that transitivity will only hold for $\text{III} \cup \Delta$, and not for III alone, is that for any x, (x, x) is always a member of any complete preference. Moving (x, x) from III to I will thus not have any influence on the outcome of the welfare function whenever all the arguments are complete.

In fact, if we augment III some more, we can establish an even stronger property. Let S be the symmetric part of $\text{I} \cup \text{II}$, i.e.,

$$S = \{ (x, y) : (x, y) \in \text{I} \cup \text{II} \text{ and } (y, x) \in \text{I} \cup \text{II} \}.$$

Define the set D to be $(X \times X) - S$.

Lemma 6.3 *The set $D \cup \Delta$ is the graph of an equivalence relation on X.*

PROOF: We have to show that $D \cup \Delta$ is reflexive, symmetric and transitive. It is obvious that the set is reflexive, and it follows immediately from the definition of S that the set must be symmetric. To show transitivity, assume that (x, y) and (y, z) are two pairs in D (we may assume x, y, and z all different, since transitivity is otherwise trivial). We need to prove that (x, z) is in D, and this amounts to showing that at least one of (x, z) and (z, x) is in III.

That (x, y) and (y, z) are in D means that at least one of (x, y) and (y, x) is in III, and at least one of (y, z) and (z, y) is in III. Since III by Lemma 6.2 is a transitive set, the only nontrivial cases are when only (x, y) and (z, y) are in III, or only (y, x) and (y, z) are in III.

First, assume that (x, y) and (z, y) are in III. If neither (x, z) nor (z, x) is in III, at least one of them must be in I, or $f(P, Q)$ would never be complete for any P. If (x, z) is in I, we arrive at a contradiction, since there is clearly a transitive P for which $(z, y) \in P$ and $(x, y) \notin P$, but this would imply by transitivity $(x, y) \in f(P, Q)$ and contradict the assumption that (x, y) is in III. Similarly, if (z, x) is in I, there exists a transitive P for which $(x, y) \in P$ and $(z, y) \notin P$, but then by transitivity $(z, y) \in f(P, Q)$, and again we have a contradiction, this time to the assumption that (z, y) is in III. We have to conclude that at least one of (x, z) and (z, x) is in III.

The case where (y, x) and (y, z) are in III can be treated in a symmetrical way, and the details are left to the reader. This completes the proof. ∎

The equivalence relation in Lemma 6.3 will generate a partition of X in the usual way. This partition has the following interpretation: We say that the first agent has *veto power* on (x, y) for a given Q if $(x, y) \in P$ implies $(x, y) \in f(P, Q)$ (but $(y, x) \notin P$ will not necessarily imply $(y, x) \notin f(P, Q)$). Lemma 6.3 then says that if two points $x, y \in X$ both are members of some set in the partition (that is, both should be members of *the same* set), then the first agent has veto power on (x, y). If x and y are members of two different sets in the partition, then the first agent has no influence on the way the preference ranks x and y. So we can divide X into a class of mutually disjoint subsets, and the agent has veto power "within" each subset, but has no influence on the ranking of points from different subsets.

The results we have developed so far will be used in the next section to prove a theorem that is, at least on the surface, similar to Theorem 4.2, even though the methods employed are very different from those used in Chapter 4.

6.3 WELFARE FUNCTIONS THAT RESPECT UNANIMITY

In this section we will assume that all preferences are linear orderings. The aim is to derive a result that says that a strategy-proof welfare function that is onto, must be dictatorial.

Let $\mathcal{L}$ be the class of linear orderings on a (not necessarily finite) choice set X. $\mathcal{L}$ is thus a subclass of Ω. If $\mathcal{M}$ is the class of all metapreferences that can be generated by measures (as defined in the previous section), let $\overline{\mathcal{M}}$ be the class of the same metapreferences restricted to $\mathcal{L}$. Thus, any metapreference in $\overline{\mathcal{M}}$ has $\mathcal{L}$ as domain.

The lemmata in Section 6.2 will still hold true if we replace Ω with $\mathcal{L}$, and $\mathcal{M}$ with $\overline{\mathcal{M}}$, as it is easy to check that the arguments and proofs nowhere presuppose the existence of preferences that are not linear orderings.

