

Arrow's Theorem Through a Fixpoint Argument

Feys, Frank; Hansen, Helle

DOI

[10.4204/EPTCS.297.12](https://doi.org/10.4204/EPTCS.297.12)

Publication date

2019

Document Version

Final published version

Published in

Proceedings TARK 2019

Citation (APA)

Feys, F., & Hansen, H. (2019). Arrow's Theorem Through a Fixpoint Argument. In *Proceedings TARK 2019* (Vol. 297, pp. 175-188). (Electronic Proceedings in Theoretical Computer Science, EPTCS). Electronic Proceedings in Theoretical Computer Science. <https://doi.org/10.4204/EPTCS.297.12>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Arrow's Theorem Through a Fixpoint Argument

Frank M. V. Feys

Helle Hvid Hansen

Faculty of Technology, Policy and Management
Delft University of Technology
Delft, The Netherlands

`f.m.v.feys@tudelft.nl`

`h.h.hansen@tudelft.nl`

We present a proof of Arrow's theorem from social choice theory that uses a fixpoint argument. Specifically, we use Banach's result on the existence of a fixpoint of a contractive map defined on a complete metric space. Conceptually, our approach shows that dictatorships can be seen as "stable points" (fixpoints) of a certain process.

Keywords Social choice theory · voting · Arrow's impossibility theorem · Banach's fixpoint theorem · dictatorship · force · fixpoint · metric

1 Introduction

Arrow's impossibility theorem, introduced in Kenneth Arrow's seminal monograph *Social Choice and Individual Values* [2], deals with the issue of aggregating the preferences of a group of individuals into a single collective preference that appropriately represents the group.

Arrow defined a social welfare function to be a function that aggregates a collection of individual preferences into a societal preference, also called social choice. Preferences are defined as linear orders, where candidates are ranked from top to bottom. He considered the following three desirable properties that a social welfare function, which can be thought of as an election scheme or voting rule, might satisfy:

- *Pareto condition*. If all voters rank some candidate higher than another one, then also the social choice should do so.
- *Independence of irrelevant alternatives*. The societal ranking of any two candidates depends only on their relative rankings in the individual rankings, and on nothing else.
- *Non-dictatoriality*. There is no dictator, meaning that there is no voter whose individual preference always coincides with the societal outcome.

Arrow's impossibility theorem states that these seemingly mild conditions of reasonableness cannot be simultaneously satisfied by any voting rule.

In this paper we provide a new proof of this theorem that uses Banach's fixpoint theorem, which states that a contractive map on a complete metric space has a unique fixpoint [3]. Our proof method shows that dictatorships can be seen as "stable points" (fixpoints) of a certain process.

We give a sketch of our proof. First we define a metric parametrized by a probability distribution on profiles. The distance between two voting rules is the probability under the given distribution that the outcome of the election is different. The use of a distribution has the benefit that it allows for profiles to be considered concurrently. Inspired by the technical notion of influence from Boolean analysis [17], we then introduce a notion that we call "force" (a related idea was already presented in [8]). The force of a voter on a given voting rule is the probability that the outcome of the election coincides with that voter's preference. We use this notion to define a map that changes voting rules by considering the least powerful

voters' votes and transferring them to the most powerful voter. This map is shown to be a contraction. The process converges towards its unique fixpoint, because of Banach's fixpoint theorem, and we show this fixpoint to be the set of dictators. One technical point that is fundamental to our approach is that we must consider voting rules only up to equivalence via permutation of the voters, leading to a quotient metric space. This requires us to see to it that the original metric extends appropriately to the quotient metric. A certain group action takes care of that. Our technical development (metric space) is parametric in the distribution, and we study all notions as much as possible in generality. In the final step, we pick a specific distribution and derive Arrow's theorem from it.

Our proof does not require us to manipulate specific profiles, like most other previous proofs, but offers an analytic perspective in terms of fixpoints and convergence rather than a combinatoric one. The use of a probability distribution for measuring distance between voting rules is crucial in that respect. Our employment of fixpoints connects Arrow's theorem to the literature of mathematical economics, where many equilibrium concepts, such as for example Nash equilibrium [16], arise as a fixpoint.

We finish the introduction by discussing related work. A vast number of proofs have appeared since Arrow's first demonstration of the impossibility theorem. Arrow's original proof proceeded by showing the existence of a "decisive" voter. This approach was then refined by others, such as Blau [6] and Kirman and Sondermann [14], who used ultrafilters. Barbera [5] replaced the notion of a decisive voter with the weaker notion of a "pivotal" voter. In later work, Geanakoplos [11] and Reny [19] sharpened the latter approach by introducing the concept of "extremely pivotal" voter in order to obtain shorter proofs. Even though the aforementioned proofs are all distinct, they all are of a similar, combinatorial nature: One proves results about properties such as decisiveness or pivotalness by defining and manipulating profiles in a clever way that allows for the desirable properties to be taken advantage of. To our knowledge, Kalai's [13] proof of Arrow's theorem based on Fourier analysis on the Boolean cube was the first approach to take a different stance. Kalai's core idea was to calculate the probability of having a Condorcet cycle under some probability distribution. One advantage of this approach is that, by using a theorem from Boolean analysis by Friedgut, Kalai, and Naor [9], it produces a robust, quantitative version of the theorem, in the following sense: The more one seeks to avoid Condorcet's paradox, the more the election scheme will look like a dictator. Our approach is similar to Kalai's in that we use quantitative methods, but it differs from it as we prove the classical version of Arrow's theorem and not a quantitative one.

Contents of this paper. In Section 2, we give a brief overview of the basic formal concepts that are used in this paper. We also formulate Arrow's theorem in a precise way. In Section 3, we present our fixpoint-based proof of Arrow's theorem. Finally, we conclude and discuss future work in Section 4. The proofs that are not given in the main body of the text can be found in the appendix.

2 Preliminaries

In this section we introduce all basic notions that will be used later.

2.1 Arrovian Framework for Social Choice

In this subsection we recall the basic Arrovian framework, going back to Arrow's original work [2].

We assume that the n voters are linearly ordered, so without loss of generality we may suppose that the set of voters is $\{1, 2, \dots, n\}$ with the natural ordering. Given a set A , we let $\mathcal{L}(A)$ be the set of strict total orders on A . For our purposes, A will be the set of candidates of the election, which we shall always

assume to be finite and at least three. We can identify an element of $\mathcal{L}(A)$ with a complete listing of the elements of A in order of preference. If $\ell \in \mathcal{L}(A)$ and $a, b \in A$, we also write $a\ell b$ to mean that $(a, b) \in \ell$. We shall use $x = (x_1, x_2, \dots, x_n)$ to refer to an element of $\mathcal{L}(A)^n$. Such an element is usually called a *profile* in social choice theory. The element x_i is the *individual preference* of voter i .

Definition 2.1 A voting rule for n voters and set of candidates A is a map $\mathcal{L}(A)^n \rightarrow \mathcal{L}(A)$.

Here the interpretation is that if the preferences of the n voters are respectively $x_1, x_2, \dots, x_n \in \mathcal{L}(A)$, then the outcome of the election under the voting rule f is $f(x_1, x_2, \dots, x_n) \in \mathcal{L}(A)$. The latter is the so-called *social choice* given how the electorate voted.

