Arrow's Paradox

Prerna Nadathur

January 1, 2010

Abstract

In this paper, we examine the problem of a "ranked" voting system and introduce Kenneth Arrow's impossibility theorem (1951). We provide a proof sketch which suggests the method of Arrow's original proof, and a formal proof which uses finitely-additive 0-1 measures and ultrafilters.

Contents

1 Introduction

1.1 A Voter's Paradox

It might reasonably be said that the problem of establishing a modern utopia is the problem of establishing a totally equitable, all-inclusive system for making social and political decisions; that is, the problem of establishing a voting system. The problem does not seem particularly vexing when we examine it in the context of a simple situation: suppose, for example, that the decision to be made is whether Candidate A or B is to be president. The solution is quite simple: we ask each member of the population to state which candidate they prefer, and from the votes cast we then determine whether more of the population prefers A or B.[1]

Suppose, however, that we consider a situation in which it is important to rank alternatives; among options A, B, and C, Citizen 1 would like to select A, but realizes that his choice is unlikely to be the popular one, and would strongly prefer B to C if he cannot have A. At first, this seems easy to resolve; we simply alter our voting system to take a preference order as a vote. 1 hands in a vote which states A > B > C. From a collection of such rankings, it seems reasonable to expect that we will be able to extract a "national" preference order which will determine the winner. This is in fact not the case, and we only need two more citizens, 2 and 3, in order to demonstrate a familiar voting paradox.[2]

Let Citizen 1 vote as above: A > B > C, Citizen 2 rank B > C > A, and Citizen 3 rank C > A > B. Citizens 1 and 3, we see, both prefer A to B, and Citizens 1 and 2 both prefer B to C; thus, a majority of our country prefers A to B and B to C, which would lead us to the national ranking A > B > C. But Citizens 2 and 3 both prefer C to A, and so our national ranking does not reflect the preferences of the majority. There is no way out of this; it appears that our country is not behaving rationally.

The simple paradox above illustrates definitively that the passage from a collection of individual rankings to a national (or global) preference order is not without complication. Simply put, voting is a problem.

1.2 The Problem, Informally

The problem can be stated informally as follows:

Assume we want to extract a "global" preference order on a given set of (three or more) options from the individual preference orders of a finite collection of individuals. Let each individual determine his or her rankings independently: what we want, then, is a "social choice function" which takes these rankings as inputs and outputs a single preference order.

If this is to be the social choice function by which decisions are made in Utopia, however, it cannot be arbitrary. It must respect modern political sensibilities about equity and inclusiveness. In particular, we would like a social choice function to obey the following constraints:

1. **Nondictatorship**. A social choice function is incompatible with a dictatorship; that is, the function cannot simply select the ranking of one particular voter as the national preference order, but must account instead for the preferences of all.

2. **Universality.** A social choice function should account for all preferences among all voters to yield a complete, unique ranking of social choices; that is, it should account for

[1] It must be noted, of course, that in the United States we have managed to complicate even this situation with the electoral college, as a result of which it is possible for A to be selected by the popular vote and yet fail to win the election.

[2] This is due to the Marquis de Condorcet.

all individual preferences, result in a complete ranking, and provide the same ranking every time voter's preferences are presented the same way.

3. **Independence of Irrelevant Alternatives.** A social choice function should provide the same (relative) ranking among a subset of options as it would in the complete system. Changes in the rankings of the irrelevant elements should not alter the relative rankings within the subset.

4. **Pareto Optimality/Unanimity.** If all "voters" rank option a before option b, the social choice function should provide a national ranking that has a ranked before b.

In his 1951 Ph.D. thesis, economist Kenneth Arrow put forward the problem of finding such a social choice function and proved, moreover, that there is no such function. That is, given a finite set of inputs, it is impossible to create a function that will respect all the constraints stated above and yield a rational, non-contradictory global preference order. Any such preference order we manage to extract must necessarily violate one of the conditions 1-4 above.

2 The Road to Formality

We provide some mathematical background before proceeding with a formal mathematical statement of Arrow's Paradox (also known as Arrow's Impossibility Theorem) and its proof.

2.1 Measuring

Definition 2.1. A *measure* on a set S is an injective function μ from the power set $\mathcal{P}(S)$ (the set of all subsets) of S to the real numbers, which obeys the following constraints:

1. Non-negativity: $\forall A \subseteq S, \mu(A) \geq 0$

2. Null empty set: $\mu(\emptyset) = 0$

3. Countable additivity: Let $\{A_n\}, n \in \mathbb{N}$ be a countable collection of pairwise disjoint subsets of S. Then

$$\mu(\bigcup_{n\in\mathbb{N}} A_n) = \sum_{n\in\mathbb{N}} \mu(A_n) \tag{1}$$

Measures, in fact, have their own problems, and it turns out that, in order to define a measure as above, we have to accept the existence of non-measurable sets. The proof of that fact is non-trivial, and moreover, of no particular relevance to the problem at hand, and so we will omit it here. We introduce measures in general, indeed, merely to provide context for the particular type of measure in which we are interested.

Definition 2.2. A *finitely-additive 0-1 measure* over a set A is a function $\mu : \mathcal{P}(A) \to \{0, 1\}$ such that

1. $\forall B \subseteq A, \mu(B) \in \{0, 1\}$
2. $\forall B_1, B_2 \subseteq A, B_1 \cap B_2 = \emptyset \Rightarrow \mu(B_1 \cup B_2) = \mu(B_1) + \mu(B_2)$
3. $\mu(\emptyset) = 0$
4. $\mu(A) = 1$

Note that the second condition above implies not just pairwise additivity of disjoint sets, but in fact the additivity of any finite collection of pairwise disjoint set. Moreover, if $B \subseteq A$ with $\mu(B) = 1$, then $A \setminus B$ (the complement of B in A) must have $\mu(A \setminus B) = 0$ since $B \cap (A \setminus B) = \emptyset$ and $\mu(A) = \mu(B \cup (A \setminus B)) = \mu(B) + \mu(A \setminus B) = 1$. Similarly, if $\mu(B) = 0$, then $\mu(A \setminus B) = 1$. That is, from each set and its complement, one must have measure 1 and the other 0.