In the previous section, we introduced the sets I, II, and III, corresponding to statements (i)–(iii) in Lemma 6.1. This lemma refers to a given preference Q, and in general, the sets I–III will be different for different values of Q. In the following proofs, we will often refer to the set III, and if the second agent's preference is fixed at Q, we write the set as $\mathrm{III}_1(Q)$ to indicate that it depends upon Q. When we look at the situation from the second agent's viewpoint, and keep the first argument fixed (at say, P), we write the set of pairs that satisfy Lemma 6.1(iii) as $\mathrm{III}_2(P)$.

Furthermore, when all preferences are linear orderings, and thus antisymmetric, it is easy to see that for any two points x and y (with $x \neq y$), if (x, y) satisfy Lemma 6.1(i), then (y, x) must satisfy (ii), and if (x, y) satisfy (ii), then (y, x) must satisfy (i). So $\mathrm{I} \cup \mathrm{II}$ will be symmetric, and the set D in Lemma 6.3 will thus be equal to III. So $\mathrm{III}_1(Q) \cup \Delta$ and $\mathrm{III}_2(P) \cup \Delta$ will both be graphs of equivalence relations. This also implies that $\mathrm{III}_1(Q)$ and $\mathrm{III}_2(P)$ are both symmetric sets.

In order to simplify the proof of the main result of this section, we shall develop three intermediary propositions. In Lemmas 6.4–6.6 it will be assumed that f is a social welfare function from $\mathcal{L} \times \mathcal{L}$ to $\mathcal{L}$ that respects unanimity and is strategy-proof for $\overline{\mathcal{M}}$.

For any linear ordering P, we shall define $-P$ to be the "opposite" ordering, i.e., $(x, y) \in -P$ if and only if $(y, x) \in P$.

Lemma 6.4 *For any linear ordering P and any two points x and y in X, we must have $(x, y) \in \mathrm{III}_1(P)$ if and only if $(x, y) \notin \mathrm{III}_2(-P)$.*

Proof: Assume $(x, y) \in \mathrm{III}_1(P)$. Thus, $(x, y) \in f(-P, P)$ if and only if $(x, y) \in -P$. But since P is antisymmetric, this means $(x, y) \in f(-P, P)$ if and only if $(x, y) \notin P$, so we cannot have $(x, y) \in \mathrm{III}_2(-P)$.

Assume $(x, y) \notin \mathrm{III}_1(P)$ and $(x, y) \notin \mathrm{III}_2(-P)$. This leads to a contradiction: Without loss of generality, assume $(x, y) \in f(-P, P)$. We must then have $(x, y) \in \mathrm{I}_1(P)$ and $(x, y) \in \mathrm{I}_2(-P)$. We will thus get $(x, y) \in f(P, P)$ and $(x, y) \in f(-P, -P)$. This is impossible since f respects unanimity, and P is antisymmetric. So $(x, y) \notin \mathrm{III}_1(P)$ must imply $(x, y) \in \mathrm{III}_2(Q)$. This completes the proof. ∎

Lemma 6.5 *For any linear ordering P, and any three points x, y, and z, where $x \neq y$ and $y \neq z$, we cannot simultaneously have $(x, y) \in \mathrm{III}_2(P)$ and $(y, z) \in \mathrm{III}_1(-P)$.*

Proof: By contradiction. Assume there are three such points so that $(x, y) \in \mathrm{III}_2(P)$ and $(y, z) \in \mathrm{III}_1(-P)$.

First assume $x = z$. This means $(x, y) \in \mathrm{III}_2(P)$ and $(y, x) \in \mathrm{III}_1(-P)$. Since $\mathrm{III}_1(-P)$ is a symmetric set, we must also have $(x, y) \in \mathrm{III}_1(-P)$. But this contradicts Lemma 6.4.

We may thus assume $x \neq z$. We must now have $(x,z) \notin \mathrm{III}_2(P)$, because $(x,z) \in \mathrm{III}_2(P)$ would imply $(y,z) \in \mathrm{III}_2(P)$ (since we assumed $(x,y) \in \mathrm{III}_2(P)$ and $\mathrm{III}_2(P) \cup \Delta$ is symmetric and transitive). By Lemma 6.4 this is impossible, since we assumed $(y,z) \in \mathrm{III}_1(-P)$.

