Finite Dictatorships and Infinite Democracies

Iian B. Smythe

October 20, 2015

Abstract

Does there exist a reasonable method of voting that when presented with three or more alternatives avoids the undue influence of "irrelevant" candidates? The famous theorem of Kenneth J. Arrow roughly says that the answer is "no", or more dramatically, "yes, but only a dictatorship". As this is a result every mathematician should know, we will present what it says, what it does not say, and its proof. We will then consider the analogous question for *infinite* collections of voters. No explicit prerequisites required, though familiarity with a lemma of Zorn may be helpful. These notes were prepared for a talk given by the author in the Olivetti Club Seminar at Cornell.

1 Voting Functions

Suppose that we are given a finite set I of *individuals* (also called an *electorate*) and a finite set A of *alternatives* (also called a *slate*). We wish to model procedures by which the members of I choose an alternative in A best reflecting their collective preferences, or *social choice*. We will restrict our attention to voting for candidates in a democratic election, in which the outcome is a set of *winners* (with possible ties), and a set of *losers*.

Denote by $\mathcal{P}_A$ the set of all linear orderings of the elements of A. We think of the elements of $\mathcal{P}_A$ as possible *preference ballots* (also called *ranked ballots*), which a voter may submit as their ballot in an election. (More generally, one could allow $\mathcal{P}_A$ to be the set of all "weak orders" on A, linear-like orderings with ties.)

The set $\mathcal{P}_A^I$ of all functions $I \to \mathcal{P}_A$ is the set of *profiles*, or choices of preference ballot for each individual. If $P \in \mathcal{P}_A^I$, and $i \in I$, we denote by $<_i^P$ the ordering in $\mathcal{P}_A$ given by $P(i)$. We denote by $\mathcal{P}(A)^+$ the set of all non-empty subsets of A.

Definition 1.1. A function $\sigma : \mathcal{P}_A^I \to \mathcal{P}(A)^+$, whose input is a profile, and output is a nonempty subset of A, is called a *voting function*. Elements of the set $\sigma(P)$ are *winners*, and elements of $A \setminus \sigma(P)$ are *losers*.

More generally, one can allow the output of a social choice function to be a weak order on A, rather than just a subset. We note that the relation $b < a$ if and only if $a \in S$ and $b \notin S$ for a nonempty set $S \subseteq A$, is itself a weak order. Even with the more restrictive definition, there are many voting functions for any given I and A. Here are a few examples:

Example 1.2. The *simple plurality method* is the voting function σ such that given a profile $P \in \mathcal{P}_A^I$, $a \in \sigma(P)$ if and only if a is ranked first by the most members of I (i.e., a is the $<_i^P$-greatest element for a plurality of $i \in I$, which makes sense since I is finite). Note that only the highest ranked entries of each preference ballot are used in this method, conforming to the usual "vote for one" ballots used in most North American elections.

Example 1.3. A *dictatorship*, with *dictator* $i_0 \in I$, is the voting function σ such that given a profile $P \in \mathcal{P}_A^I$, $a \in \sigma(P)$ if and only if a is ranked first by i_0 (i.e., a is the $<_{i_0}^P$-greatest element).

What makes a voting method "reasonable" for elections in a democracy? The following criteria, at the very least, would seem necessary:

Definition 1.4. A voting function $\sigma : \mathcal{P}_A^I \to \mathcal{P}(A)^+$ satisfies *unanimity* (or the *Pareto criterion*) if whenever P is a profile such that some alternative a is the $<_i^P$-greatest element for all $i \in I$, we have that $\sigma(P) = \{a\}$, i.e., a is the unique winner.

Definition 1.5. A voting function $\sigma : \mathcal{P}_A^I \to \mathcal{P}(A)^+$ satisfies *monotonicty* if whenever P and Q are two profiles such that Q differs from P only in that some $i \in I$ has moved some $a \in A$ up in their preference order, then $a \in \sigma(P)$ implies $a \in \sigma(Q)$.

Definition 1.6. A voting function $\sigma : \mathcal{P}_A^I \to \mathcal{P}(A)^+$ is *non-dictatorial* if it is not a dictatorship.

