Computational Social Choice

Lecture Notes

based on a lecture by Prof. F. Brandt

Marvin Jahn

mail@marvin-jahn.de

March 1, 2022

Contents

These are (unofficial) lecture notes for the lecture *Computational Social Choice* held by Prof. F. Brandt at the Technical University Munich in the winter semester 2021/22.

Some proofs, many examples and some complexity-theoretic results (hardness results) are omitted.

1 Basics of Social Choice Theory

1.1 Motivation

Lec 1 2021-10-19

Social choice theory tries to aggregate possibly conflicting preferences into collective choices in a "good" way. Such a mechanism is important in practice. Indeed, its main application is to identify voting systems that yield "reasonable" results for an election, but it can also be used to find suitable coalitions or determine a "fair" allocation of resources.

1.2 Rational Choices

Lec 2 2021-10-26

We start by formalizing what it means to make rational choices.

Definition 1.1. Let U be a finite set with $m \in \mathbb{N}$ elements (*alternatives*). We call the set $F(U) := \mathcal{P}(U) \setminus \{\emptyset\}$ the set of **feasible sets**. A **choice function** is a function $S\colon F(U) \to F(U)$ with $F(A) \subset A$ for $A \in F(U)$.

We always implicitly assume that our set U of alternatives is finite.

We can think of the input of the function as the available alternatives. By allowing this function to return a set of alternatives, we allow for indifference between certain alternatives.

Also note that a choice function only needs to be specified on sets with at least two elements.

A rational decision maker should have preferences over all alternatives (in the whole universe) that are independent from their feasibility.

Definition 1.2. A **preference relation** is a complete binary relation on the set of alternatives U.

Any binary relation $R \subset U \times U$ admits a disjoint decomposition into two parts:

- *asymmetric part* P: $(a, b) \in P \subset U \times U$ if and only if $(a, b) \in R$ and $(b, a) \notin R$.
- *symmetric part* I: $(a, b) \in I \subset U \times U$ if and only if $(a, b) \in R$ and $(b, a) \in R$.

We usually write $\geq$ for R, $>$ for P and $\sim$ for I. We may think of a preference relation as a table in which each line may contain multiple alternatives and the top alternatives are the ones most liked.

Definition 1.3. Let R be a binary relation on U and $A \in F(U)$. The set of maximal elements of A w.r.t. the strict part of R is called **maximal set** $\mathrm{Max}(R, A)$.

If R is a complete relation (e.g. a preference relation), this is equivalent to the set of greatest elements of A.

Example 1.4. The complete relation $a > b > c > a$ on three alternatives $U = \{a, b, c\}$ has an empty maximal set. This corresponds to the cyclic graph with three vertices.

Definition 1.5. A binary relation $\geq$ on U is called

(a) **transitive**, if for all $x, y, z \in U$, $x \geq y$ and $y \geq z$ implies $x \geq z$.

(b) **quasi-transitive**, if its asymmetric part is transitive; i.e. for all $x, y, z \in U$, $x > y$ and $y > z$ implies $x > z$.

(c) **acyclic**, if for all $x_1, \dots, x_n \in U$, $x_1 > x_2, \dots, x_{n-1} > x_n$ implies $x_1 \geq x_n$.

Any (finite) binary relation R on a set U can be represented by a directed graph, where the vertices are the elements of U and an edge $u \to v$ exists if and only if $u \geq v$. If R is complete or asymmetric, then the subgraph given by $u \to v$ if and only if $u > v$ already completely determines R.
Also note that the asymmetric part of an antisymmetric relation R is just the same relation R but with the diagonal $\{(x, x) : x \in R\}$ removed.

Proposition 1.6. For any binary relation $\geq$ on a (finite) set U, it holds

$$\geq \text{ transitive } \implies \geq \text{ quasi-transitive } \implies \geq \text{ acyclic.}$$

If $\geq$ is asymmetric, these are equivalences.