We must also have $(x,z) \notin \mathrm{III}_1(-P)$, because $(x,z) \in \mathrm{III}_1(-P)$ would imply $(x,y) \in \mathrm{III}_1(-P)$ (since we assumed $(y,z) \in \mathrm{III}_2(P)$ and $\mathrm{III}_1(-P) \cup \Delta$ is symmetric and transitive). By Lemma 6.4 this is again impossible, since we assumed $(x,y) \in \mathrm{III}_2(P)$.

But $(x,z) \notin \mathrm{III}_2(P)$ and $(x,z) \notin \mathrm{III}_1(-P)$ contradicts Lemma 6.4, and this completes the proof. ∎

Lemma 6.6 *Let f be a social welfare function from $\mathcal{L} \times \mathcal{L}$ to $\mathcal{L}$ that respects unanimity, is strategy-proof for $\overline{\mathcal{M}}$, and is nondictatorial. Then there must exist some P in $\mathcal{L}$ such that*

$$f(P,-P) \neq P \quad \text{and} \quad f(P,-P) \neq -P.$$

Proof: We will assume that there is no P with this property, and then show that f is dictatorial.

Assume there is no such P. This means that $\mathcal{L}$ can be partitioned into two classes $\mathcal{L}_1$ and $\mathcal{L}_2$ defined by $P \in \mathcal{L}_1$ if $f(P,-P) = P$, and $P \in \mathcal{L}_2$ if $f(P,-P) = -P$.

Assume neither of $\mathcal{L}_1$ and $\mathcal{L}_2$ are empty. Pick a preference P from $\mathcal{L}_1$ and a preference Q from $\mathcal{L}_2$. It is not hard to see that as long as the choice set X contains more than two alternatives (as we postulated on p. 92), it must be possible to pick P and Q such that $-P \neq Q$.

By the definition of $\mathcal{L}_1$, we have $f(P, -P) = P$. Clearly, $\mathrm{III}_2(P)$ must be empty, since the second argument $-P$ disagree with the social outcome P for every pair of alternatives. By Lemma 6.4, this means that $\mathrm{III}_1(-P)$ must be equal to $X \times X$. But then we must have $f(Q, -P) = Q$.

Since Q is in $\mathcal{L}_2$, we have $f(Q, -Q) = -Q$. By a symmetric argument, we see that $\mathrm{III}_2(Q)$ must equal $X \times X$, and this implies $f(Q, -P) = -P$.

However, if $f(Q, -P) = Q$ and $f(Q, -P) = -P$, this implies $Q = -P$. But we picked P and Q such that $-P \neq Q$. We are thus forced to conclude that one of $\mathcal{L}_1$ and $\mathcal{L}_2$ is empty.

Assume without loss of generality that $\mathcal{L}_1 = \varnothing$. This implies that for all Q in $\mathcal{L}$, we have $f(Q, -Q) = -Q$. By the same argument as above, $\mathrm{III}_2(Q)$ must equal $X \times X$ (for all Q). But then we have, for any Q in $\mathcal{L}$, $f(Q, P) = P$ for all P in $\mathcal{L}$. This shows that the second agent is a dictator, and the proof is complete. ∎

It is now easy to show that strategy-proofness is equivalent to the existence of a dictator:

Theorem 6.7 *Let f be a social welfare function from $\mathcal{L} \times \mathcal{L}$ to $\mathcal{L}$ that is onto. Then f is strategy-proof for $\overline{\mathcal{M}}$ if and only if it is dictatorial.*

PROOF: Assume there is such a function f that is strategy-proof and not dictatorial. By the remark on page 66, we can assume that f respects unanimity (although the development in Chapter 4 requires continuous aggregation maps, the particular argument in the remark does not depend upon continuity, so it will also apply here, where no assumptions about continuity of the social welfare function are made).

Since f is not dictatorial, by Lemma 6.6 there will be a preference P in $\mathcal{L}$ such that $f(P,-P) \neq P$ and $f(P,-P) \neq -P$. There will then be two points x and y so that $(x,y) \in P$ and $(x,y) \in f(P,-P)$ (if there were no such points, $f(P,-P)$ would equal $-P$). Since $f(P,-P) \neq P$, there must also be two points w and z so that $(w,z) \in P$ and $(w,z) \notin f(P,-P)$.