A particularly bad type of voting rule is the following:

Definition 2.2 The n projection maps $\mathcal{L}(A)^n \rightarrow \mathcal{L}(A)$ are called dictatorships. The i -th dictator is denoted as Dict_i , where $i = 1, 2, \dots, n$. We furthermore define $\text{DICT}^n = \{\text{Dict}_i \mid i = 1, 2, \dots, n\}$.

Note that $\text{Dict}_i(x) = x_i$ for all $x = (x_1, x_2, \dots, x_n) \in \mathcal{L}(A)^n$. We write DICT^n shorthand as DICT .

Definition 2.3 Let $f: \mathcal{L}(A)^n \rightarrow \mathcal{L}(A)$ be a voting rule for n voters and set of candidates A . We say that f satisfies the

- Pareto property if for all $a, b \in A$ and all $x = (x_1, x_2, \dots, x_n) \in \mathcal{L}(A)^n$, if $a x_i b$ for all i , then $a f(x) b$.
- independence of irrelevant alternatives (IIA) property in case for all $a, b \in A$ and for all profiles $x = (x_1, x_2, \dots, x_n) \in \mathcal{L}(A)^n$ and $x' = (x'_1, x'_2, \dots, x'_n) \in \mathcal{L}(A)^n$, if $(a x_i b \text{ iff } a x'_i b)$ for all i , then $(a f(x) b \text{ iff } a f(x') b)$.

The Pareto property speaks for itself: It states that for any candidates a and b , if the *whole* electorate prefers a to b , then surely also in the social choice a ought to be preferred to b . The independence of irrelevant alternatives property is a bit more difficult to grasp. It says that the relative social ranking of two alternatives only depends on their relative individual rankings. In that sense there should be no dependence on any other alternative (hence “irrelevant”).

Definition 2.4 We let $\text{PIIA}^n = \{f: \mathcal{L}(A)^n \rightarrow \mathcal{L}(A) \mid f \text{ satisfies Pareto and IIA}\}$.

We write PIIA^n shorthand as PIIA . Note that $\text{DICT}^n \subseteq \text{PIIA}^n$.

Arrow set out to investigate whether there was a non-dictatorial voting rule satisfying both the Pareto condition as well as IIA. The answer turned out to be negative.

Theorem (Arrow [2]) A voting rule for at least three candidates that satisfies the Pareto property and IIA must be a dictatorship, i.e., $\text{PIIA}^n = \text{DICT}^n$ for all $n \geq 1$.

2.2 Metric Space Basics

We assume basic knowledge of metric spaces such as introduced in, e.g., [1]. Here we just recall a few basic definitions. A *metric space* (X, d) is a set X equipped with a metric.

A crucial element in our proof is Banach’s fixpoint theorem [3]. We recall that a function $F: X \rightarrow X$ on a metric space X is *contractive* if there is a $C < 1$ such that we have $d(F(x_1), F(x_2)) \leq C \cdot d(x_1, x_2)$ for all $x_1, x_2 \in X$. A *fixpoint* of F is an element x^* such that $F(x^*) = x^*$. For any map $F: X \rightarrow X$, by $F^{(n)}$ we mean the n -fold composition $F \circ \dots \circ F$.

Theorem 2.5 (Banach Fixpoint Theorem) Let (X, d) be a non-empty complete metric space. If $F: X \rightarrow X$ is contractive then F has a unique fixpoint x^* . For all $x \in X$, $x^* = \lim_{n \rightarrow \infty} F^{(n)}(x)$.

When X is a compact metric space it suffices to show that $d(F(x_1), F(x_2)) < d(x_1, x_2)$ for all $x_1, x_2 \in X$ to conclude that F is a contraction. This is in particular true if X is finite. Recall that if there is an n such that $F^{(n)}$ is a contraction, then F has a unique fixpoint (see, e.g., [18]). We shall use this fact later.

Let (X, d) be any metric space for which on the underlying set X an equivalence relation $\sim$ is defined. For reasons that will become clear later, we wish to extend the metric on X to a metric on $X/\sim$. In general, the following construction exists (see, e.g., [7] for the details). A *chain* C between points $x \in X$ and $y \in X$ is a sequence of points $x = a_0 \sim b_0, a_1 \sim b_1, \dots, a_n \sim b_n = y$, and the *length* of such a chain is defined as $\text{length}(C) = \sum_{i=0}^{n-1} d(b_i, a_{i+1})$ (this sum is understood to be zero when $n = 0$). Define

$$d_{\sim}([x], [y]) = \inf\{\text{length}(C) \mid C \text{ is a chain between } x \text{ and } y\}. \quad (1)$$

It is easy to see that $d_{\sim}$ is well-defined and satisfies all axioms of a pseudometric (meaning that the distance between two distinct points can be zero). Note though that $d_{\sim}([x], [y]) = 0 \Rightarrow [x] = [y]$ is not generally true. However, in case d is such that the infimum is always attained (which happens in particular in case X is finite), then $d_{\sim}$ is, in fact, a metric:

Lemma 2.6 *Let (X, d) be a metric space and $\sim$ an equivalence relation on X . If the infimum in $d_{\sim}$ is always attained, then $d_{\sim}$ is a metric on $X/\sim$.*

Proof. Indeed, if $d_{\sim}([x], [y]) = 0 = d_{\sim}([x], [y]) = \text{length}(C)$ for some chain C between x and y given by $x = a_0 \sim b_0, a_1 \sim b_1, \dots, a_n \sim b_n = y$, then $d(b_i, a_{i+1}) = 0$ for all i . Since d is a metric, we get $b_i = a_{i+1}$ for all i . Thus, $x = a_0 \sim a_1 \sim \dots \sim a_n \sim y$, which implies $x \sim y$ by transitivity. Hence, $[x] = [y]$. $\square$

Nonetheless, the defining formula for $d_{\sim}$ is difficult to work with. The following result gives a sufficient condition for the formula to reduce to an easier one. It uses the notion of action and orbit under a group action (see, e.g., [20]). A *group action* of a group G on a set X is a map $\varphi: G \times X \rightarrow X$ such that $\varphi(e, x) = x$ for all $x \in X$, where e is the identity element of G , and $\varphi(g, \varphi(g', x)) = \varphi(gg', x)$ for all $g, g' \in G$ and $x \in X$. The *orbit* of an $x \in X$ under the group action φ is the set $\{\varphi(g, x) \mid g \in G\}$. We also recall that the set of isometries on a metric space forms a group, with function composition as group law. With a group of isometries we mean a subgroup of that group.

Proposition 2.7 *If (X, d) is a metric space endowed with an equivalence relation $\sim$ where the equivalence classes are the orbits of the action of a group of isometries on (X, d) , then it holds for all $[x], [y] \in X/\sim$ that $d_{\sim}([x], [y]) = \inf\{d(x', y') \mid x' \in [x], y' \in [y]\}$.*

Proof. This is part of Theorem 2.1 from [7]. $\square$

3 Fixpoint Argument Proof

In this section, we give the new proof of Arrow's theorem. The general idea is to define a contraction that has the set of dictators as its unique fixpoint.