Proof. The first part is straightforward to show, the second follows directly from the previous observation. □

Lemma 1.7. Let R be a preference relation on U.
Then $\mathrm{Max}(R, A) \neq \emptyset$ for all $A \in F(U)$ if and only if R is acyclic.
In particular, R induces a choice function

$$\mathrm{Max}(R, -) \colon F(U) \to F(U)$$

if and only if it is acyclic.

Proof. If $x_1 > x_2 > \dots > x_n$, then $\mathrm{Max}(R, \{x_1, \dots, x_n\}) = \{x_1\}$.
For the other direction, let $A \in F(U)$ and pick an arbitrary element $a_1 \in A$. For $i \in \{1, \dots, |A|\}$, if $a_i \in \mathrm{Max}(R, A)$, then $\mathrm{Max}(R, A) \neq \emptyset$ and we are done. Otherwise, there exists $a_{i+1} \in A$ with $a_{i+1} > a_i$ and because R is acyclic, we have $a_{i+1} \in A \setminus \{a_1, \dots, a_i\}$. Since A is finite, iterating this must eventually terminate. □

Definition 1.8. A choice function $S \colon F(U) \to F(U)$ is called **rationalizable**, if there exists a binary relation R on U, such that $S = \mathrm{Max}(R, -)$.
The **base relation** of S is the relation $R_S \subset U \times U$ given by

$$(x, y) \in R_S \;:\Longleftrightarrow\; x \in S(\{x, y\}).$$

Lemma 1.9. A choice function $S \colon F(U) \to F(U)$ is rationalizable if and only if it is rationalized by its base relation; i.e. $S = \mathrm{Max}(R_S, -)$.

Proof. Let S be rationalizable by R. Then $R = R_S$, since

$$(x, y) \in R \iff x \in S(\{x, y\}) \iff (x, y) \in R_S.$$

□

By Lemma 1.7 and Lemma 1.9, the base relation of a rationalizable preference relation is acyclic.

Example 1.10. A choice function $S\colon F(\{a,b,c\}) \to F(\{a,b,c\})$ with $S(\{a,b,c\}) = \{a\}$ and $S(\{a,b\}) = \{b\}$ cannot be rationalized. This follows from the previous lemma.

Our next goal is to find *consistency* criteria which are necessarily satisfied by rationalizable choice functions.

Definition 1.11. A choice function $S\colon F(U) \to F(U)$ satisfies **contraction** (short: α) if for all $A, B \in F(U)$, $B \subset A$ implies $S(A) \cap B \subset S(B)$.

In words, a choice function satisfies contraction, if any alternative $u \in U$ that is chosen in some set A is also chosen for any subset $B \subset A$ with $u \in B$.

Lemma 1.12. If a choice function $S\colon F(U) \to F(U)$ satisfies α, then its base relation R_S is acyclic.

Proof. Let $x_1 > x_2 > \cdots > x_n$ and $A := \{x_1, \ldots, x_n\}$. If $x_1 \notin S(A)$, then there must exist some $i \in \{2, \ldots, n\}$, such that $x_i \in S(A)$. But because S satisfies α, this implies $x_i \geq x_{i-1}$, which contradicts our assumption and thus shows $x_1 \in S(A)$. Applying the α property to $\{x_1, x_n\}$, we conclude $x_1 \geq x_n$. □

Definition 1.13. A choice function $S\colon F(U) \to F(U)$ satisfies **expansion** (short γ) if for all $A, B \in F(U)$, we have $S(A) \cap S(B) \subset S(A \cup B)$.

A. K. Sen showed the following central result in 1971.

Theorem 1.14. A choice function $S\colon F(U) \to F(U)$ is rationalizable if and only if it satisfies α and γ.

Proof. If S is rationalizable, then $S = \operatorname{Max}(R, -)$ and $\operatorname{Max}(R, -) : F(U) \to F(U)$ obviously satisfies α and γ for any binary relation R.
On the other hand, assume that S satisfies α and γ and let $A \in F(U)$. We need to show that $\operatorname{Max}(R_S, A) = S(A)$.
If $x \in \operatorname{Max}(R_S, A)$, then $(x, a) \in R_S$ for all $a \in A$; i.e. $x \in S(\{x, a\})$ for all $a \in A$. Iterating the γ property, this implies $x \in S(A)$.
Similarly, $x \in S(A)$ satisfies $x \in S(\{x, a\})$ for any $a \in A$ by the α property, so $x \in \operatorname{Max}(R_S, A)$. □

We give an equivalent characterization of α and γ and thus of rationalizability.