Since $(x,y) \in f(P,-P)$ and $(x,y) \notin -P$, we must have $(x,y) \notin \mathrm{III}_2(P)$, and thus (by Lemma 6.4), $(x,y) \in \mathrm{III}_1(-P)$.

First assume that $z = x$. This means $(w,x) \in P$ and $(w,x) \notin f(P,-P)$. We must thus have $(w,x) \notin \mathrm{III}_1(-P)$, so by Lemma 6.4, $(w,x) \in \mathrm{III}_2(P)$. But this contradicts Lemma 6.2, since $(x,y) \in \mathrm{III}_1(-P)$.

We may thus assume $z \neq x$. Since $(x,y) \in \mathrm{III}_1(-P)$, this implies $(z,x) \notin \mathrm{III}_2(P)$, as $(z,x) \in \mathrm{III}_2(P)$ would contradict Lemma 6.5.

By Lemma 6.4, we must then have $(z,x) \in \mathrm{III}_1(-P)$. By Lemma 6.5 we would get $(w,z) \notin \mathrm{III}_2(P)$, so, by Lemma 6.4, $(w,z) \in \mathrm{III}_1(-P)$. But this is impossible, since $(w,z) \in P$ and $(w,z) \notin f(P,-P)$. ∎

6.4 CONCLUDING REMARKS AND FURTHER RESEARCH

In Sections 6.2 and 6.3 we assumed that metapreferences could be generated by measures that might include atoms. In many situations, it will be reasonable to disregard such measures. A subject for further research would be to take an even more general approach, and allow for the possibility that some sets in Ω are always assigned zero measure (for instance, all the countable sets). Call these sets *insignificant sets*. When we have decided on a class of insignificant sets, we can then define a class $\mathcal{N}$ of metapreferences, consisting of all metapreferences that can be generated by measures, under the restriction that these measures assign zero to the insignificant sets.

Equation (6.1) implies that for any P and Q,

$$\mu([P \triangle f(P,Q)] - [P \triangle f(R,Q)]) = 0 \quad \text{for all } R \in \Omega$$

for any measure μ that generates metapreferences in $\mathcal{N}$. In other words, the argument to μ in this equation must be an insignificant set for all R. If this was not the case, we would always be able to find a preference R and some measure μ so that μ assigns a mass to the set

$$E = [P \triangle f(P,Q)] - [P \triangle f(R,Q)]$$

that is arbitrarily large relative to the mass of sets disjoint from E. This would violate the condition expressed by (6.1).

In the following, when we say that a statement about points in $X \times X$ is valid *almost surely* (a.s.), we mean that the set of points that do not satisfy the statement is an insignificant set.

It is not hard to see that for any P, Q, and R, the two sets

$$E_1 = [P - f(P,Q)] \cap f(R,Q),$$
$$E_2 = [f(P,Q) - P] - f(R,Q)$$

will both be subsets of E, and thus insignificant sets. This means that, for any R and Q, the following two statements will hold:

(i) If $(x,y) \in f(R,Q)$, then we have $(x,y) \in f(P,Q)$ or

$$(x,y) \notin f(P,Q) \text{ only if } (x,y) \notin P,$$

a.s., for all P.

(ii) If $(x,y) \notin f(R,Q)$, then we have $(x,y) \notin f(P,Q)$ or

$$(x,y) \in f(P,Q) \text{ only if } (x,y) \in P,$$

a.s., for all P.

The argument is simple: The set of pairs that do not satisfy (i) will be the set of (x,y) for which $(x,y) \in f(R,Q)$, $(x,y) \notin f(P,Q)$, and $(x,y) \in P$, but this is just the set E_1 above. In the same manner, those pairs that do not satisfy (ii) belong to E_2.