Note that simple plurality satisfies all of the above. However, a common complaint against simple plurality is that in an election between two major candidates, it may allow "irrelevant" third-party candidates to influence which of the two major candidates is declared a winner. For a familiar example:

Example 1.7. In the 2000 US Presidential Election in the state of Florida, George W. Bush received 2,912,790 votes, Al Gore received 2,912,253 votes (a margin of 537 votes, officially), and Ralph Nader received 97,488 votes. It it is reasonable to infer from polling and political ideology that the vast majority of Nader's supporters preferred Gore to Bush. If those voters had voted for Gore instead, the ultimate outcome of the election would have been a Gore victory and Bush defeat, despite no change in any voters preference between Bush and Gore.

In [1], Arrow isolated a criterion which would avoid this scenario.

Definition 1.8. A voting function $\sigma : \mathcal{P}_A^I \to \mathcal{P}(A)^+$ is *independent* (or satisfies *independence of irrelevant alternatives*) if whenever a and b are distinct alternatives, and P and Q are two profiles in which $<_i^P$ and $<_i^Q$ agree about a and b for all $i \in I$, then $a \in \sigma(P)$ and $b \notin \sigma(P)$ implies $b \notin \sigma(Q)$.

The example above shows that simple plurality, with three or more candidates, fails independence.

2 Arrow's Theorem

Does there exist a voting function which is unanimous, monotone, independent, and non-dictatorial? Arrow's theorem says "no".

Theorem 2.1 (Arrow [1] [2])**.** *Let I be a finite electorate, and A a finite set of at least three alternatives. Then, any voting function $\sigma : \mathcal{P}_A^I \to \mathcal{P}(A)$ which satisfies unanimity, monotonicity and independence must be a dictatorship.*

We will follow, with minor changes for our voting-specific setting, the proof of this result given in [5]. We note that monotonicity can be omitted in the above theorem; we use it merely to simplify the proof. To this end, fix such I, A, and a unanimous, monotone, independent voting function σ. The crucial notion is as follows:

Definition 2.2. For $X \subseteq I$ and $a, b \in A$ distinct, we say that X is *decisive for a against b* if whenever a profile P is such that $\forall i \in X(b <_i^P a)$, we have that $b \notin \sigma(P)$.

That is, the coalition of voters X is decisive for a against b if they can force b to lose by placing a above b.

Definition 2.3. For $X \subseteq I$, we say that X is a *dictating set* if for all $a, b \in A$ distinct, X is decisive for a against b.

Observe that unanimity is exactly the statement that I itself is a dictating set, and a dictatorship occurs exactly when there is a singleton dictating set $\{i\}$. Furthermore, unanimity and having more than one alternative implies that $\emptyset$ is never decisive.

Lemma 1. *Let $a, b \in A$ be distinct, and $X \subseteq I$. Suppose that there is a profile P such that*

$$\forall i \in X(b <_i^P a) \quad \wedge \quad \forall j \notin X(a <_j^P b) \quad \wedge \quad a \in \sigma(P) \quad \wedge \quad b \notin \sigma(P).$$

Then, X is decisive for a against b.

Proof. Let Q be another profile in which $\forall i \in X(b <_i^Q a)$. We must show that $b \notin \sigma(Q)$. Observe that by moving b up in the preference lists of those not in X, one at a time, we can go from Q to a profile Q' where $\forall i \in X(b <_i^Q a)$, and $\forall j \notin X(a <_j^Q b)$. If $b \in \sigma(Q)$, then by monotonicity, $b \in \sigma(Q')$, but comparing Q' and P, we see that independence implies $b \notin \sigma(Q')$. Thus, $b \notin \sigma(Q)$. □

The intuition behind part of the next lemma is that if X is decisive for a against b, then X is decisive for a against any alternative, and X is decisive for any alternative against b.

Lemma 2. *Let $a, b, c \in A$ be distinct, and $X \subseteq I$ decisive for a against b. If $X = Y \cup Z$, with $Y \cap Z = \emptyset$, then either Y is decisive for a against c, or Z is decisive for c against b. In particular, X is decisive both for a against c, and c against b.*

Proof. Let $X = Y \cup Z$ as above. Consider a profile P such that

$$\begin{cases} c <_i^P b <_i^P a & \text{if } i \in Y \\ b <_i^P a <_i^P c & \text{if } i \in Z \\ a <_i^P c <_i^P b & \text{otherwise,} \end{cases}$$

and all other alternatives are ranked below a, b and c, by all voters. Since everyone in both Y and Z has $b <_i^P a$, and since X is decisive for a against b, we have $b \notin \sigma(P)$.