- **Contraction** (α): $\forall A, B \in F(U)\colon S(A \cup B) \cap A \cap B \subset S(A) \cap S(B)$.
- **Expansion** (γ): $\forall A, B \in F(U)\colon S(A \cup B) \cap A \cap B \supset S(A) \cap S(B)$.
- **Rationalizability**: $\forall A, B \in F(U)\colon S(A \cup B) \cap A \cap B = S(A) \cap S(B)$.

Definition 1.15. A choice function S satisfies **strong expansion** (short: β^+) if for all $A, B \in F(U)$, $B \subset A$ and $S(A) \cap B \neq \emptyset$, we have $S(B) \subset S(A)$.

Intuitively, a choice function satisfies β^+ if whenever an element $x \in A$ is chosen in $S(A)$, then it is also chosen in all supersets whose winners contain some element of A.

Lemma 1.16. Any choice function that satisfies β^+ also satisfies γ.

K. Arrow showed in 1959:

Theorem 1.17. A choice function $S\colon F(U) \to F(U)$ is rationalizable by a transitive relation if and only if it satisfies α and β^+.

1.3 Social Choice Functions

Lec 3 2021-11-09

Throughout this section, let $N := \{1, \dots, n\}$ denote the finite set of "voters" with $n \geq 2$ and let $R(U)$ be the set of all transitive and complete relations over the set of alternatives U.

Definition 1.18. The elements of $R(U)^n$ are called **preference profiles**.
A **social choice function** (**SCF**) is a function

$$\phi\colon R(U)^n \to \{\text{choice function } S\colon F(U) \to F(U)\},$$

which assigns to every preference profile R and every collection of alternatives A a set of chosen alternatives ("winners") $\phi(R, A) \subset A$.

Notions like rationalizability and consistency carry over from the previous section by saying that a social choice function satisfies such a property if all the choice functions in its image satisfy it.

Definition 1.19. A SCF ϕ is called **anonymous**, if it is invariant under permutations $\pi\colon N \to N$; i.e. if $R_i = R'_{\pi(i)}$ for all $i \in N$ implies that $\phi(R) = \phi(R')$.

Intuitively, an anonymous SCF is one where all voters are equal.

Definition 1.20. A SCF ϕ is called **neutral**, if for any $A, B \in F(U)$, $R \in R(U)^n$ and bijection $\pi\colon A \to B$, we have $\pi(\phi(R, A)) = \phi(R', B)$ for all $R' \in R(U)^n$ such that $\pi(R'|_A) = R|_A$. Here $\pi(R'|_A)$ denotes the relation R'' on A for which $x >_{R''} y$ if and only if $\pi(x) >_{R'} \pi(y)$ for $x, y \in A$.

This definition is quite strong, it not only corresponds to the intuitive notion of a SCF being invariant under renaming the alternatives (for example, the SCF "A is always the unique winner" is not neutral) but it also implies independence of preferences over alternatives that are not contained in the feasible set (take $A = B$ and $\pi = \mathrm{id}_A$).

Definition 1.21. The finite product $\prod_{i=1}^n A_i$ of sets with relations (A_i, R_i) is equipped with the **$\wedge$-order**[1] R:

$$(x, y) \in R \;:\Longleftrightarrow\; \forall\, i \in \{1, \dots, n\} : (x, y) \in R_i.$$

Definition 1.22. Let R be a preference profile and $x, y \in U$ two alternatives. With the construction above applied to the strict parts of the R_i, we get a strict preorder $>$ on U.

(a) x **Pareto-dominates** y if $x > y$.

(b) y is **Pareto-optimal**, if it is a maximal element w.r.t. $>$. Otherwise, it is called **Pareto-dominated**.