There is clearly some similarity between (i) and (ii) above, and Lemma 6.1 in Section 6.2. The statements (i) and (ii) are together considerably weaker than Lemma 6.1, though. Let the set $f(R,Q)$ be called F_{i}, and let F_{ii} be the complement of F_{i}. The two statements (i) and (ii) apply to F_{i} and F_{ii}, respectively, and thus partition $X \times X$ into two subsets. However, to get something comparable in strength to Lemma 6.1, we ought to subdivide F_{i} and F_{ii} further into two subsets each; say, F_{i}^a and F_{i}^b (and F_{ii}^a and F_{ii}^b), so that $(x,y) \in F_{\mathrm{i}}^a$ implied $(x,y) \in f(P,Q)$ a.s. for all P, and $(x,y) \in F_{\mathrm{i}}^b$ implied $(x,y) \in f(P,Q)$ if and only if $(x,y) \in P$ a.s. for all P (and similarly with F_{ii}^a and F_{ii}^b).

Whether such a partition is possible is an open question. The problem is essentially that an arbitrary union of insignificant sets is not necessarily insignificant. With the setup in Section 6.2, only the empty set is insignificant, and since an arbitrary union of empty sets is empty, this accounts for the strength of Lemma 6.1.

The reason we would like a result comparable to Lemma 6.1 also in the case with nontrivial insignificant sets, is obviously that this might provide us with a result similar to Theorem 6.7. It would not be unreasonable to restrict the space of preferences to complete and transitive preferences without "thick" indifference sets, i.e., preferences where the indifference sets are insignificant sets. Such preferences would be generalizations of linear orderings, since linear orderings are what we would get in the special case where only the empty set is insignificant. If we then define a dictator to be an agent such that the social preference always differs from the agent's own preference on an insignificant set only, it might be possible to prove a result along the same lines as Theorem 6.7, but where linear orderings are replaced with preferences without thick indifference sets. This would significantly expand the range of situations where the result might be relevant.

The subjects addressed by this dissertation can be divided into two categories: We have investigated metrics on preference spaces and the associated topologies, and we have analyzed strategy-proofness of general aggregation maps, and social welfare functions in particular. In view of the relative simplicity of the results, it appears that the approach we have chosen — measuring the symmetric set difference between preference graphs — is a natural and fruitful one. In particular, the measure-based metrics have a property that we have barely touched upon: Applied to real world situations, it is actually possible to generate a statistical estimate of the distance between two preferences by looking at how the preferences rank alternatives from a *finite subset* of the choice space.

It is possible that this property might lead towards something like an econometric, or statistical, theory of social choice. Social choice theory has (so far) mainly been concerned with characterizing and analyzing "absolute" properties of choice or welfare functions, under the implicit assumption that all of the individual preferences in the domain have an equal status, without regard to the fact that some of these preferences (or combinations of them) may in some sense be unreasonable or less likely to occur. An alternative approach would be to look at how such functions perform "on the average". The measure-based metrics, in view of their natural interpretation as being generated by probability measures, might provide the foundation for such an analysis. While beyond the scope of this thesis, developments along these lines would make an interesting future research topic.

A | Mathematical prerequisites

This appendix lists the definitions of some common mathematical terms that occur in the previous chapters. It is not intended to be a tutorial, and can definitely not be used as one. It is merely included as a reference, to help refresh the memory of readers who are already acquainted with the subject matter.

A.1 GENERAL TOPOLOGY

A **pseudometric** is a real-valued function d on pairs of elements of a space M satisfying

(i) $d(x, y) \geq 0$.

(ii) $d(x, y) = 0$ if $x = y$.

(iii) $d(x, y) = d(y, x)$.

(iv) $d(x, y) + d(y, z) \geq d(x, z)$.

If (ii) holds with "if and only if" instead of "if", the pseudometric is called a **metric**. We say that M is a **metric space** with metric d. Let

x be any point of a metric space with metric d, and let ϵ be a positive number. The ϵ-neighborhood $S(x,r)$ of the point x is the set of all points y in M such that $d(x,y) < \epsilon$. The class of all ϵ-neighborhoods in M forms a basis for a topology. Such a topology is called a **metric topology**. For a topological space X, if there exists a metric such that the basis of ϵ-neighborhoods yields the original topology, X is called a **metrizable space.**

Let M be a metric space with metric d. A sequence $\{x_n\}$ of points in M is called a **Cauchy sequence** provided that for any positive number ϵ, there is an integer N_ϵ sufficiently large that $d(x_m, x_n) < \epsilon$ whenever m and n exceed N_ϵ. A metric space is **complete** if every Cauchy sequence of points in M has a limit point in M.