Definition 2.3. We call μ a *principal measure* if there exists $x \in A$ such that for all $B \subseteq A, \mu(B) = 1$ if and only if $x \in B$.

A principal measure, of course, is a dictatorship, and we may further be alarmed by the following fact:

Proposition 2.4. *If A is finite, and μ is a finitely additive 0-1 measure on A, then μ is a principal measure.*

Proof. Let μ be a finitely-additive 0-1 measure on a finite set A, where A has n elements, $n \in \mathbb{N}$.

Since A is finite, it is equal to the disjoint union of its singleton subsets, $\{x_1\}, \{x_2\}, \{x_3\}, \ldots, \{x_n\}$, where $x_i \in A$ for $1 \leq i \leq n$ and the x_i are unique. By finite additivity, we know that

$$1 = \mu(A) = \mu(\{x_1\}) + \mu(\{x_2\}) + \mu(\{x_3\}) + \ldots + \mu(\{x_n\}) \tag{2}$$

By definition, $\mu(\{x_i\})$ is either 0 or 1 for each i. Since $\Sigma_{1 \leq i \leq n} \mu(\{x_i\}) = 1$, it is clear that there is a unique $j \in \{1, 2, 3, \ldots, n\}$ such that $\mu(\{x_j\}) = 1$

Let B be a subset of A. Then B is a disjoint union of some of the $\{x_i\}$. Thus, by finite additivity, $\mu(B)$ is the sum of $\mu(\{x_i\})$ over the relevant i. Thus, $\mu(B) = 1$ if and only if $x_j \in B$. □

We can, however, avoid principal measures in the infinite case, by adding the following constraint to our definition of finitely additive 0-1 measures:

5. $\forall a \in A, \mu(\{x\}) = 0$. That is, all singleton sets have measure 0.

In fact, this implies that every finite subset has measure 0, since, by finite additivity, the measure of a finite set is the sum of the measures of its singleton subsets.

We will also need, in passing, the following infamous result:

Theorem 2.5. ***Zorn's Lemma.*** *Let P be a partially ordered set. If every chain (totally ordered subset) of P has an upper bound, then P has a maximal element.*

This is equivalent to the Axiom of Choice, so we omit proof here.

Definition 2.6. Let A be a set. Let $\mathcal{F}$ be a subset of $\mathcal{P}(A)$; that is, $\mathcal{F}$ is a set of some subsets of A. $\mathcal{F}$ is said to have the *finite intersection property* if the intersection of any finite number of elements of $\mathcal{F}$ is nonempty.

The finite intersection property gives us a neat way of constructing finitely additive 0-1 measures.

Proposition 2.7. *Let A be a set. If $\mathcal{F} \subseteq \mathcal{P}(A)$ has the finite intersection property, then there exists a finitely additive 0-1 measure μ on A such that every set in $\mathcal{F}$ has measure 1.*

Proof. Let Φ be the set of all subsets of $\mathcal{P}(A)$ with the finite intersection property. (We will refer to such subsets as f.i.p. subsets.) Then $\mathcal{F} \in \Phi$. Let $\Phi' \subseteq \Phi$ be the set of sets in Φ which contain $\mathcal{F}$. Φ' is a partially ordered set under the inclusion relation.

Every chain in Φ' is bounded: consider the union of all the elements of the chain is an upper bound for the chain. So, by Zorn's Lemma, Φ' has a maximal element; that is, there is a maximal f.i.p. subset of $\mathcal{P}(A)$ containing $\mathcal{F}$. Call this maximal element $\mathcal{F}_{\max}$. We define μ as follows:

$$\mu(B) = \begin{cases} 1 & \text{if } B \in \mathcal{F}_{\max} \\ 0 & \text{otherwise} \end{cases}$$

Since $\mathcal{F} \subseteq \mathcal{F}_{\max}$, it is evident that $\mu(F) = 1 \; \forall F \in \mathcal{F}$.

Claim 2.8. *μ is a finitely additive 0-1 measure.*

By definition, $\mu : \mathcal{P}(A) \to \{0, 1\}$. We show first that $\mu(\emptyset) = 0$:

Suppose not. Then $\mu(\emptyset) = 1$, which implies that $\emptyset \in \mathcal{F}_{\max}$. Consider $F \cap \emptyset$ for any $F \in \mathcal{F}_{\max}$. $F \cap \emptyset = \emptyset$ since, for all sets F, $\emptyset \subseteq F$. But then $F \cap \emptyset$ is a finite intersection of elements of $\mathcal{F}_{\max}$ that is empty, contradicting the fact that $\mathcal{F}_{\max}$ has the finite intersection property. Hence $\emptyset \notin \mathcal{F}_{\max}$, and thus $\mu(\emptyset) = 0$.

Next, $\mu(A) = 1$. Suppose not. Then $\mu(A) = 0$, so $A \notin \mathcal{F}_{\max}$. However, $\forall S \in \mathcal{P}(A)$, $A \cap S \neq \emptyset$, so $\mathcal{F}_{\max} \cup \{A\}$ has the finite intersection propety. Moreover, it contains $\mathcal{F}_{\max}$, contradicting the fact that $\mathcal{F}_{\max}$ is a *maximal* f.i.p. set. Thus $A \in \mathcal{F}_{\max}$ and so $\mu(A) = 1$.

To see that μ is finitely additive, consider $B_1, B_2 \subseteq A$ such that $B_1 \cap B_2 = \emptyset$. It is clear that at most one of B_1, B_2 can be in $\mathcal{F}_{\max}$, since $\mathcal{F}_{\max}$ has the finite intersection property.

Case 1. Assume $B_1 \in \mathcal{F}_{\max}$. Then $B_2 \notin \mathcal{F}_{\max}$, so $\mu(B_1) = 1$ and $\mu(B_2) = 0$.

We claim that $B_1 \cup B_2 \in \mathcal{F}_{\max}$. Suppose not. $B_1 \subseteq B_1 \cup B_2$, so $B_1 \cap F \subseteq (B_1 \cup B_2) \cap F$ for all $F \in \mathcal{F}_{\max}$. In particular, since $B_1 \cap F \neq \emptyset$ for all $F \in \mathcal{F}_{\max}$, we know that $(B_1 \cup B_2) \cap F \neq \emptyset$ for all $F \in \mathcal{F}_{\max}$. Then $\mathcal{F}_{\max} \cup \{B_1 \cup B_2\}$ has the finite intersection property and contains $\mathcal{F}_{\max}$, contradicting maximality of $\mathcal{F}_{\max}$. Hence $B_1 \cup B_2 \in \mathcal{F}_{\max}$.