3.1 Strategy and Outline

We wish to use Banach's fixpoint theorem. In order to do so, we need to view our voting rules as part of some metric space. For technical reasons, we shall focus on voting rules that satisfy the Pareto property. Thus, we let M be the set of all voting rules $\mathcal{L}(A)^n \rightarrow \mathcal{L}(A)$ that are Pareto.

The idea to view voting rules as a metric space where the metric is based on a distribution is inspired by earlier work from the Boolean analysis approach to social choice theory, such as Mossel [15]. Viewing

the distance between voting rules as the probability that they produce a different outcome under a given distribution has the advantage that profiles can be considered synchronously. By framing voting in terms of probabilities, this naturally puts us on the way to a quantitative stance.

Definition 3.1 Given a probability distribution μ on $\mathcal{L}(A)^n$, we define a map $d_\mu: M \times M \rightarrow \mathbb{R}$ by $d_\mu(f, g) = \Pr_{x \sim \mu}[f(x) \neq g(x)]$.

That is, $d_\mu(f, g)$ is the probability under μ that f and g produce a different outcome. For any set X , by 1_X we mean the indicator function on X . Note that

$$d_\mu(f, g) = \Pr_{x \sim \mu}[f(x) \neq g(x)] = \sum_{x \in \mathcal{L}(A)^n} \mu(x) 1_{f(x) \neq g(x)}.$$

Throughout this paper, μ is a distribution on $\mathcal{L}(A)^n$, but for brevity we often simply say that “ μ is a distribution”. For example, taking μ to be the uniform distribution is known in the literature as the *impartial culture assumption* [10].

Under a mild assumption on μ we obtain our desired metric space.

Proposition 3.2 If μ is a probability distribution on $\mathcal{L}(A)^n$ with full support, then d_μ is a metric on M .

Proof. Clearly $d_\mu(f, f) = 0$ for all $f \in M$. Since μ has full support, $d_\mu(f, g) = \Pr_{x \sim \mu}[f(x) \neq g(x)] = 0$ implies $f = g$. Symmetry follows immediately. To prove the triangle inequality, let $f, g, h \in M$. Note that for any x , $f(x) \neq h(x)$ implies $f(x) \neq g(x)$ or $g(x) \neq h(x)$. Thus,

$$\Pr_{x \sim \mu}[f(x) \neq h(x)] \leq \Pr_{x \sim \mu}[f(x) \neq g(x) \text{ or } g(x) \neq h(x)],$$

and the latter is at most $\Pr_{x \sim \mu}[f(x) \neq g(x)] + \Pr_{x \sim \mu}[g(x) \neq h(x)]$. $\square$

Note that the metric space (M, d_μ) is complete and compact for any μ , since it is finite.

The proof idea is to define a contraction that has the set of dictators as its unique fixpoint. We start by defining a map Φ on M . Essentially, given a voting rule f , we will define $\Phi(f)$ as the voting rule in which the “least forceful” voter’s vote is replaced by the “most forceful” voter’s vote. This idea is inspired by, though different from, the notion of *influence* that is well-studied in Boolean analysis [17]. The influence of the i -th voter is defined as the probability that the i -th vote affects the outcome (assuming there are only two candidates). In political science and voting theory, influence also goes by the name *voting power*. Voting power can be thought of as the ability of a legislator, by his vote, to affect the passage or defeat of a measure [4].

To make the aforementioned idea formal, we define what we mean by the force of a voter.

Definition 3.3 Given a probability distribution μ on $\mathcal{L}(A)^n$ and voter i , we define the force of voter i on voting rule f under distribution μ as

$$\mathcal{F}_\mu^i[f] = \Pr_{x \sim \mu}[f(x) = x_i].$$

The set of most forceful voters is $\text{MostForce}_\mu[f] = \text{argmax}_{i \in \{1, 2, \dots, n\}} \mathcal{F}_\mu^i[f]$.

The set of least forceful voters is $\text{LeastForce}_\mu[f] = \text{argmin}_{i \in \{1, 2, \dots, n\}} \mathcal{F}_\mu^i[f]$.

The notion of force was first introduced in [8].

Definition 3.4 Let μ be a given distribution with full support. Given a voting rule $f \in M$, we define a new voting rule $\Phi^\mu(f) \in M$ by $\Phi^\mu(f)(x_1, \dots, x_n) = f(y_1, \dots, y_n)$, where

$$y_i = \begin{cases} x_i & \text{if } i \notin \text{LeastForce}_\mu[f], \\ x_{\min \text{MostForce}_\mu[f]} & \text{else.} \end{cases}$$

This defines a map $\Phi^\mu : M \rightarrow M$, also written for short as Φ if μ is understood from the context.

So what happens when passing from f to $\Phi(f)$ is that all voters with least force have their vote replaced by the vote of the first voter that has maximal force. Note that $\Phi(f)$ is Pareto if f is Pareto, so Φ is well-defined on M . The choice to consistently pick specifically the first voter in Definition 3.4 is arbitrary. We might as well always pick the last voter, for instance. Our choice is merely a technicality to ensure that the votes are transferred to one specific voter.

Proposition 3.5 *For all i we have $\Phi(\text{Dict}_i) = \text{Dict}_i$.*

Proof. Voter i is the unique most forceful voter of Dict_i . It is irrelevant how others vote, and we have $\Phi(\text{Dict}_i)(x_1, \dots, x_n) = x_i$ for all $(x_1, \dots, x_n)$. Hence, $\Phi(\text{Dict}_i) = \text{Dict}_i$. $\square$

Note that Φ cannot be a contraction, since *every* dictator is a fixed point of Φ , contradicting the uniqueness of the fixed point in Banach's fixpoint theorem. The idea is therefore to consider all dictators to be “the same”, i.e., to consider them to be equivalent under some equivalence relation.

Given an equivalence relation $\sim$, we write $[x]$ for the equivalence class of x under $\sim$. We shall write $M_\sim$ for $M/\sim$ throughout.

The remainder of the proof is organized as follows. We define an equivalence relation $\sim$ using a notion of permutation-invariance of voting rules as well as our notion of force. We then show that DICT is an equivalence class of $\sim$. The use of permutations gives rise to a notion of permutation-invariance for distributions. We then show that for all distributions μ on $\mathcal{L}(A)^n$ that have full support and are permutation-invariant, the following hold.

- (i) The map $\Phi^\mu : M^\mu \rightarrow M^\mu$ is compatible with $\sim$, where M^μ is the metric space resulting from μ , and hence extends to a map $\Phi_\sim^\mu : M_\sim^\mu \rightarrow M_\sim^\mu$.
- (ii) The map $\Phi_\sim^\mu$ has DICT as unique fixpoint.

Finally, we prove Arrow's theorem by showing the existence of a concrete distribution μ satisfying the abovementioned requirements.

3.2 Equivalence Relation on M and Metric on $M_\sim$

We first have to find an appropriate equivalence relation on M . To do that, we use the following notion.