Let M and N be metric spaces with metrics d and ρ, respectively. A transformation $f\colon M \to N$ is continuous provided that for each point x in M and each positive real number ϵ, there exists a positive real number $\delta(x,\epsilon)$, in general depending on both x and ϵ, such that $\rho[f(x), f(y)] < \epsilon$ whenever $d(x,y) < \delta(x,\epsilon)$. If the number $\delta(x,\epsilon)$ can be chosen to be independent of the point x, we say that the transformation f is **uniformly continuous.**

A set Y of points in a space S is said to be **dense** in S if the closure of Y is S. A space is **separable** if it has a countable dense subset.

A space S is **Hausdorff** if, given two points of S, there exist two *disjoint* open sets, each containing just one of the two points.

Let S be a space that is not compact, and let ω be any abstract element not in S. The **one-point compactification** $\hat{S}$ of the space S consists of the points of $S \cup \{\omega\}$ with a basis for a topology of $\hat{S}$ consisting of (a) all open sets of S, and (b) all subsets U of $\hat{S}$ such that $\hat{S} - U$ is a closed compact subset of S.

A.2 MEASURE THEORY

A class $\mathcal{R}$ of subsets of a space X is called a **ring** if it has the following properties:

(i) $\varnothing \in \mathcal{R}$.

(ii) If $A \in \mathcal{R}$ and $B \in \mathcal{R}$, then $A - B \in \mathcal{R}$.

(iii) if $A \in \mathcal{R}$ and $B \in \mathcal{R}$, then $A \cup B \in \mathcal{R}$.

A ring $\mathcal{R}$ is called an **algebra** if it has the additional property that $X \in \mathcal{R}$.

A class $\mathcal{R}$ is called a **σ-algebra** if it is an algebra with the additional property:

$$\text{If } A_n \in \mathcal{R} \text{ for } n = 1, 2, \ldots, \text{ then } \bigcup_{n=1}^{\infty} A_n \in \mathcal{R}.$$

A **Borel-σ-algebra** on X is the smallest σ-algebra that contains all the open subsets of X (and thus, all the closed ones as well). A **Borel-measure** is a measure with a domain that includes a Borel-σ-algebra.

An extended real-valued function is a function with values in $\mathbb{R}$ or one of the two values ∞ and $-\infty$. A **set function** μ is an extended real-valued function defined on a class $\mathcal{D}$ of sets. Let μ be a set function defined on a ring $\mathcal{R}$. We say that μ is countably additive if

$$\mu\left(\bigcup_{n=1}^{\infty} E_n\right) = \sum_{n=1}^{\infty} \mu(E_n)$$

whenever the E_n are mutually disjoint sets of $\mathcal{R}$ such that their union is also in $\mathcal{R}$.

The **superior limit** of a sequence of sets $\{A_n\}$, written A^*, is the set of all points that are members of an infinite number of the sets in $\{A_n\}$. The **inferior limit** of a sequence of sets $\{A_n\}$, written A_*, is the set of all points that are members of all but a finite number of the sets in $\{A_n\}$. If the superior limit is the same as the inferior limit, we say that $\{A_n\}$ has a **limit**, and write $\lim\{A_n\} = A^* = A_*$.

A set function μ is **continuous** if $\mu(\lim\{A_n\}) = \lim_{n\to\infty} \mu(A_n)$ whenever $\{A_n\}$ has a limit.

A **measure** is a set function μ having the following properties:

(i) The domain $\mathcal{A}$ of μ is a σ-algebra.

(ii) μ is nonnegative on $\mathcal{A}$.

(iii) μ is countably additive on $\mathcal{A}$.

(iv) $\mu(\varnothing) = 0$.

A real-valued set function is called **finitely additive** if

$$\mu\left(\bigcup_{j=1}^{m} E_j\right) = \sum_{j=1}^{m} \mu(E_j)$$

for any finite collection of mutually disjoint sets E_j, $j = 1, \ldots, m$.

If μ is a measure and $\mu(X) < \infty$, we say that μ is a **finite measure**.