Then $\mu(B_1) + \mu(B_2) = 1 + 0 = 1 = \mu(B_1 \cup B_2)$.

Case 2. Assume $B_1, B_2 \notin \mathcal{F}_{\max}$. Then $\mu(B_1) = \mu(B_2) = 0$.

We claim that $B_1 \cup B_2 \notin \mathcal{F}_{\max}$. To show this, we need the following lemma:

Lemma 2.9. $F_1, F_2 \in \mathcal{F}_{max} \Rightarrow F_1 \cap F_2 \in \mathcal{F}_{max}$.

Proof. Suppose not. Then it must be the case that $\{F_1 \cap F_2\} \cup \mathcal{F}_{\max}$ does not have the finite intersection property, or we would contradict maximality of $\mathcal{F}_{\max}$. Hence $\exists F_3 \in \mathcal{F}_{\max}$ such that $F_3 \cap (F_1 \cap F_2) = \emptyset$. But $F_1, F_2, F_3 \in \mathcal{F}_{\max}$, so they cannot have empty intersection. This is a contradiction; thus $F_1 \cap F_2 \in \mathcal{F}_{\max}$. □

Since (from above) $B_1, B_2 \notin \mathcal{F}_{\max}$, $\exists F_1, F_2 \in \mathcal{F}_{\max}$ such that $B_1 \cap F_1 = \emptyset$ and $B_2 \cap F_2 = \emptyset$. But then $B_1 \cap (F_1 \cap F_2) = \emptyset$ and $B_2 \cap (F_1 \cap F_2) = \emptyset$, so $(B_1 \cup B_2) \cap (F_1 \cap F_2) = \emptyset$. By the lemma, $F_1 \cap F_2 \in \mathcal{F}_{\max}$, so $\mathcal{F}_{\max}$ contains a set which has empty intersection with $B_1 \cup B_2$. Hence $B_1 \cup B_2 \notin \mathcal{F}_{\max}$. Thus $\mu(B_1) + \mu(B_2) = 0 + 0 = 0 = \mu(B_1 \cup B_2)$.

This shows that μ is finitely additive, and thus a finitely-additive 0-1 measure on A. □

Corollary 2.10. *For all $B \subseteq A$, exactly one of B and $A \setminus B$ must be in $\mathcal{F}_{max}$.*

Proof. Evidently, both B and its complement $A \setminus B$ cannot be in $\mathcal{F}_{\max}$, since $B \cap (A \setminus B) = \emptyset$.

Suppose neither B nor $A \setminus B$ is in $\mathcal{F}_{\max}$. Then, by maximality of $\mathcal{F}_{\max}$, $\exists F_1, F_2 \in \mathcal{F}_{\max}$ such that $F_1 \cap B = \emptyset$ and $F_2 \cap (A \setminus B = \emptyset$. As in the proof above, we have that $(B \cup (A \setminus B)) \cap (F_1 \cap F_2) = \emptyset$. But $B \cup (A \setminus B) = A$, so $A \cap (F_1 \cap F_2) = \emptyset$. Since $F_1, F_2 \in \mathcal{F}_{\max} \Rightarrow F_1, F_2 \in \mathcal{P}(A) \Rightarrow F_1, F_2 \subseteq A$, we have that either $F_1 = \emptyset, F_2 = \emptyset$ or $F_1 \cap F_2 = \emptyset$. But $F_1, F_2, F_1 \cap F_2 \in \mathcal{F}_{\max}$, and, from above, we know that $\emptyset \notin \mathcal{F}_{\max}$. Hence at least one of $B, A \setminus B$ must be in $\mathcal{F}_{\max}$.

Since both B and its complement cannot be in $\mathcal{F}_{\max}$, and at least one of them must be, we conclude that exactly one of B and its complement is in $\mathcal{F}_{\max}$. □

2.2 Filtering

We now introduce some more tools for dealing with sets, subsets, and sets of subsets.

Definition 2.11. Let A be a set. Let $\mathcal{F} \subseteq \mathcal{P}(A)$. $\mathcal{F}$ is called a *filter* if it is

1. upward closed: $\forall C \in \mathcal{P}(A), \forall B \in \mathcal{F}, B \subseteq C \Rightarrow C \in \mathcal{F}$.
2. closed under finite intersections: $\forall B_1, B_2 \in \mathcal{F}, B_1 \cap B_2 \in \mathcal{F}$.

Note that, as a consequence of the first condition, any nontrivial filter $\mathcal{F}$ (i.e. neither the empty filter nor all of $\mathcal{P}(A)$) will satisfy the conditions i) $A \in \mathcal{F}$ and ii) $\emptyset \notin \mathcal{F}$.

Definition 2.12. $\mathcal{F}$ is an *ultrafilter* if $\mathcal{F}$ is a filter with the property that, for all $B \subseteq A$, exactly one of B and $A \setminus B$ is in $\mathcal{F}$.

Recall that the measure we constructed using the finite intersection property satisfied the property that, from each set and its complement, exactly one was in $\mathcal{F}_{\max}$. This suggests that $\mathcal{F}_{\max}$ was actually an ultrafilter, and leads us to the following conjecture.

Conjecture 2.13. *Let μ be a finitely-additive 0-1 measure on a set A. Then $\mu^{-1}(1)$ is an ultrafilter.*

Proof. Consider $\mu^{-1}(1)$. This is the set of elements of $\mathcal{P}(A)$ that have measure 1 under the measure μ. Lemma 2.6 showed that if μ is a finitely-additive 0-1 measure, then from every set and its complement, one has measure 1 and the other measure 0. Hence, $\forall B \subseteq A$, exactly one of B and $A \setminus B$ must be in $\mu^{-1}(1)$. It remains to show that $\mu^{-1}(1)$ is a filter.