Case 1: Suppose that $c \in \sigma(P)$. Then, we have produced a profile such that $\forall i \in Z(b <_i^P c)$, $\forall j \notin Z(c <_j^P b)$, $c \in \sigma(P)$ and $b \notin \sigma(P)$. By Lemma 1, this shows that Z is decisive for c against b.

Case 2: Suppose that $c \notin \sigma(P)$. Then, we have produced a profile such that $\forall i \in Y(c <_i^P a)$, $\forall j \notin Y(a <_j^P c)$, and $c \notin \sigma(P)$. Moreover, by unanimity, we must also have that $a \in \sigma(P)$. By Lemma 1, this shows that Y is decisive for a against c.

The remaining observation is seen by setting one of Y or Z to be X. □

Lemma 3. *If $X \subseteq I$ is decisive for a against b, then X is decisive b against a.*

Proof. Suppose X is decisive for a against b, and choose an alternative c distinct from a and b. Lemma 2 says that X is decisive for a against c. By another application of Lemma 2, X is decisive for b against c. A third application of Lemma 2 gives that X is decisive for b against a, as claimed. □

Lemma 4. *If $X \subseteq I$ is decisive for any distinct elements $a, b \in A$, then X is a dictating set.*

Proof. Suppose that X is decisive for a against b. Lemma 4 implies that X is also decisive for b against a. Suppose that $x, y \in A$ are distinct alternatives. We claim that X is decisive for x against y. If $a = y$, then this follows immediately from Lemma 2. If $a \neq y$, Lemma 2 says that X is decisive for a against y, and by Lemma 2 again, decisive for x against y. □

Lemma 5. *If $X \subseteq I$ is a dictating set, and $X = Y \cup Z$ is a partition of X, then either Y is a dictating set, or Z is a dictating set.*

Proof. Suppose that X is a dictating set, and $X = Y \cup Z$ is a partition of X. Choose three distinct alternatives, a, b and c. Since X is a dictating set, X is decisive for a against b. Then, by Lemma 1, either Y is decisive for a against c, or Z is decisive for c against b. Then, Lemma 4 implies that one of Y or Z is a dictating set. □

Arrow's Theorem follows immediately by repeatedly dividing the dictating set I into pieces two pieces, until we are left with a singleton $\{i\}$, in which case i is "the dictator". □

3 Infinite Electorates

We now consider the case when the electorate I is not necessarily finite. For simplicity, we continue to assume that the set of all alternatives A has at least three elements, but remains finite. (This assumption is not necessary in the general setting where the output is a weak order on A.)

Definition 3.1. An *ultrafilter* $\mathcal{U}$ on I is collection of subsets of I such that
(i) $I \in \mathcal{U}$ and $\emptyset \notin I$,
(ii) if $X \in \mathcal{U}$ and $X \subseteq Y$, then $Y \in \mathcal{U}$,
(iii) if $X, Y \in \mathcal{U}$, then $X \cap Y \in \mathcal{U}$,
(iv) for any $X \subseteq I$, either $X \in \mathcal{U}$ or $I \setminus X \in \mathcal{U}$.
Any collection satisfying (i)-(iii) is called a *filter* on I.

One can show that a filter is maximal (with respect to containment amongst filters) if and only if it is an ultrafilter.

Example 3.2. Fix $i \in I$, and let $\mathcal{U} = \{X \subseteq I : i \in X\}$. It is easy to see that $\mathcal{U}$ is an ultrafilter, called a *principle ultrafilter*.

Example 3.3. Let $\mathcal{F}$ be the collection of all subsets of I whose complements are finite. It is easy to see that $\mathcal{F}$ is a filter. Since the union of an increasing $\subseteq$-chain of filters on I is again a filter on I, and application of Zorn's Lemma shows that $\mathcal{F}$ is contained in a maximal filter (i.e., ultrafilter) $\mathcal{U}$. Moreover, since $\mathcal{F}$ contains the complements of singletons, $\mathcal{U}$ cannot be principle.