Definition 3.6 *Given a permutation $\pi : n \rightarrow n$ of the voters, we define $\tilde{\pi}(x)$ as $(x_{\pi(1)}, \dots, x_{\pi(n)})$ for all tuples $x = (x_1, \dots, x_n) \in \mathcal{L}(A)^n$. If f is a voting rule, we define a new voting rule, written as $f \circ \tilde{\pi}$, by $(f \circ \tilde{\pi})(x) = f(\tilde{\pi}(x))$ for all $x \in \mathcal{L}(A)^n$.*

A permutation $\pi : n \rightarrow n$ is in what follows often written without types simply as π . Observe that $\tilde{\pi} : \mathcal{L}(A)^n \rightarrow \mathcal{L}(A)^n$ is a bijection because π is a bijection.

Definition 3.7 *We define the relation $\sim$ on M by letting $f \sim g$ iff (i) $f = g$, or (ii) there exists a permutation $\pi : n \rightarrow n$ such that $f = g \circ \tilde{\pi}$ and f has a unique voter with maximal force.*

In this way we obtain our desired equivalence relation.

Proposition 3.8 *The relation $\sim$ is an equivalence relation on M .*

Proof. Clearly $\sim$ is reflexive, as $= \subseteq \sim$. To show symmetry, let $f \sim g$, so that $f = g$, or $f = g \circ \vec{\pi}$ for some π where f has a unique maximal force voter. The former case $f = g$ is trivial. In the latter case, clearly $g = f \circ \vec{\pi}^{-1}$, and because f has a unique maximal force voter, the relation $f = g \circ \vec{\pi}$ implies that the same holds for g . Now to show transitivity, let $f \sim g$ and $g \sim h$. Thus, $f = g$, or $f = g \circ \vec{\pi}$ for some π where f has a unique maximal force voter. Also, $g = h$, or $g = h \circ \vec{\tau}$ for some τ where g has a unique maximal force voter. There are four cases. We will only show two of them, as they are all easy.

Suppose $f = g$ and $g = h \circ \vec{\tau}$ where g has a unique maximal force voter. Then since g has a unique maximal force voter, the same holds for h , so the relation $f = h \circ \vec{\tau}$ implies that this is also true for f . If $f = g \circ \vec{\pi}$ and $g = h \circ \vec{\tau}$, where both f and g have a unique maximal force voter, we have that $f = h \circ (\vec{\tau} \circ \vec{\pi})$. $\square$

Note that $\text{DICT} \subseteq M$, so it makes sense to speak about $[\text{Dict}_i]$, and that the equivalence classes of $\sim$ are of the following two types:

- If f does not have a unique maximal force voter, then $[f] = \{f\}$.
- If f does have a unique maximal force voter, then $[f] = \{f \circ \vec{\pi} \mid \pi: n \rightarrow n\}$.

Proposition 3.9 *The set DICT is an equivalence class of $\sim$.*

Proof. For any i , we show that we have $[\text{Dict}_i] = \text{DICT}$. Let $f \in M$ be such that $f \sim \text{Dict}_i$. If $f = \text{Dict}_i$, surely $f \in \text{DICT}$. Otherwise, we have that $f = \text{Dict}_i \circ \vec{\pi}$ for some permutation π . Suppose $\pi(i) = j$. A trivial computation then shows that $f = \text{Dict}_j$, so $f \in \text{DICT}$. We now show that for any j , $\text{Dict}_i \sim \text{Dict}_j$. Indeed, if π is the permutation that switches i and j and is the identity elsewhere, then $\text{Dict}_i = \text{Dict}_j \circ \vec{\pi}$. Furthermore, it is clear that any dictator has a unique voter with maximal force. $\square$

Let μ be a distribution with full support such that $d = d_\mu$ is a metric on M (by Proposition 3.2), and consider the quotient metric $d_\sim$ defined in (1) for the relation $\sim$ from Definition 3.7. Since M is finite, we obtain from Lemma 2.6 that $d_\sim$ is a metric on $M_\sim = M/\sim$. However, as we saw in that same subsection, the defining formula for $d_\sim$ is convoluted, so we set out to identify a subgroup of isometries that will allow us to apply Proposition 2.7, as that would give us a more convenient formula to work with. To do that, we shall need the following notion.

Definition 3.10 *A distribution μ on $\mathcal{L}(A)^n$ is n -permutation-invariant if for all $\pi: n \rightarrow n$, $\mu \circ \vec{\pi} = \mu$.*

The condition $\mu = \mu \circ \vec{\pi}$ for all permutations π demands that $\mu(x) = \mu(x')$ whenever x' and x are rearrangements of one another. Mathematically, it ensures that the probability distribution μ is well-defined up to permutation equivalence of profiles.

For each permutation $\pi: n \rightarrow n$, let $J_\pi: M \rightarrow M: f \mapsto f \circ \vec{\pi}$. This map is well-defined, as $f \circ \vec{\pi}$ is Pareto if f is.

Proposition 3.11 *Let μ be a distribution on $\mathcal{L}(A)^n$ with full support and such that μ is n -permutation-invariant. Then J_π is an isometry of M , for each π .*

Proof. From the bijectivity of π , it follows that J_π is bijective. Let π be a permutation and $f, g \in M$. Then we need to show that $d_\mu(f, g) = d_\mu(f \circ \vec{\pi}, g \circ \vec{\pi})$.

We calculate:

$$\begin{aligned}
 \Pr_{x \sim \mu}[(f \circ \tilde{\pi})(x) \neq (g \circ \tilde{\pi})(x)] &= \sum_{x \in \mathcal{L}(A)^n} \mu(x) \cdot 1_{(f \circ \tilde{\pi})(x) \neq (g \circ \tilde{\pi})(x)} \\
 &= \sum_{x \in \mathcal{L}(A)^n} \mu(\tilde{\pi}^{-1}(x)) \cdot 1_{(f \circ \tilde{\pi})(\tilde{\pi}^{-1}(x)) \neq (g \circ \tilde{\pi})(\tilde{\pi}^{-1}(x))} \\
 &= \sum_{x \in \mathcal{L}(A)^n} \mu(\tilde{\pi}^{-1}(x)) \cdot 1_{f(x) \neq g(x)}.
 \end{aligned}$$

Hence, $\Pr_{x \sim \mu}[(f \circ \tilde{\pi})(x) \neq (g \circ \tilde{\pi})(x)] = \Pr_{x \sim \mu \circ \tilde{\pi}^{-1}}[f(x) \neq g(x)]$. Note that $\mu \circ \tilde{\pi}^{-1}: \mathcal{L}(A)^n \rightarrow [0, 1]$ is indeed a probability distribution on $\mathcal{L}(A)^n$. The proof is complete after observing that $\mu \circ \tilde{\pi}^{-1} = \mu$. $\square$

Let $G = \{J_\pi \mid \pi \text{ is a permutation}\}$. From Proposition 3.11 it follows that G is a subgroup of the group of isometries of M . We define a map $\cdot: G \times M \rightarrow M$ by

$$J_\pi \cdot f = \begin{cases} J_\pi(f) & \text{if } f \text{ has a unique maximal force voter,} \\ f & \text{otherwise.} \end{cases}$$

Proposition 3.12 *The operation $\cdot$ is a group action of the group G on M and its orbits coincide with the equivalence classes under $\sim$.*

Proof. Note that $\text{id}: M \rightarrow M$ is the identity of G , and $\text{id} \cdot f = f$ for all $f \in M$. Also, $(g \circ h)(f) = g(h(f))$ for all $g, h \in G$ and $f \in M$. Indeed, this easily follows by making the case distinction whether f has a unique maximal force voter.