For upward closure, let $B \subseteq A$ such that $B \in \mu^{-1}(1)$, and let $C \subseteq A$ such that $B \subseteq C$. $B \cap (C \setminus B) = \emptyset$, and since $C = B \cup (C \setminus B)$, $\mu(C) = \mu(B \cup (C \setminus B)) = \mu(B) + \mu(C \setminus B) = 1 + \mu(C \setminus B)$.

Since $\mu(C)$ can be at most 1 and $\mu(C \setminus B)$ cannot be negative, $\mu(C \setminus B) = 0$ and C must have measure 1. Thus, $C \in \mu^{-1}(1)$, and $\mu^{-1}(1)$ is upward closed.

To show that $\mu^{-1}(1)$ is closed under finite intersections, let $B_1, B_2 \subseteq A$ such that $B_1, B_2 \in \mu^{-1}(1)$. Then $\mu(B_1) = \mu(B_2) = 1$.

$B_1 \subseteq B_1 \cup B_2$; hence, by upward closure, $B_1 \cup B_2$ has measure 1. We want to show that $B_1 \cap B_2$ has measure 1.

Assume not. Then $\mu(B_1 \cap B_2) = 0$. Consider $B_1 \setminus (B_1 \cap B_2)$, $B_2 \setminus (B_1 \cap B_2)$, and $B_1 \cap B_2$. These are all pairwise disjoint. In particular, $B_1 = (B_1 \setminus (B_1 \cap B_2)) \cup (B_1 \cap B_2)$. Thus,

$$\begin{aligned}
1 &= \mu(B_1) \\
&= \mu((B_1 \setminus (B_1 \cap B_2)) \cup (B_1 \cap B_2)) \\
&= \mu((B_1 \setminus (B_1 \cap B_2))) + \mu(B_1 \cap B_2) \\
&= \mu(B_1 \setminus (B_1 \cap B_2)) + 0 \\
&= \mu(B_1 \setminus (B_1 \cap B_2))
\end{aligned}$$

Similarly, $\mu(B_2 \setminus (B_1 \cap B_2)) = 1$.

$(B_1 \setminus (B_1 \cap B_2)) \cup (B_2 \setminus (B_1 \cap B_2)) \subset B_1 \cup B_2$, hence $\mu((B_1 \setminus (B_1 \cap B_2)) \cup (B_2 \setminus (B_1 \cap B_2))) \leq \mu(B_1 \cup B_2)$.

$(B_1 \setminus (B_1 \cap B_2)) \cap (B_2 \setminus (B_1 \cap B_2)) = \emptyset$, so $\mu((B_1 \setminus (B_1 \cap B_2)) \cup (B_2 \setminus (B_1 \cap B_2))) = \mu(B_1 \setminus (B_1 \cap B_2)) + \mu(B_2 \setminus (B_1 \cap B_2)) = 1 + 1 = 2$. But $\mu(B_1 \cup B_2) = 1$ and $1 < 2$, so we have a contradiction. Hence $\mu(B_1 \cap B_2) \neq 0$, and so $\mu(B_1 \cap B_2) = 1$, and $\mu^{-1}(1)$is closed under finite intersection.

Thus $\mu^{-1}(1)$ is an ultrafilter. □

In fact, there is a one-to-one correspondence between ultrafilters and finitely-additive 0-1 measures. We have shown that any finitely-additive 0-1 measure μ on a set A gives $\mu^{-1}(1)$ an ultrafilter on A. It remains only to be shown, then, that an ultrafilter $\mathcal{U}$ gives a finitely-additive 0-1 measure μ.

We use the following as a lemma:

Theorem 2.14. *For all $\mathcal{F} \subseteq \mathcal{P}(A)$ with the finite intersection property, there exists an ultrafilter $\mathcal{U}$ such that $\mathcal{F} \subseteq \mathcal{U}$.*

Proof. Recall the construction of $\mathcal{F}_{\text{max}}$ in Proposition 2.10. We claim that $\mathcal{F}_{\text{max}}$ is an ultrafilter.

For upward closure, let $B \in \mathcal{F}_{\text{max}}$. Let $C \subseteq A$ such that $B \subseteq C$. Then $\mathcal{F}_{\text{max}} \cup \{C\}$ has the finite intersection property, and so by maximality of $\mathcal{F}_{\text{max}}$, $C \in \mathcal{F}_{\text{max}}$.

We showed that $\mathcal{F}_{\text{max}}$ is closed under finite intersections in Lemma 2.12. So $\mathcal{F}_{\text{max}}$ is a filter.

Then Corollary 2.13 shows that $\mathcal{F}_{\text{max}}$ is an ultrafilter.

Hence, for all $\mathcal{F} \subseteq \mathcal{P}(A)$ with the finite intersection property, there exists an ultrafilter $\mathcal{U} = \mathcal{F}_{\text{max}}$ such that $\mathcal{F} \subseteq \mathcal{U}$. □

Recall that, using $\mathcal{F} \subseteq \mathcal{U}$, we defined a function:

$$\mu(B) = \begin{cases} 1 & \text{if } B \in \mathcal{F}_{\text{max}} \\ 0 & \text{otherwise} \end{cases}$$

μ is a finitely-additive 0-1 measure; in particular, $\mu(F) = 1$ for all $F \in \mathcal{F}$.

In particular, then, since all ultrafilters are also filters, we can construct a finitely-additive 0-1 measure $\mu_{\mathcal{U}}$ corresponding to any ultrafilter $\mathcal{U}$ such that $\mu_{\mathcal{U}}(U) = 1$ for all $U \in \mathcal{U}$.

Thus, we have a one-to-one correspondence between finitely-additive 0-1 measures and ultrafilters on a set A.

Now let A be infinite. Consider the set of subsets of A which are the complements of the singleton subsets of A; that is, the set of all subsets of A which contain all but one element of A. Call this $\mathcal{F}$.

Claim 2.15. *$\mathcal{F}$ has the finite intersection property.*

Proof. Consider an arbitrary finite collection $\{F_1, F_2, \ldots, F_n\}$ of elements of $\mathcal{F}$, where $F_1 = A \setminus \{a_1\}, F_2 = A \setminus \{a_2\}, \ldots, F_n = A \setminus \{a_n\}$, and $a_i \in A \; \forall i \in \{1, \ldots, n\}$.