Theorem 3.4 (Fishburn [3], Kirman–Sondermann [4])**.** *Let $\mathcal{U}$ be an ultrafilter on I. Then, the map $\sigma_{\mathcal{U}} : \mathcal{P}_A^I \to \mathcal{P}(A)$ defined by*

$$a \in \sigma_{\mathcal{U}}(P) \quad \iff \quad \{i \in I : a \textit{ is the } <_i^P\textit{-greatest element}\} \in \mathcal{U}$$

is a well-defined voting function which is monotone, unanimous and independent. Moreover, $\sigma_{\mathcal{U}}$ is dictatorial if and only if $\mathcal{U}$ is principle, and distinct ultrafilters yield distinct voting functions.

Proof. To see that $\sigma_{\mathcal{U}}$ is well defined, given P, note that the sets

$$U_a = \{i \in I : a \text{ is the } <_i^P\text{-greatest element}\}$$

partition I, as a ranges over A. Since we assume that A is finite, and $\mathcal{U}$ is an ultrafilter, exactly one of these sets must be in $\mathcal{U}$, say U_{a_0}, in which case $\sigma_{\mathcal{U}}(P) = \{a_0\}$.

To see that $\sigma_{\mathcal{U}}$ is unanimous, suppose that for some $a, b \in A$, $\forall i \in I(a <_i^P b)$, then clearly $\{i \in I : a \text{ is the } <_i^P\text{-greatest element}\} = \emptyset \notin \mathcal{U}$, so $a \notin \sigma_{\mathcal{U}}(P)$. Similarly, monotonicty follows from the fact that $\mathcal{U}$ is closed upwards.

To see that $\sigma_{\mathcal{U}}$ is independent, suppose that P and Q agree on the distinct alternatives $a, b \in A$, and that $a \in \sigma(P)$. Then,

$$\begin{aligned} \{i \in I : b \text{ is the } <_i^Q\text{-greatest element}\} &\subseteq \{i \in I : a <_i^Q b\} \\ &= \{i \in I : a <_i^P b\} \notin \mathcal{U}, \end{aligned}$$

showing that $b \notin \sigma_{\mathcal{U}}(Q)$.

Suppose $\sigma_{\mathcal{U}}$ is dictatorial, say with dictator i_0. Let P be a profile in which a is i_0's first choice, and some $b \neq a$ is the first choice of all other voters. Since i_0 is the dictator, $\sigma_{\mathcal{U}}(P) = \{a\}$, and in particular, $\{i_0\} = \{i \in I : a \text{ is the } <_i^P\text{-greatest element}\} \in \mathcal{U}$. Thus, $\mathcal{U}$ is principle. The converse is similar.

Lastly, if $\mathcal{U}$ and $\mathcal{V}$ are distinct ultrafilters, say with $X \in \mathcal{U}$ and $X^c \in \mathcal{V}$, then we can define a profile P such that all of the voters in X rank a first, and all of the voters in X^c rank b first. Clearly, $\sigma_{\mathcal{U}}(P) = \{a\}$ while $\sigma_{\mathcal{V}}(P) = \{b\}$. □

Thus, non-principle ultrafilters yield non-dictatorial voting functions, satisfying monotonicity, unanimity and independence! In fact, the converse of the above theorem is true as well; all monotone, unanimous and independent voting functions arise as $\sigma_{\mathcal{U}}$ for some ultrafilter $\mathcal{U}$. First, we need more useful criterion for being an ultrafilter:

Lemma 3.5. *A collection $\mathcal{U}$ of subsets of I is an ultrafilter if and only if $I \in \mathcal{U}$, $\emptyset \notin \mathcal{U}$, $\mathcal{U}$ is closed finite intersections, and for all $X \in \mathcal{U}$, if $X = Y \cup Z$ is a partition, then one of Y or Z is in $\mathcal{U}$.*

Proof. ($\Rightarrow$): Given $X \in \mathcal{U}$, suppose $X = Y \cup Z$ is a partition. Since $\mathcal{U}$ is an ultrafilter, one of Y or Y^c is in $\mathcal{U}$, and by intersecting with X, we obtain that Y or Z is in $\mathcal{U}$. The other conditions are automatic.
($\Leftarrow$): Note that the hypotheses imply that for any $X \subseteq I$, one of X or X^c is in $\mathcal{U}$. It remains to check that $\mathcal{U}$ is closed upwards. Let $X \in \mathcal{U}$ and $X \subseteq Y \subseteq I$. If $Y \notin \mathcal{U}$, then $Y^c \in \mathcal{U}$, but then $Y^c \cap X = \emptyset \in \mathcal{U}$, a contradiction. □