The orbit of an $f \in M$ under this group action is equal to $G \cdot f = \{J_\pi \cdot f \mid \pi\}$. Now, $J_\pi \cdot f$ is $J_\pi(f) = f \circ \tilde{\pi}$ if f has a unique maximal force voter, and f otherwise. Hence, $G \cdot f = \{J_\pi \cdot f \mid \pi\} = [f]$. Thus, the orbits of the group action coincide with the equivalence classes under $\sim$. $\square$

Proposition 3.13 *Let μ be a distribution with full support and with $\mu = \mu \circ \tilde{\pi}$ for all permutations π . Let $d_\sim$ be the metric on $M/\sim$ based on the metric $d = d_\mu$ on M (see Lemma 2.6). Then for all $[f], [g] \in M/\sim$ we have $d_\sim([f], [g]) = \min\{d(f', g') \mid f' \in [f]_\sim, g' \in [g]_\sim\}$.*

Proof. This follows from Proposition 2.7, Proposition 3.11, and Proposition 3.12. $\square$

3.3 Extending Φ to a Map with Unique Fixpoint

From now on, unless specifically mentioned otherwise, we shall always assume that μ is a distribution on $\mathcal{L}(A)^n$ with full support and with $\mu = \mu \circ \tilde{\pi}$ for all permutations π , such that Proposition 3.13 applies. Such a μ obviously exists, for instance the uniform distribution is an example, but one can easily obtain many other examples simply by giving one representative of each permutation equivalence class of profiles (meaning profiles x and y are equivalent iff there is a permutation π such that $\pi(x) = y$) a non-zero weight. Later we will exploit this property.

Our aim in this subsection is to extend the map $\Phi: M \rightarrow M$ to a map $M_\sim \rightarrow M_\sim$ where, as we recall, $M_\sim = M/\sim$. The following is a technical result that we shall use later.

Lemma 3.14 *If $f = g \circ \tilde{\pi}$ for π , then $\mathcal{F}_\mu^i[g] = \mathcal{F}_\mu^{\pi(i)}[f]$ for each $i = 1, 2, \dots, n$.*

Proof. For any i , let $p_i: \mathcal{L}(A)^n \rightarrow \mathcal{L}(A)$ be the i -th projection map. Then we have

$$\mathcal{F}_\mu^i[g] = \Pr_{x \sim \mu}[g(x) = x_i] = \sum_x \mu(x) \cdot 1_{g(x)=p_i(x)} = \sum_x \mu(\vec{\pi}(x)) \cdot 1_{g(\vec{\pi}(x))=p_i(\vec{\pi}(x))}.$$

Note that $p_i(\vec{\pi}(x)) = x_{\pi(i)}$ for each x , so $\mathcal{F}_\mu^i[g] = \sum_x \mu(\vec{\pi}(x)) \cdot 1_{g(\vec{\pi}(x))=x_{\pi(i)}}$. By assumption, $\mu \circ \vec{\pi} = \mu$ and $f = g \circ \vec{\pi}$, so the conclusion follows. $\square$

We now show that Φ respects the equivalence relation $\sim$, so that it extends to the quotient.

Proposition 3.15 *For all $f, g \in M$, if $f \sim g$ then $\Phi(f) \sim \Phi(g)$.*

Proof. Let $f, g \in M$ be such that $f \sim g$. If $f = g$, then clearly $\Phi(f) = \Phi(g)$. Now suppose there is a π such that $f = g \circ \vec{\pi}$, and that f (and hence also g) has a unique voter with maximal force. We need to show that $\Phi(f) \sim \Phi(g)$. In fact, we will show that $\Phi(f) = \Phi(g) \circ \vec{\pi}$. From Lemma 3.14 we know that $\mathcal{F}_\mu^i[g] = \mathcal{F}_\mu^{\pi(i)}[f]$ for each i . This implies that π maps the first (and only) most forceful voter of g to the first (and only) most forceful voter of f . Similarly, π maps the least forceful voters in g to the least forceful voters in f . Applying the definition of Φ (see Definition 3.4), we see that $\Phi(f) = \Phi(g) \circ \vec{\pi}$. $\square$

Definition 3.16 *We define $\Phi_\sim: M_\sim \rightarrow M_\sim$ by $\Phi_\sim([f]) = [\Phi(f)]$.*

We want to point out that, although we have not written it explicitly, $\Phi_\sim$ (and Φ alike) does depend on a chosen distribution μ . By Proposition 3.15, $\Phi_\sim$ is well-defined. We also define for each $i = 1, 2, \dots, n$ a map $s_i: \mathcal{L}(A)^n \rightarrow \mathcal{L}(A)^n$ by $s_i(x_1, \dots, x_n) = (x_i, \dots, x_i)$.

The following is a technical result.

Lemma 3.17 *For any $f \in M$ and i , voter i is the unique voter with maximal force on $f \circ s_i$. Moreover, $[f \circ s_i] = \{f \circ s_k \mid k = 1, 2, \dots, n\}$ for each i .*

Proof. Since f is Pareto, we have

$$\mathcal{F}_\mu^j[f \circ s_i] = \Pr_{x \sim \mu}[f(s_i(x)) = x_j] = \Pr_{x \sim \mu}[x_i = x_j]$$

for any j . Now as μ is assumed to have full support, $\Pr_{x \sim \mu}[x_i = x_j] = 1$ iff $x_i = x_j$ for all $x \in \mathcal{L}(A)^n$, i.e., iff $i = j$. This proves that i is the unique voter with maximal force.

From the first part we know that $f \circ s_i$ has a unique voter with maximal force. Hence we have $[f \circ s_i] = \{(f \circ s_i) \circ \vec{\pi} \mid \pi\} = \{f \circ s_k \mid k = 1, 2, \dots, n\}$. $\square$

Proposition 3.18 *The map $\Phi_\sim: M_\sim \rightarrow M_\sim$ has a unique fixpoint.*

Proof. Since $M_\sim$ is finite, it suffices to show $d_\sim(\Phi_\sim^{(n)}([f]), \Phi_\sim^{(n)}([g])) < d_\sim([f], [g])$ for all $[f], [g] \in M_\sim$. So let $[f], [g] \in M_\sim$. Every iteration of Φ on f , at least one voter loses their vote as it is taken over by the most forceful voter. Thus, there is an i such that $\Phi^{(n)}(f)(x_1, \dots, x_n) = f(x_i, \dots, x_i)$ for all $(x_1, \dots, x_n)$, or in other words, $\Phi^{(n)}(f) = f \circ s_i$. There is similarly a j such that $\Phi^{(n)}(g) = g \circ s_j$.