A contains infinitely many elements. In particular A contains an element a_{n+1} such that $a_{n+1} \neq a_i \; \forall i \in \{1, \ldots, n\}$.

Each F_i contains every element of A except a_i, respectively. Hence, $a_{n+1} \in F_i \; \forall i \in \{1, \ldots, n\}$ since $a_{n+1} \neq a_i$. In particular, then, $a_{n+1} \in \cap_{i \in \{1,\ldots,n\}} F_i$ and so $\cap_{i \in \{1,\ldots,n\}} F_i \neq \emptyset$. The F_i comprise an arbitrary finite collection of elements of $\mathcal{F}$; hence $\mathcal{F}$ has the finite intersection property. □

This particular f.i.p. subset of an infinite set A, in fact, allows us to construct a nonprincipal finitely-additive 0-1 measure on A:

Theorem 2.16. *Let A be an infinite set. Then there exists a nonprincipal finitely-additive 0-1 measure μ on A. Equivalently, there exists a nonprincipal ultrafilter $\mathcal{U}$ on A; in particular, $\mathcal{F} \subseteq \mathcal{U}$ for the f.i.p subset $\mathcal{F}$ as described above.*

Proof. By theorem 2.17, we know that there exists an ultrafilter $\mathcal{U}$ on A such that $\mathcal{F} \subseteq \mathcal{U}$; in particular, $\mathcal{U}$ is the maximal f.i.p. subset of $\mathcal{P}(A)$ which contains $\mathcal{F}$. Thus, we know there exists a finitely-additive 0-1 measure measure μ on A such that

$$\mu(B) = \begin{cases} 1 & \text{if } B \in \mathcal{U} \\ 0 & \text{otherwise} \end{cases}$$

We claim that this μ is nonprincipal: that is, we claim that $\mathcal{U}$ does not contain any singleton subsets of A.

Suppose not. Then, for some $a \in A$, $\{a\} \in \mathcal{U}$. $\mathcal{F} \subseteq \mathcal{U}$ by construction. By definition, $A \setminus \{a\} \in \mathcal{F}$, so $A \setminus \{a\} \in \mathcal{U}$. It is clear that $\{a\} \cap (A \setminus \{a\}) = \emptyset$; however, $\mathcal{U}$ has the finite intersection property, so this is a contradiction. Hence there cannot be any $a \in A$ such that $\{a\} \in \mathcal{U}$, and thus, for all $a \in A$, $\mu(\{a\}) = 0$.

Thus, μ is a nonprincipal finitely-additive 0-1 measure on A. □

3 Dystopia

3.1 The Problem, Formally

We now have the mathematical tools we need to prove Arrow's Impossibility Theorem. We begin by stating the theorem formally.

Let C be a set of options, and n the number of voters. The set of all full linear orderings of C is $\mathcal{S}_{|C|}$, or the symmetry group of C. A social choice function will be a function $F : \mathcal{S}_{|C|}^n \to \mathcal{S}_{|C|}$ that takes as its input the "votes" (i.e. n elements of $\mathcal{S}_{|C|}$) and outputs a single preference order (element of $\mathcal{S}_{|C|}$).

We require F, further, to satisfy the following three conditions:

1. **Pareto Optimality.** Let $a, b \in C$. We will write $a > b$ to indicate that a is preferred to b. Let the input set of F be $(v_1, \ldots v_n)$. If $v_1, \ldots, v_n$ all have $a > b$, then $F(v_1, \ldots, v_n)$ must have $a > b$.

2. **Nondictatorship.** There is no $i \in \{1, \ldots, n\}$ such that it is always true that $v_i = F(v_1, \ldots, v_n)$.

3. **Independence of Irrelevant Alternatives (IIA).** Given two sets of inputs $(v_1, \ldots, v_n)$ and $(u_1, \ldots u_n)$, and $a, b \in A$, if $a < b$ in v_i and u_i for all i, then $a < b$ in $F(v_1, \ldots, v_n)$ and $F(u_1, \ldots, u_n)$.

Theorem 3.1. ***Arrow's Impossibility Theorem.*** *If $|C| \geq 3$ and $n < \omega$, then there does not exist a social choice function F which satisfies the above three conditions.*

We will show, in particular, that a social choice function F satisfying Pareto optimality and IIA necessitates a dictatorship.

3.2 A Proof, Informally

For clarity's sake, we state first an informal proof that does not necessitate the use of the mathematical tools we introduced in section 2.

Proposition 3.2. *Let $C = \{a, b, c\}$, and n be the (finite) number of voters. Let F be a social choice function satisfying Pareto optimality and IIA. Then F has a dictator.*

Claim 3.3. *There exists a "pivotal" voter for b.*

Proof. Let $(v_1, \ldots, v_n)$ be an input set in which v_i ranks b lowest among the choices for all $i \in \{1, \ldots, n\}$. By Pareto optimality, then, $F(v_1, \ldots, v_n)$ ranks b the lowest. Now let $(u_1, \ldots, u_n)$ be an input set in which u_i ranks b highest among the choices for all $i \in \{1, \ldots, n\}$. Again by Pareto optimality, $F(v_1, \ldots, v_n)$ ranks b highest.

Now consider the following rankings:

$$\begin{aligned} F_0 &= F(v_1, v_2, \ldots, v_n), \\ F_1 &= F(u_1, v_2, v_3, \ldots, v_n), \\ F_2 &= F(u_1, u_2, v_3, \ldots, v_n), \ldots, \\ F_i &= F(u_1, \ldots, u_i, v_{i+1}, \ldots, v_n), \ldots, \\ F_n &= F(u_1, u_2, \ldots, u_n). \end{aligned}$$

There is clearly some "pivotal" (i.e. first) number j in $\{1, \ldots, n\}$ where F_j moves b off of the bottom.

Indeed, we claim that for this number j, when voter j moves b from lowest ranked to highest ranked, society moves b from lowest to highest as well, and not to an intermediate point.

Assume that this is not the case: Then, after j has moved b to the top (i.e. when voters 1 through j have b at the top and voters $j + 1$ through n have b at the bottom), society would have some option preferred to b (this can without loss of generality be assumed to be a), and one less preferred than b (again, WLOG, c).