A measure μ with domain $\mathcal{A}$ is said to be **complete** if for any two sets N, E the following holds: If $N \subset E$, $E \in \mathcal{A}$ and $\mu(E) = 0$, then $N \in \mathcal{A}$.

A real-valued set function λ is said to be **absolutely continuous** with respect to a measure μ if for any $\epsilon > 0$ there exists a number $\delta > 0$ such that, for any measurable set E with $\mu(E) < \delta$, $|\lambda(E)| < \epsilon$.

An **atom** of a measure μ is a set E different from $\varnothing$ such that if $F \subset E$, then either $F = 0$ or $F = E$. A measure with no atoms is **nonatomic**.

Let μ and ν be finite measures on two spaces X and Y, respectively. For any set E in $X \times Y$, let $E_x = \{ y \in Y : (x, y) \in E \}$ and $E^y = \{ x \in X : (x, y) \in E \}$. Define a set function λ by

$$\lambda(E) = \int \nu(E_x)\, d\mu = \int \mu(E^y)\, d\nu$$

whenever all the E_x and E_y involved are measurable. Then λ is a measure, and it is called the **product** of the measures μ and ν, and we write $\lambda = \mu \times \nu$.

Fubini's theorem says that if h is a nonnegative, measurable function on $X \times Y$, then

$$\int h\, d(\mu \times \nu) = \iint h\, d\mu\, d\nu = \iint h\, d\nu\, d\mu.$$

A.3 HOMOTOPY THEORY

A subset R of a space S is a **retract** of S provided that there is a continuous map $r: S \to R$, such that $r(x) = x$ for each point x in R. Such a map r is called a **retraction**.

Two maps f and g of a space X into a space Y are **homotopic** if there is a map $h: X \times [0,1] \to Y$ such that for each point x in X,

$$h(x,0) = f(x) \quad \text{and} \quad h(x,1) = g(x).$$

The map h is called a **homotopy** between f and g.

The relation between maps from X to Y of being homotopic is an equivalence relation. The family of all continuous maps from X to Y can thus be decomposed by the homotopy relation into disjoint **homotopy classes**.

Two spaces X and Y are of the same **homotopy type** (are **homotopy equivalent**) if there exist maps $f: X \to Y$ and $g: Y \to X$ such that the composite maps $f \circ g$ and $g \circ f$ are homotopic, respectively, to the identity maps $i: Y \to Y$ and $i: X \to X$.

A space X is **contractible** if the identity map $i(x) = x$ of X onto itself is homotopic to the constant map $c(X) = p$ (where p is a point in X).

Bibliography

[1] K. ARROW. *Social Choice and Individual Values.* Cowles Foundation for Research in Economics Monograph 12, Yale University, New Haven, CT, 1963.

[2] G. AUMANN. "Über Räume mit Mittelbildungen." *Mathematische Annalen*, 119:210–215, 1943.

[3] P. BACON. "Unicoherence in means." *Colloquium Mathematicum*, 21:211–215, 1970.

[4] N. BAIGENT AND P. HUANG. "Topological social choice: Reply to LeBreton and Uriarte." *Social Choice and Welfare*, 7:141–146, 1990.

[5] K. P. BOGART. "Preference structures I: distances between transitive preference relations." *Journal of Mathematical Sociology*, 3:49–67, 1973.

[6] W. BOSSERT AND T. STORCKEN. "Strategy-proofness of social welfare functions: The use of the Kemeny distance between preference orderings." *Social Choice and Welfare*, 9:345–360, 1992.

[7] G. CHICHILNISKY. "Manifolds of preferences and equilibria." Technical Report 27, Project on Efficiency of Decision Making in Economic Systems, Harvard University, 1976.

[8] G. CHICHILNISKY. "On fixed point theorems and social choice paradoxes." *Economics Letters*, 3:347–351, 1979.

[9] G. CHICHILNISKY. "Social choice and the topology of spaces of preferences." *Advances in Mathematics*, 37:165–176, 1980.

[10] G. CHICHILNISKY. "Structural instability of decisive majority rules." *Journal of Mathematical Economics*, 9:207–221, 1982.

[11] G. CHICHILNISKY. "The topological equivalence of the pareto condition and the existence of a dictator." *Journal of Mathematical Economics*, 1982.