We have

$$d_\sim(\Phi_\sim^{(n)}([f]), \Phi_\sim^{(n)}([g])) = d_\sim([\Phi^{(n)}(f)], [\Phi^{(n)}(g)]) = d_\sim([f \circ s_i], [g \circ s_j]).$$

Applying Proposition 3.13 and Lemma 3.17, we obtain $d_\sim(\Phi_\sim^{(n)}([f]), \Phi_\sim^{(n)}([g])) = \min_{k,l} d(f \circ s_k, g \circ s_l)$. Since f and g are Pareto, we have

$$(f \circ s_k)(x_1, \dots, x_n) = f(x_k, \dots, x_k) = x_k \quad \text{and} \quad (g \circ s_l)(x_1, \dots, x_n) = g(x_l, \dots, x_l) = x_l$$

for all $(x_1, \dots, x_n) \in \mathcal{L}(A)^n$. Thus, $\min_{k,l} d(f \circ s_k, g \circ s_l) = \min_{k,l} \Pr_{x \sim \mu}[x_k \neq x_l] = 0$, and from this it follows that for all $d_\sim([f], [g]) \neq 0$ we have $d_\sim(\Phi_\sim^{(n)}([f]), \Phi_\sim^{(n)}([g])) = 0 < d_\sim([f], [g])$. So $\Phi_\sim$ collapses all equivalence classes after n steps. Note that $\Phi_\sim^{(n)}([f]) = \text{DICT}$ for all $f \in M$. $\square$

We now show that this fixpoint is the equivalence class of dictators.

Proposition 3.19 *It holds that $\Phi_\sim(\text{DICT}) = \text{DICT}$.*

Proof. From Proposition 3.5 we know that $\Phi(\text{Dict}_i) \sim \text{Dict}_i$ for each i . This implies the equality $\Phi_\sim(\text{DICT}) = \Phi_\sim([\text{Dict}_i]) = [\Phi(\text{Dict}_i)] = [\text{Dict}_i] = \text{DICT}$, where we used Proposition 3.9. $\square$

We shall now develop a technical result needed for proving Arrow's theorem. For any i , we define $v_i: \mathcal{L}(A)^n \rightarrow \mathcal{L}(A)^{n-1}$ by $v_i(x) = x_{-i}$, where x_{-i} is the same vector as x but with the i -th component left out (so $x = (x_i, x_{-i})$ for all x and i). Given any distribution μ on $\mathcal{L}(A)^{n-1}$ and $i = 1, 2, \dots, n$, we can associate with μ a real-valued map $\mu^{[i]}$ on $\mathcal{L}(A)^n$ by defining

$$\mu^{[i]}(x) = \frac{\sum_{\tau: n \rightarrow n} \mu(v_i(\vec{\tau}(x)))}{n! |\mathcal{L}(A)|}. \quad (2)$$

Lemma 3.20 *Let μ be a distribution on $\mathcal{L}(A)^{n-1}$ and $i \in \{1, 2, \dots, n\}$. Then $\mu^{[i]}$ is a distribution on $\mathcal{L}(A)^n$ and $\mu^{[i]}$ is n -permutation-invariant. If moreover μ has full support, then also $\mu^{[i]}$ has full support.*

Proof. To show that $\mu^{[i]}$ is a distribution, note that for a fixed permutation $\tau: n \rightarrow n$,

$$\sum_{x \in \mathcal{L}(A)^n} \mu(v_i(\vec{\tau}(x))) = \sum_{x \in \mathcal{L}(A)^n} \mu(v_i(x)) = \sum_{x_i} \sum_{x_{-i}} \mu(x_{-i}) = \sum_{x_i} 1 = |\mathcal{L}(A)|.$$

Thus, as the number of permutations $n \rightarrow n$ is $n!$, we get

$$\sum_{x \in \mathcal{L}(A)^n} \sum_{\tau} \mu^{[i]}(x) = \sum_{\tau} \sum_{x \in \mathcal{L}(A)^n} \mu(v_i(\vec{\tau}(x))) = n! |\mathcal{L}(A)|.$$

We now show that $\mu^{[i]}$ is n -permutation-invariant. Let $\pi: n \rightarrow n$. We show that $\mu^{[i]} \circ \vec{\pi} = \mu^{[i]}$. For any $x \in \mathcal{L}(A)^n$, we have

$$(n! |\mathcal{L}(A)|) ((\mu^{[i]} \circ \vec{\pi})(x)) = \sum_{\tau: n \rightarrow n} \mu(v_i(\vec{\tau}(\vec{\pi}(x)))) = \sum_{\sigma: n \rightarrow n} \mu(v_i(\vec{\sigma}(x))) = (n! |\mathcal{L}(A)|) (\mu^{[i]}(x)).$$

The last claim is trivial, so the proof is complete. $\square$

The following is a technical lemma.

Lemma 3.21 *Let $n \geq 2$, and let $g: \mathcal{L}(A)^{n-1} \rightarrow \mathcal{L}(A)$ be a voting rule for $n-1$ voters. If $f: \mathcal{L}(A)^n \rightarrow \mathcal{L}(A)$ is such that $f(x_1, \dots, x_{n-1}, x_n) = g(x_1, \dots, x_{n-1})$, and μ is a $n-1$ -permutation-invariant distribution on $\mathcal{L}(A)^{n-1}$, then $\mathcal{F}_{\mu^{[n]}}^n[f] \leq 2/(n |\mathcal{L}(A)|)$, and for each $i = 1, 2, \dots, n-1$ that $\mathcal{F}_{\mu^{[n]}}^i[f] \geq \frac{1}{n} \mathcal{F}_{\mu}^i[g]$.*

Proof. Let $i = 1, 2, \dots, n-1$. We show that $\mathcal{F}_{\mu^{[n]}}^i[f] \geq \frac{1}{n} \mathcal{F}_{\mu}^i[g]$. More precisely, we will show that

$$\mathcal{F}_{\mu^{[n]}}^i[f] = \frac{1}{n} \mathcal{F}_{\mu}^i[g] + \frac{1}{n |\mathcal{L}(A)|} \sum_{j=1}^{n-1} \mu(x_{-j}) 1_{g(x_1, \dots, x_{n-1})=x_i}.$$

We have

$$\mathcal{F}_{\mu^{[n]}}^i[f] = \frac{1}{n! |\mathcal{L}(A)|} \sum_{x_1, \dots, x_n} \sum_{\tau} \mu(x_{\tau(1)}, \dots, x_{\tau(n-1)}) 1_{g(x_1, \dots, x_{n-1})=x_i}.$$

For each $i = 1, 2, \dots, n$, let $V_i = \{\tau : n \rightarrow n \mid \tau(i) = n\}$. To start, note that

$$\begin{aligned} & \sum_{x_1, \dots, x_{n-1}, x_n} \sum_{\tau \in V_n} \mu(x_{\tau(1)}, \dots, x_{\tau(n-1)}) 1_{g(x_1, \dots, x_{n-1})=x_i} \\ (\text{since } \mu \text{ is } n-1\text{-perm.inv.}) &= \sum_{x_1, \dots, x_{n-1}, x_n} \sum_{\tau \in V_n} \mu(x_1, \dots, x_{n-1}) 1_{g(x_1, \dots, x_{n-1})=x_i} \\ (\text{since } |V_n| = (n-1)!) &= (n-1)! \sum_{x_1, \dots, x_{n-1}, x_n} \mu(x_1, \dots, x_{n-1}) 1_{g(x_1, \dots, x_{n-1})=x_i} \\ &= (n-1)! \sum_{x_1, \dots, x_{n-1}} \sum_{x_n} \mu(x_1, \dots, x_{n-1}) 1_{g(x_1, \dots, x_{n-1})=x_i} \\ &= (n-1)! |\mathcal{L}(A)| \sum_{x_1, \dots, x_{n-1}} \mu(x_1, \dots, x_{n-1}) 1_{g(x_1, \dots, x_{n-1})=x_i} \\ &= (n-1)! |\mathcal{L}(A)| \mathcal{F}_{\mu}^i[g]. \end{aligned} \tag{3}$$