In this situation, if each voter moves his preference for c *directly* above a, then society ranks $c > a$ by Pareto optimality. By IIA, the relative change in the rankings of a and c should not affect the relative ranking of a and b. Moreover, since a is preferred to b already, the new "societal" preference must have $c > a > b$. But now our global preference has gone from having $b > c$ to having $c > b$ *without* any voter changing their relative ranking of b and c. This violates IIA; hence, b could never have been in the intermediate position. Voter j, then, is the pivotal voter for b in the sense that it is when voter j changes his ranking of b from lowest to highest that society moves b from lowest ranked to highest ranked. □

Claim 3.4. *Voter j, in fact, is the dictator between a and c.*

Proof. Consider the set of inputs where voters 1 through $j - 1$ rank b the highest, voter j ranks $a > b > c$, and everyone else ranks b lowest. As far as the relative rankings of a and b go, this set of inputs is the same as the set $(u_1, u_2, \ldots, u_{j-1}, v_j, v_{j+1}, \ldots, v_n)$, and we know from above that $F(u_1, u_2, \ldots, u_{j-1}, v_j, v_{j+1}, \ldots, v_n)$ ranks $a > b$ (since in fact we know that $F(u_1, u_2, \ldots, u_{j-1}, v_j, v_{j+1}, \ldots, v_n)$ has b ranked lowest). Similarly, as far at the relative rankings of b and c are concerned, this set of inputs is the same as the set $(u_1, u_2, \ldots, u_j, v_{j+1}, \ldots, v_n)$, and we know from above that $F(u_1, u_2, \ldots, u_j, v_{j+1}, \ldots, v_n)$ ranks $b > c$ (since in fact it has b ranked highest). Thus this particular set of inputs results in the global ordering $a > b > c$. If voter j had instead ranked $c > b > a$, we would be forced to conclude, by a similar argument, that the global ordering was $c > b > a$. Thus voter j dictates the relative global ranking of a and c. □

Claim 3.5. *There exists at most one dictator.*

Proof. By arguments analogous to those in Claims 3.2 and 3.3, we can show that there exist dictators over the $(a, b), (a, c)$, and (b, c) pairs; let these dictators be voters i, j, and k, respectively. Suppose voter i ranks $a > b > c$, voter j ranks $c > b > a$, and voter k ranks $b > c > a$. Then, since i dictates (a, b), we must have $a > b$, and since k dictates (b, c), we must have $b > c$. This forces a global preference order of $a > b > c$. However, j is the dictator for (a, c) and mandates $c > a$; this is a contradiction. Hence we cannot have independent dictators for each of the three pairings. Since we know from Claims 3.2 and 3.3 that there does indeed exist a dictator for each pairing, we are forced to assume that two of these dictators are the same person. Without loss of generality, we may assume that the dictator for the (a, b) and (b, c) pairings is the same person (let it be voter i). Continue to assume that voter j dictates the (a, c) ranking. Suppose voters i and j rank as before (i ranks $a > b > c$ and j ranks $c > b > a$). We must, as before, have $a > b$ and $b > c$ which mandates a global ranking of $a > b > c$, as before. However, voter j mandates the ranking of $c > a$, which is a contradiction. Thus, we see that we cannot even have two dictators; the dictators for each of the three pairings must in fact be the same person. There can exist only one dictator for the entire ranking. □

We chose the number n of voters at random; hence, to generalize this proof, we need only to generalize the number of options to some arbitrary number m, where $m \in \mathbb{N}$. It is perhaps not particularly difficult to imagine how one might generalize the above argument in this way; we need merely show the existence of a dictator for each pair of options, and then an extension of the argument in Claim 3.4 will reduce the number of possible dictators to one.[3] However, we can provide a more elegant proof using the mathematical tools developed previously, and will do so in the following section.

3.3 The Proof, Formally

Hopefully, section 2 has illuminated the method for proof. In this case, we aim to show that if a social choice function F exists which respects Pareto optimality and IIA, then it must have a dictator. This would initially suggest that we want to show that F can give us a finitely-additive 0-1 measure on the set A of voters; however, this seems a bit awkward to

[3] Arrow's original proof of the theorem developed along these lines.

manage, given that F is a function from $S^n_{|C|} \to S_{|C|}$ and hence clearly is not such a measure (although, by the result of the theorem, it must somehow induce one).

We have shown previously that there is a one-to-one correspondence between finitely-additive 0-1 measures and ultrafilters; ultrafilters allow us to work directly with subsets, so we use these instead.

For the sake of elegance, we here restate some notation.

Let C be the set of candidates, or options, with $|C| = m$, where $m \in \mathbb{N}$ and $m \geq 3$. Let the set of voters be A, with $|A| = n$, for some $n \in \mathbb{N}$.

Instead of defining a social choice function directly, let us first define a preference function p which assigns to each voter $i \in A$ a ranking of the elements of C. That is, $p : A \times C \to \{1, 2, \ldots, m\}$ such that $p(i, c)$ is the number at which voter i ranks candidate c. Piecewise, we have functions $p_i : C \to \{1, 2 \ldots, m\}$ such that $p_i(c) = p(i, c)$ for all $i \in A, c \in C$.

Given a preference state p, then, voter i ranks candidate a higher than candidate b if and only if $p(i, a) > p(i, b)$. Evidently, the number of possible preference states is $(m!)^n$; we will call the set of preference states P.

It is clear that p is equivalent to the vector $(v_1, \ldots, v_n)$ which we defined previously, in which v_i represented voter i's ranking of the candidates. We will find the new notation easier to work with.

The social choice function f can be defined as a function which assigns to each preference state p a ranking of candidates (society's ranking). That is, we can think of f as a function which takes as its inputs a preference state p and a candidate i and outputs a number between 1 and m to indicate the candidate's placement in society's ranking under the preference state p. Thus, $f : P \times C \to \{1, 2, \ldots, m\}$ such that for all $a, b \in C$ and all $p \in P$, $f(p, a) > f(p, b)$ indicates that the social function ranks a higher than b when the preferences of the voters are as in preference state p.

(Note that we could also think of f as a function which took only a preference state p as an input and gave as output an ordering of the elements of C. The first way will be more convenient.)