As $>$ is a transitive relation on the finite set U, there always is at least one Pareto-optimal alternative, which allows us to define a SCF based on Pareto-optimality.

Definition 1.23. The **Pareto SCF** $\mathrm{pareto}(R, A)$ returns all Pareto-optimal alternatives in A.

[1] In the language of category theory, this is just the product in the category of sets with relations.

It is anonymous and neutral.

Definition 1.24. A SCF ϕ satisfies **Pareto-opimality**, if $\phi(R, A) \subset \text{pareto}(R, A)$ for all $A \in F(U)$, $R \in R(U)^n$.

A Pareto-optimal SCF only returns Pareto-optimal alternatives.

Definition 1.25. A SCF ϕ is called **resolute**, if $|\phi(R, A)| = 1$ for all $A \in F(U)$, $R \in R(U)^n$.

One reason why resoluteness is not always desirable is that there do not always exist anonymous and neutral resolute SCFs. In fact, Moulin showed in 1983:

Theorem 1.26. There is an anonymous, neutral, Pareto-optimal and resolute SCF for strict preferences with n voters on m alternatives if and only if no $q \in \mathbb{N}$ with $2 \leq q \leq m$ divides n.

Definition 1.27. Given a set A equipped with a relation $R \subset A \times A$, its power set $\mathcal{P}(A)$ can be equipped with the **power set relation**[2] $\mathcal{P}(R)$

$$(X, Y) \in \mathcal{P}(R) \;:\Longleftrightarrow\; X = \{x\} \wedge Y = \{y\} \wedge (x, y) \in R.$$

It is clear that there are other ways to extend a relation to its power set.

Definition 1.28. A SCF ϕ is called **manipulable** by voter $i \in N$, if there exists $A \in F(U)$ and two preference profiles $R, R' \in R(U)^n$ whose j-th components agree for all $j \neq i$, such that i strictly prefers $\phi(R', A)$ to $\phi(R, A)$ w.r.t. the power set relation $\mathcal{P}(P_i)$. A SCF is called **strategyproof**, if it is not manipulable by any voter.

Intuitively, this means that a manipulable SCF allows some voter to misrepresent their preferences in order to obtain a preferred outcome. This assumes that the voter precisely knows how the other voters will vote.

Definition 1.29. A SCF ϕ is called **monotonic**, if for any voter $i \in N$, alternative $a \in U$ and $A \in F(U)$ the following is true: For two preference profiles R, R' that agree on all but the i-th component and such that

$$\begin{aligned} \forall\, x, y \in U \setminus \{a\} \colon (x, y) \in R_i &\iff (x, y) \in R'_i, \\ (a, y) \in R_i &\implies (a, y) \in R'_i, \\ (a, y) \in P_i &\implies (a, y) \in P'_i, \end{aligned}$$

we have $a \in \phi(R', A)$ whenever $a \in \phi(R, A)$. It is called **positive responsive** if the additional assumption $R_i|_A \neq R'_i|_A$ implies $\{a\} = \phi(R', A)$ whenever $a \in \phi(R, A)$.

The idea behind the monotonicity property is as follows: If a voter increases their opinion of an alternative that was already a winner before, then this alternative should still be a winner. If the SCF is even positive responsive, then this process should even distinguish the alternative as the unique winner.

Intuitivley, a monotonic SCF on two alternatives $\{a, b\}$ roughly distinguishes three cases: For few votes for a, b usually wins (the “usually” refers to rules like “b always wins”). Increasing the number of votes for a, we may get into the “indifference” case, where both alternatives win. Finally, if we further increase the number of votes for a, we may enter the third case in which a becomes the unique winner.

[2]In the language of category theory, we obtain an endofunctor on the category of sets with relations.

Definition 1.30. A SCF ϕ satisfies **independence of infeasible alternatives** (**IIA**), if for all $A \in F(U)$, and $R, R' \in R(U)^n$ which agree on A, we have $\phi(R, A) = \phi(R', A)$.

Note that neutrality implies IIA and that any SCF on two alternatives trivially satisfies IIA.