Also, for each $i \neq n$, note that $V_i = \bigcup_{j=1}^{n-1} H_j^i$ where

$$H_j^i = \{\tau \mid \tau(i) = n \text{ and } \tau(n) = j\}.$$

It is clear that $|H_j^i| = (n-2)!$. If $\tau \in H_j^i$, then clearly $\mu(x_{\tau(1)}, \dots, x_{\tau(n-1)}) = \mu(x_{-j})$. Now fix $x_1, \dots, x_{n-1}, x_n$, and a $j \in \{1, 2, \dots, n-1\}$. Then

$$\begin{aligned} \sum_{i=1}^{n-1} \sum_{\tau \in H_j^i} \mu(x_{\tau(1)}, \dots, x_{\tau(n-1)}) 1_{g(x_1, \dots, x_{n-1})=x_i} &= \sum_{i=1}^{n-1} \sum_{\tau \in H_j^i} \mu(x_{-j}) 1_{g(x_1, \dots, x_{n-1})=x_i} \\ &= (n-1) |H_j^i| \mu(x_{-j}) 1_{g(x_1, \dots, x_{n-1})=x_i} \\ &= (n-1)(n-2)! \mu(x_{-j}) 1_{g(x_1, \dots, x_{n-1})=x_i} \\ &= (n-1)! \mu(x_{-j}) 1_{g(x_1, \dots, x_{n-1})=x_i}. \end{aligned}$$

Summing this expression over all $j = 1, 2, \dots, n-1$, we get

$$(n-1)! 1_{g(x_1, \dots, x_{n-1})=x_i} \sum_{j=1}^{n-1} \mu(x_{-j}).$$

Thus,

$$\begin{aligned} \mathcal{F}_{\mu^{[n]}}^i[f] &= \frac{1}{n! |\mathcal{L}(A)|} \sum_{x_1, \dots, x_n} \sum_{\tau} \mu(x_{\tau(1)}, \dots, x_{\tau(n-1)}) 1_{g(x_1, \dots, x_{n-1})=x_i} \\ &= \frac{1}{n! |\mathcal{L}(A)|} \sum_{x_1, \dots, x_n} \sum_{\tau \in V_n} \mu(x_{\tau(1)}, \dots, x_{\tau(n-1)}) 1_{g(x_1, \dots, x_{n-1})=x_i} \\ &\quad + \frac{1}{n! |\mathcal{L}(A)|} \sum_{x_1, \dots, x_n} \sum_{\tau \in \bigcup_{i=1}^{n-1} V_i} \mu(x_{\tau(1)}, \dots, x_{\tau(n-1)}) 1_{g(x_1, \dots, x_{n-1})=x_i}. \end{aligned}$$

Plugging in our calculations from above, we conclude that

$$\mathcal{F}_{\mu^{[n]}}^i[f] = \frac{1}{n} \mathcal{F}_{\mu}^i[g] + \frac{1}{n|\mathcal{L}(A)|} \sum_{j=1}^{n-1} \mu(x_{-j}) 1_{g(x_1, \dots, x_{n-1})=x_i}.$$

If we repeat the steps from above, but with i replaced by n , then the whole reasoning is the same, except in (3): there, we get

$$(n-1)! \sum_{x_1, \dots, x_{n-1}} \sum_{x_n} \mu(x_1, \dots, x_{n-1}) 1_{g(x_1, \dots, x_{n-1})=x_n} = (n-1)! \sum_{x_1, \dots, x_{n-1}} \mu(x_1, \dots, x_{n-1}) = (n-1)!.$$

This lets us conclude that

$$\mathcal{F}_{\mu^{[n]}}^n[f] = \frac{1}{n|\mathcal{L}(A)|} + \frac{1}{n|\mathcal{L}(A)|} \sum_{j=1}^{n-1} \mu(x_{-j}) 1_{g(x_1, \dots, x_{n-1})=x_n}.$$

Since $\sum_{j=1}^{n-1} \mu(x_{-j}) 1_{g(x_1, \dots, x_{n-1})=x_n} \leq 1$ we have $\mathcal{F}_{\mu^{[n]}}^n[f] \leq 2/(n|\mathcal{L}(A)|)$. $\square$

Let $\varepsilon > 0$ be such that $\varepsilon < 1 - 2/|\mathcal{L}(A)|$. Note that this is possible precisely because $|\mathcal{L}(A)| > 2$; this is where we use the assumption that $|A| \geq 3$, i.e., there are at least three candidates. Fix any $y \in \mathcal{L}(A)$. Let $G = \{(y, \dots, y)\}$. We define a particular distribution μ_* on $\mathcal{L}(A)^{n-1}$, as follows: μ_* gives weight $1 - \varepsilon$ to $(y, \dots, y)$ and spreads the remaining ε out over all other profiles. That is, for any $x \in \mathcal{L}(A)^{n-1}$, we let

$$\mu_*(x) = \begin{cases} 1 - \varepsilon & \text{if } x \in G, \\ \frac{\varepsilon}{|\mathcal{L}(A)|^{n-1} - 1} & \text{if } x \notin G. \end{cases}$$

Note that μ_* is $n-1$ -permutation-invariant. Also note that for a Pareto $g: \mathcal{L}(A)^{n-1} \rightarrow \mathcal{L}(A)$ we have $g(x, x, \dots, x) = x$ for all $x \in \mathcal{L}(A)$.

Lemma 3.22 *Let $g: \mathcal{L}(A)^{n-1} \rightarrow \mathcal{L}(A)$ be a voting rule for $n-1$ voters that is Pareto, and let $f: \mathcal{L}(A)^n \rightarrow \mathcal{L}(A)$ be $f(x_1, \dots, x_{n-1}, x_n) = g(x_1, \dots, x_{n-1})$. Then it holds that voter n is the unique least forceful voter of f with respect to $\mu_*^{[n]}$.*

Proof. By Lemma 3.21 it suffices to show that $\mathcal{F}_{\mu_*}^i[g] > 2/|\mathcal{L}(A)|$ for each $i = 1, 2, \dots, n-1$.

We calculate:

$$\mathcal{F}_{\mu_*}^i[g] = \sum_{(x_1, \dots, x_{n-1})} \mu_*(x_1, \dots, x_{n-1}) 1_{g(x_1, \dots, x_{n-1})=x_i} \geq \sum_{(x_1, \dots, x_{n-1}) \in G} \mu_*(x_1, \dots, x_{n-1}) 1_{g(x_1, \dots, x_{n-1})=x_i},$$

and this equals $\mu_*(y, \dots, y) = 1 - \varepsilon$ since g is Pareto. By choice of ε , we have $1 - \varepsilon > 2/|\mathcal{L}(A)|$. $\square$

Finally we arrive at the proof of Arrow's theorem.