As defined, then, f satisfies Pareto optimality if and only if, for all $p \in P$ and $a, b \in C$, it is the case that, for all $i \in A, p(i, a) > p(i, b)$ implies that $f(p, a) > f(p, b)$. IIA is satisfied if and only if whenever $p, p' \in P$, and $a, b \in C$ are such that $p_i(a) > p_i(b)$ if and only if $p'_i(a) > p'_i(b)$ for all $i \in A$ and $f(p, a) > f(p, b)$, then we have $f(p', a) > f(p', b)$.

In this framework, then, a dictator is a voter k such that for all preference states p, and all candidates a, $f(p, a) = p(k, a)$.

In addition, we specify *monotonicity*. A social choice function f satisfies this condition if and only if, whenever $p, p' \in P$ and $a, b \in C$ have $f(p, a) > f(p, b)$ and $f(p', c) = f(p, c)$ for all $c \in C$ such that $c \neq a, b$, and $p(i, a) > p(i, b)$ implies that $p'(i, a) > p'(i, b)$, then we have $f(p'a) > f(p', b)$. In other words, if the social choice function for preference state p ranks candidate a higher than candidate b, and the only difference between p and p' is that some of the voters who previously had $b > a$ now switch to $a > b$, then the social choice function must rank a higher than b for the preference state p' as well. In particular, if the set of voters who ranked a higher than b was dominant enough to secure $a > b$ in society's ranking, then increasing that set of voters should not change the relative ranking of a and b. We can see that monotonicity in essence follows from Pareto optimality and IIA.

Let f be a monotonic social choice function which satisfies Pareto optimality and IIA. For any preference state p and pair of candidates a, b, let $R(p, a, b) = \{i \in A : p(i, a) > p(i, b)\}$. That is, let $R(p, a, b)$ be the set of voters who prefer a to b.

Definition 3.6. A set X is an *oligarchy for a over b* if, for all $p \in P$, $X \subseteq R(p, a, b) \Rightarrow f(p, a) > f(p, b)$. X is simply an *oligarchy* if X is an oligarchy for a over b for all pairs of candidates a, b.

Let $\mathcal{U}(a, b)$ denote the set of oligarchies for a over b, and $\mathcal{U}$ denote the set of all oligarchies. We would like to show that $\mathcal{U}$ is an ultrafilter on A.

Lemma 3.7. *For all preference states $p, p' \in P$ and all pairs of candidates $a, b \in C$*

$$f(p, a) > f(p, b) \text{ and } R(p, a, b) \subseteq R(p', a, b) \Rightarrow f(p', a) > f(p', b) \tag{3}$$

That is, if the social choice function for preference state p ranks $a > b$, and if also the set of voters who rank a higher than b in preference state p is a subset of the set of voters who rank a higher than b in preference state p', then the social choice function for p' ranks $a > b$.

Proof. Let $p, p' \in P$ and $a, b \in C$ such that

$$f(p, a) > f(p, b) \text{ and } R(p, a, b) \subseteq R(p', a, b) \tag{4}$$

Let $p^* \in P$ be the preference state which has the following, for $i \in A$ and $c \in C$ such that $c \neq a, b$:

$$p^*(i, a) = p'(i, a)$$
$$p^*(i, b) = p'(i, b)$$
$$p^*(i, c) = p(i, c)$$

Now, $p, p^* \in P$ and $a, b \in C$ have

$$f(p, a) > f(p, b),$$
$$f(p^*, c) = f(p, c)$$

for all $c \in C$ such that $c \neq a, b$, and

$$p(i, a) > p(i, b) \Rightarrow p^*(i, a) > p^*(i, b) \tag{5}$$

since $p^* = p'$ for a, b and $R(p, a, b) \subseteq R(p', a, b)$.

By monotonicity, then, we have $f(p^*, a) > f(p^*, b)$. Since $p^* = p'$ for a, b, by IIA we have $f(p', a) > f(p', b)$. $\square$

Lemma 3.8. *If $p \in P, a, b \in C$ such that $f(p,a) > f(p,b)$, and $X = R(p,a,b)$, then $X \in \mathcal{U}(a,b)$. That is, if the social choice function ranks $a > b$ for some preference state p and X is the set of voters who rank a higher than b in preference state p, then X is an oligarchy for a over b.*

Proof. This follows directly from Lemma 3.6. □

Lemma 3.9. *If X is an oligarchy for a over b, then X is an oligarchy. That is, $\mathcal{U}(a,b) = \mathcal{U}$.*

Proof. It suffices to show that, for all $a \neq b$ and $a' \neq b'$ in C, $\mathcal{U}(a,b) = \mathcal{U}(a',b')$.

Fix $a, b \in C$ and let $a' \in C$ such that $a' \neq b$. Let $X \in \mathcal{U}(a,b)$ and $X \subseteq R(p,a',b)$.

Choose a preference state p' such that $R(p',a',b) = R(p,a,b)$; that is, so that the set of voters who prefer a' to b in preference state p' is the same as the set of voters who prefer a to b in preference state p. Since $X \in \mathcal{U}(a,b)$, we have $f(p,a) > f(p,b)$, and so, for preference state p' we must have $f(p',a') > f(p',b)$. Then $X \in \mathcal{U}(a',b)$ since, for arbitrary $p^* \in P$, $X \subseteq R(p^*,a',b) \Rightarrow f(p^*,a') > f(p^*,b)$. Thus, $X \in \mathcal{U}(a,b) \Rightarrow X \in \mathcal{U}(a',b)$; that is, $\mathcal{U}(a,b) \subseteq \mathcal{U}(a',b)$.

A similar argument (reversing the roles of a, a') gives us that $\mathcal{U}(a',b) \subseteq \mathcal{U}(a,b)$, so $\mathcal{U}(a',b) = \mathcal{U}(a,b)$. The same argument will give us that $\mathcal{U}(a,b) = \mathcal{U}(a,b')$ for $b' \neq a$. The name of the first candidate is arbitrary, as long as the two candidates are not the same; hence $\mathcal{U}(a',b) = \mathcal{U}(a',b')$. Then $\mathcal{U}(a,b) = \mathcal{U}(a',b) = \mathcal{U}(a',b')$. Hence an oligarchy is independent of the pair of candidates, and if X is an oligarchy for a over b, X is an oligarchy. □

Proposition 3.10. *$\mathcal{U}$ is an ultrafilter.*

Proof. By Lemma 3.8, it suffices to show, that for some $a, b \in C, \mathcal{U}(a,b)$ is an ultrafilter.