[12] G. CHICHILNISKY. "Social choice and the closed convergence topology." *Social Choice and Welfare*, 8:307–317, 1991.

[13] G. CHICHILNISKY. "On strategic control." *The Quarterly Journal of Economics*, 108:285–290, 1993.

[14] G. CHICHILNISKY AND G. HEAL. "Necessary and sufficient conditions for a resolution of the social choice paradox." *Journal of Economic Theory*, 31:68–87, 1983.

[15] G. CHICHILNISKY AND G. HEAL. "Patterns of power: Bargaining and incentives in two-person games." *Journal of Public Economics*, 23:333–349, 1984.

[16] J. S. CHIPMAN, L. HURWICZ, M. K. RICHTER, AND H. F. SONNENSCHEIN. *Preferences, Utility, and Demand.* Harcourt Brace Jovanovich, New York, 1971.

[17] G. DEBREU. "Neighboring economic agents." In *La Decision.* Editions du Centre National de la Recherche Scientifique, Paris, 1969.

[18] G. DEBREU. "Smooth preferences." *Econometrica*, 40:603–615, 1972.

[19] B. ECKMANN. "Räume mit Mittelbildungen." *Comm. Mathematici Helvetici*, 28:329–340, 1954.

[20] B. ECKMANN, T. GANEA, AND P. J. HILTON. "Generalized means." In *Studies in Mathematical Analysis and Related Topics*. Stanford University Press, Stanford, 1962.

[21] P. C. EKLOF. "Ultraproducts for algebraists." In J. Barwise, editor, *Handbook of Mathematical Logic*. North-Holland, 1977.

[22] A. FRIEDMAN. *Foundations of Modern Analysis*. Holt, Rinehart and Winston, Inc., New York, 1970.

[23] A. GIBBARD. "Manipulation of voting schemes: a general result." *Econometrica*, 41:587–601, 1973.

[24] P. HALMOS. *Measure Theory*. Van Nostrand, Inc., Princeton, 1950.

[25] F. HAUSDORFF. *Set Theory*. Chelsea Publ. Comp., New York, 1962.

[26] G. HEAL. "Contractibility and public decision-making." In P. K. Pattanaik and M. Salles, editors, *Social Choice and Welfare*, pages 103–121. North-Holland, 1983.

[27] W. HILDENBRAND. "On economies with many agents." *Journal of Economic Theory*, 2:161–188, 1970.

[28] Y. KANNAI. "Continuity properties of the core of a market." *Econometrica*, 38:791–815, 1970.

[29] J. G. KEMENY AND J. L. SNELL. *Mathematical Models in the Social Sciences.* Blaisdell, New York, 1962.

[30] M. LE BRETON. "Kannai topology with locally saturated preferences." D.P. 8608, Université de Rennes I, 1986.

[31] M. LE BRETON AND J. R. URIARTE. "On the robustness of the impossibility result in the topological approach to social choice." *Social Choice and Welfare*, 7:131–140, 1990.

[32] M. LE BRETON AND J. R. URIARTE. "Topological social choice: A rejoinder." *Social Choice and Welfare*, 7:147–148, 1990.

[33] C. R. F. MAUNDER. *Algebraic Topology.* Van Nostrand Reinhold Company, London, 1970.

[34] K. R. MOUNT AND S. REITER. "Construction of a continous utility function for a class of preferences." *Journal of Mathematical Economics*, pages 227–245, 1976.

[35] W. NEUEFEIND. "On continous utility." *Journal of Economic Theory*, 5:174–176, 1972.

[36] P. K. PATTANAIK. *Strategy and Group Choice.* North Holland, Amsterdam, 1978.

[37] B. PELY. *Game Theoretic Analysis of Voting in Committees.* Cambridge University Press, Cambridge, 1984.

[38] M. A. SATTERTHWAITE. "Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions." *Journal of Economic Theory*, 35:1–18, 1975.

[39] A. K. SEN. "Social choice theory." In K. J. Arrow and M. D. Intriligator, editors, *Handbook of Mathematical Economics, vol. III.* North-Holland, Amsterdam, 1986.

[40] J. R. URIARTE. "Topological structure of a space of continous preferences as a space of retractions and the aggregation problem." *Mathematical Social Sciences*, 13:259–272, 1987.