Proof of Arrow's theorem. Assume towards a contradiction that $n \geq 2$ and there is a $g \in \text{PIIA}^{n-1} \setminus \text{DICT}^{n-1}$. We define $f: \mathcal{L}(A)^n \rightarrow \mathcal{L}(A)$ by $f(x_1, \dots, x_{n-1}, x_n) = g(x_1, \dots, x_{n-1})$. Note the following.

- (1) $f \in \text{PIIA}$ since it satisfies IIA and Pareto, as g does.
- (2) f is not a dictator: clearly none of the first $n-1$ voters can be the dictator for that would imply that g were a dictator, and if the n -th voter were the dictator then $g(x_1, \dots, x_{n-1}) = x_n$ for all $(x_1, \dots, x_{n-1}, x_n)$, in contradiction with the fact that g is a function.

Thus, $f \in \text{PIA} \setminus \text{DICT}$.

We take μ to be the $\mu_*^{[n]}$ that we just introduced. By construction, $\mu_*^{[n]}$ has full support, satisfies $\mu_*^{[n]} = \mu_*^{[n]} \circ \tilde{\pi}$ for all permutations π (by Lemma 3.20), and voter n is the unique voter with least force on f with respect to $\mu_*^{[n]}$ (by Lemma 3.22). This proves that $\Phi^{\mu_*^{[n]}}(f) = f$. Therefore, $\Phi^{\mu_*^{[n]}}([f]) = [f]$. Proposition 3.18 and Proposition 3.19 imply that $[f] = \text{DICT}$. In particular, $f \in \text{DICT}$. Contradiction. $\square$

4 Conclusion

The main goal of this paper has been to show that Arrow's impossibility theorem can be proved using Banach's fixpoint theorem. Our approach involved coming up with an appropriate equivalence relation, and then defining a contraction on the resulting equivalence space whose unique fixpoint is the set of dictators. The concept of force of a voter, as well as thinking about voting rules as elements of a metric space based on a probability distribution, are inspired by the Boolean analysis approach to social choice, a line of research initiated by [13] and further developed by others [15].

Our proof of Arrow's theorem is different in spirit from most of the previous proofs in that it does not involve manipulations of specific profiles. The Pareto property is fundamental in our analysis and we used it ubiquitously, as our original metric space consists of all Pareto voting rules. Interestingly, however, in our proof we did not use the IIA property explicitly, we only used it by noting that it was preserved under a certain operation. This makes us wonder if it would be possible to get other impossibility results by considering other properties (that preserve the same operation).

An advantage of our perspective on Arrow's theorem is that it establishes a link between this pivotal result of mathematical economics and a concept often surfacing in that area: fixpoints. The notion of fixpoint also connects the theorem better to the area of computer science, where fixpoints are omnipresent and sometimes can lead to algorithms, although this does not seem to be the case here. A possible direction for future work is to analyze if similar results in the area, such as the Gibbard-Satterthwaite theorem [12], can be proved via a related fixpoint argument. It would also be worthwhile to study the relationship between our notion of power and notions like decisiveness or pivotalness that other proofs use.

Acknowledgments. We would like to thank Thomas Santoli for pointing out in a discussion that the permutation structure gives rise to a group action, which ultimately led to Proposition 3.13, and Ronald de Wolf for inspiring us to formalize and study the notion of force.

References

- [1] Aleksandr V. Arhangel'skij & Lev S. Pontryagin (1990): *General Topology: Basic Concepts and Constructions. Dimension Theory. I*. Springer-Verlag.
- [2] Kenneth J. Arrow (1951): *Social Choice and Individual Values*. New York.
- [3] Stefan Banach (1922): *Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales*. *Fundamenta Mathematicae* 3(1), pp. 133–181, doi:10.4064/fm-3-1-133-181.
- [4] John F. Banzhaf III (1964): *Weighted Voting Doesn't Work: A Mathematical Analysis*. *Rutgers L. Rev.* 19, p. 317.
- [5] Salvador Barbera (1980): *Pivotal Voters: A New Proof of Arrow's Theorem*. *Economics Letters* 6(1), pp. 13–16, doi:10.1016/0165-1765(80)90050-6.

- [6] Julian H. Blau (1972): *A Direct Proof of Arrow's Theorem*. *Econometrica: Journal of the Econometric Society*, pp. 61–67, doi:10.2307/1909721.
- [7] Francesca Cagliari, Barbara Di Fabio & Claudia Landi (2015): *The Natural Pseudo-distance as a Quotient Pseudo-metric, and Applications*. In: *Forum Mathematicum*, 27, De Gruyter, pp. 1729–1742.
- [8] Frank M. V. Feys (2015): *Fourier Analysis for Social Choice*. Master's thesis, Universiteit van Amsterdam, the Netherlands.
- [9] Ehud Friedgut, Gil Kalai & Assaf Naor (2002): *Boolean Functions Whose Fourier Transform is Concentrated on the First Two Levels*. *Advances in Applied Mathematics* 29(3), pp. 427–437, doi:10.1016/S0196-8858(02)00024-6.
- [10] Mark B. Garman & Morton I. Kamien (1968): *The Paradox of Voting: Probability Calculations*. *Behavioral Science* 13(4), pp. 306–316, doi:10.1002/bs.3830130405.
- [11] John Geanakoplos (2005): *Three Brief Proofs of Arrow's Impossibility Theorem*. *Economic Theory* 26(1), pp. 211–215, doi:10.1007/s00199-004-0556-7.
- [12] Allan Gibbard (1973): *Manipulation of Voting Schemes: A General Result*. *Econometrica* 41(4), pp. 587–601, doi:10.2307/1914083.
- [13] Gil Kalai (2002): *A Fourier-theoretic Perspective on the Condorcet Paradox and Arrow's Theorem*. *Advances in Applied Mathematics* 29(3), pp. 412–426, doi:10.1016/S0196-8858(02)00023-4.
- [14] Alan P. Kirman & Dieter Sondermann (1972): *Arrow's Theorem, Many Agents, and Invisible Dictators*. *Journal of Economic Theory* 5(2), pp. 267–277, doi:10.1016/0022-0531(72)90106-8.
- [15] Elchanan Mossel (2012): *A Quantitative Arrow Theorem*. *Probability Theory and Related Fields* 154(1-2), pp. 49–88, doi:10.1007/s00440-011-0362-7.
- [16] John F. Nash (1950): *Equilibrium Points in n -Person Games*. *Proceedings of the National Academy of Sciences* 36(1), pp. 48–49, doi:10.1073/pnas.36.1.48.
- [17] Ryan O'Donnell (2014): *Analysis of Boolean Functions*. Cambridge University Press, doi:10.1017/CBO9781139814782.
- [18] Vittorino Pata (2014): *Fixed Point Theorems and Applications*. Politecnico di Milano.
- [19] Philip J. Reny (2001): *Arrow's Theorem and the Gibbard-Satterthwaite Theorem: A Unified Approach*. *Economics Letters* 70(1), pp. 99–105, doi:10.1016/S0165-1765(00)00332-3.
- [20] Joseph J. Rotman (2012): *An Introduction to the Theory of Groups*. 148, Springer Science & Business Media.