To show upward closure, let $X \in \mathcal{U}(a,b)$ and let $Y \subseteq A$ such that $X \subseteq Y$. Let p be any preference state such that $Y \subseteq R(p,a,b)$. Then $X \subseteq R(p,a,b)$, and so $X \in \mathcal{U}(a,b) \Rightarrow f(p,a) > f(p,b)$. Hence $Y \in \mathcal{U}(a,b)$.

For finite intersections, let $X, Y \in \mathcal{U}(a,b)$. Let $a, b \in C$ and p a preference state such that $X \cap Y \subseteq R(p,a,b)$.

Since $|C| \geq 3$, $\exists c \in C$ such that $c \neq a, b$. Suppose $p' \in P$ is such that

$$\begin{aligned} p'(i,a) > p'(i,c) \forall i \in X \setminus Y, \\ p'(i,a) > p'(i,c) > p'(i,b) \forall i \in X \cap Y, \\ p'(i,c) > p'(i,b) \forall i \in Y \setminus X, \\ R(p',a,b) = R(p,a,b) \end{aligned}$$

Then $f(p',a) > f(p',c)$ since $X \in \mathcal{U}$ by Lemma 3.8.

Similarly, $f(p',c) > f(p',b)$ since $Y \in \mathcal{U}$. Hence,

$$f(p',a) > f(p',c) > f(p',b) \tag{6}$$

and so, by IIA

$$f(p,a) > f(p,b) \tag{7}$$

But then $X \cap Y \in \mathcal{U}(a,b)$, and $\mathcal{U}(a,b)$ is closed under finite intersection.

It remains to be shown that, for all $X \subseteq A$, exactly one of X and $A \setminus X$ is in $\mathcal{U}(a,b)$. Let $X \subseteq A$. Let p be any preference state such that X unanimously has $a > b$. If $f(p,a) > f(p,b)$, then $X \in \mathcal{U}(a,b)$. In particular, if $f(p,a) > f(p,b)$ whenever $X \subseteq R(p,a,b)$, then $f(p',a) > f(p',b)$ for the preference state p' in which $X = R(p',a,b)$ and $A \setminus X = R(p',b,a)$. Thus $A \setminus X \notin \mathcal{U}(b,a)$, and so, by Lemma 3.8, $A \setminus X \notin \mathcal{U}(a,b)$.

Let p remain any preference state such that X unanimously has $a > b$. Suppose $f(p,a) \not> f(p,b)$. Then $f(p,b) > f(p,a)$, and so, some subset Y of $A \setminus X$ which has $b > a$ is in $\mathcal{U}(b,a)$. But then, by upward closure, $A \setminus X \in \mathcal{U}(b,a)$, and so $A \setminus X \in \mathcal{U}(a,b)$.

Hence, $X \in \mathcal{U}(a,b)$ implies $A \setminus X \notin \mathcal{U}(a,b)$, and $X \notin \mathcal{U}(a,b)$ implies $A \setminus X \in \mathcal{U}(a,b)$. Since we can regard $X \subseteq A$ as $A \setminus (A \setminus X)$, this shows that from every set and its complement, exactly one is in $\mathcal{U}(a,b)$.

Thus $\mathcal{U}(a,b)$ is an ultrafilter, and so $\mathcal{U}$ is an ultrafilter. □

We have shown, then, that the set of subsets of A which are able to determine the election form an ultrafilter. From section 2 we know that $\mathcal{U}$ gives us a finitely-additive 0-1 measure μ on A such that $B \subseteq A$ has measure 1 if and only if $B \in \mathcal{U}$. That is, $\mu(B) = 1$ if and only if B is an oligarchy. But A is finite; hence, by Proposition 2.5, μ is a principal measure. That is, $\exists i \in A$ such that $\mu(\{i\}) = 1$ and $\mu(B) = 1$ if and only if $i \in B$. In particular, then, there is a single voter in A who forms an oligarchy (in fact, a tyranny!). Thus, our social choice function f has a dictator.

It may be worth noting, that, as we showed in section 2, if our set A of voters was infinite, there exist nonprincipal 0-1 measures on A, and so, in particular, it is possible to define a social choice function that satisfies all three of Pareto optimality, IIA, and nondictatorship. Unfortunately, it would be a very strange sort of Utopia indeed which contained an infinite number of votes (let alone voters), and so this is largely of academic, rather than practical, interest.

4 Conclusion

One of the consequences of Arrow's theorem is that we cannot in any politically (or perhaps merely democratically) reasonable fashion pass from a collection of individual rankings to a global preference order. Philosophically, this suggests that social and political choices for a community cannot be made by simply aggregating the choices of individuals; that is, in particular, it appears that we cannot consider the preferences of a (finite) community to be in some way equivalent to the sum of the preferences of its members.

This raises a number of interesting questions about the efficacy of a "majority-rules" style democracy. If the preferences of a community cannot be reduced to the preferences of the individuals who make up the community, then it seems reasonable to infer from this that what is in the individual interest of a majority of the members of a community is not necessarily what is in the interest of the community as a whole. This is, moreover, reminiscent of Rousseau's conception of the *general will*, which, at least to some extent, informs modern conceptions of governance and democracy. In order to protect our ideas about democratic decision-making (represented here, crudely, by Pareto optimality, monotonicity and the irrelevance of independent alternatives), perhaps we must conclude that a community must govern itself as a community; collectives, in particular, must make collective decisions, and not rely on the benevolence of the individual.

References

[1] Arrow, K.J., "A Difficulty in the Concept of Social Welfare." *Journal of Political Economy*, 58(4). August, 1950, pp. 328-346.

[2] Babai, Laszlo. Lecture Notes: Transfinite Combinatorics. University of Chicago. 2007, 2009.

[3] Geanakoplos, John. "Three Brief Proofs of Arrow's Impossibility Theorem." Cowles Foundation for Research in Economics, Yale University. 2001.