Theorem 3.6 (Kirman–Sondermann [4])**.** *If $\sigma : \mathcal{P}_A^I \to \mathcal{P}(A)$ is a voting function which is monotone, unanimous and independent, then there is a (unique) ultrafilter $\mathcal{U}$ on I such that $\sigma = \sigma_{\mathcal{U}}$.*

Proof. Given such a σ, let $\mathcal{U}$ be the collection of all dictating sets in I. First, we need to verify that $\mathcal{U}$ is an ultrafilter on I. Unanimity implies both that $I \in \mathcal{U}$ and $\emptyset \notin \mathcal{U}$. Lemma 5 above, which did not require that I was finite, implies that for all $X \in \mathcal{U}$, if $X = Y \cup Z$ is a partition, then one of Y or Z is in $\mathcal{U}$. It remains to check that $\mathcal{U}$ is closed under intersection.

Let $X, Y \in \mathcal{U}$. Define sets

$$V_1 = X \cap Y, \quad V_2 = X \cap Y^c, \quad V_3 = X^c \cap Y, \quad V_4 = (X \cup Y)^c.$$

Note that these partition I. For distinct $a, b, c \in A$, define a profile P such that

$$\begin{cases} c <_i^P a <_i^P b & \text{if } i \in V_1 \\ a <_i^P b <_i^P c & \text{if } i \in V_2 \\ b <_i^P c <_i^P a & \text{if } i \in V_3 \\ b <_i^P a <_i^P c & \text{if } i \in V_4 \end{cases}$$

with all other alternatives below a, b and c. Since $X = V_1 \cup V_2$, it follows that $a \notin \sigma(P)$. Since $Y = V_1 \cup V_3$, it follows that $c \notin \sigma(P)$. Unanimity implies that we must have $b \in \sigma(P)$. But then, Lemma 1 (which did not require I to be finite) implies that $V_1 = X \cap Y$ is decisive for b over c, and thus by Lemma 4, a dictating set. That $\sigma = \sigma_{\mathcal{U}}$ is immediate from the definition of "dictating". □

4 Infinite democracies?

While the results of the previous section seem to indicate the possibility of "ideal" voting methods on infinite electorates, there are (at least) two caveats to this view. First, non-principle ultrafilters are strange objects, and in fact, can never be constructed in a definable fashion (the deep set-theoretic reason for this being that they yield non-measurable subsets of 2^I, which when I is countably infinite, can never be "definable" in any reasonable sense). A much simpler reason is as follows:

Lemma 4.1. *Let $\mathcal{U}$ be an ultrafilter on an infinite I. Suppose that I is equipped with a standard atomless probability measure μ. Then, for any $\epsilon > 0$, there is $U \in \mathcal{U}$ such that $\mu(U) < \epsilon$.*

Proof. Since μ is atomless and finite valued, it is an easy exercise to build a partition of measurable sets $I = X_1 \cup X_2 \cup \cdots \cup X_n$ such that for each $1 \leq i \leq n$, $\mu(X_i) < \epsilon$. One of the X_i must be in $\mathcal{U}$. □

Corollary 4.2 (Kirman–Sondermann [4])**.** *If I is an infinite set equipped with a standard atomless probability measure μ, A a finite set of at least three alternatives, and $\sigma : \mathcal{P}_A^I \to \mathcal{P}(A)^+$ a unanimous, monotone, independent voting function, then for any $\epsilon > 0$, there is a measurable dictatorial set $U \subseteq I$ with $\mu(U) < \epsilon$.*

References

[1] K. J. Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):pp. 328–346, 1950.

[2] K. J. Arrow. *Social Choice and Individual Values*. Yale University Press, 1952.

[3] P. C. Fishburn. Arrow's impossibility theorem: Concise proof and infinite voters. *Journal of Economic Theory*, 2(1):103–106, 1970.

[4] A. P. Kirman and D. Sondermann. Arrow's theorem, many agents, and invisible dictators. *Journal of Economic Theory*, 5(2):267–277, 1972.

[5] A. D. Taylor and A. M. Pacelli. *Mathematics and Politics: Strategy, Voting, Power, and Proof*. Springer, 2008.