By definition of a SCF ϕ, we get a map $\phi\colon F(U) \to (R(U)^n \to F(U))$. Note that

$$R \sim R' \;:\Longleftrightarrow\; \forall\, i : R_i|_A = R'_i|_A$$

defines an equivalence relation on $R(U)^n$ and that the projection

$$R(U)^n \to R(A)^n,\;\; R = (R_i) \mapsto (R_i|_A)$$

induces an isomorphism (of sets) $R(U)^n/\sim \cong R(A)^n$.
Then ϕ satisfies IIA if and only if $\phi(A)\colon R(U)^n \to F(U)$ decends to a map $\phi(A)\colon R(A)^n \cong R(U)^n/\sim \to F(U)$ for all $A \in F(U)$. Roughly, this means that when e.g. applied to $A = \{a, b\}$ and a preference profile R like

n	m
a	b
b,c	c
d	a
	d

ϕ actually takes

n	m
a	b
b	a

as input.

Suppose that for each nonempty subset $V \subset U$, we have a map $\phi_V\colon R(V)^n \to R(V)$. Then we may construct a SCF $\phi(A)$ for $A \in F(U)$ as the composition

$$R(U)^n \xrightarrow{\;\pi\;}\!\!\!\!\rightarrow R(A)^n \xrightarrow{\;\phi_V\;} F(A) \xhookrightarrow{\;\iota\;} F(U),$$

which by construction satisfies IIA.

Example 1.31. The **Borda rule** is a SCF that for $R \in R(U)^n$ and $A \in F(U)$ calculates the "score" of an alternative $a \in U$ by assigning no points for every time a voter listed the alternative at the last place, 1 point for every second last place and so on. Restricting to A, it returns those alternatives that have maximal score.

For example, the Borda rule when applied to $U = \{a, b, c\}$, $A = \{a, c\}$ and the preference profile

1	2
a	b
b	c
c	a

returns $\{a, c\}$. This example also shows that the Borda rule does not satisfy IIA, because by moving b up in the left column, the score of a decreases, so c becomes the unique winner.
Since we can define the Borda rule for every subset $V \subset U$, we can apply the above construction to obtain a SCF ϕ, which when applied to $A = \{a, c\}$ simply restricts to the subtable

1	2
a	c
c	a

before counting points. Here we also see that ϕ is not rationalizable, because $\phi(A) = \{a, c\}$, but $c > a$ in the base relation.
The same construction can be applied to majority rule.

Theorem 1.32. A resolute SCF on two alternatives is strategyproof if and only if it is monotonic. Furthermore, any (not necessarily resolute) SCF on two alternatives that is monotonic is also strategyproof.

May showed the following important theorem in 1952.

Theorem 1.33 (May's Theorem). Majority rule is the only SCF on two alternatives that is anonymous, neutral and positive responsive.

Intuitively, it is comprehensible that majority rule is the most decisive (the fewest amount of ties) SCF out of all anonymous, neutral and monotonic SCFs. For example, the Pareto SCF satisfies those three properties, but is not positive responsive.
The theorem shows that for two alternatives, majority rule is the "best" voting rule. In particular, many other voting rules agree with majority rule for the case of only two alternatives.
Also note that any SCF on two alternatives is trivially rationalizable.

Corollary 1.34. Let ϕ be an anonymous, neutral and positive responsive SCF. Then for any subset A with two elements, $\phi(-, A) = M(-, A)$, where M denotes the majority rule SCF.

Proof. Since ϕ is neutral, it also satisfies IIA, so it descends to a function $\phi(A)\colon R(A)^n \to F(U)$, which induces a SCF on $A \subset U$

$$\phi_A\colon R(A)^n \to \{\text{choice function } S\colon F(A) \to F(A)\}.$$

This SCF takes $R \in R(A)^n$ and $B \in F(A) \subset F(U)$, extends R to a complete relation R' on U in such a way that $R'|_A = R$ and then returns the result of $\phi(R', B)$. It is anonymous, neutral and positive responsive as the same holds for ϕ. May's Theorem (Theorem 1.33) implies that ϕ_A is just majority rule, so for $B = A$, we conclude $\phi(R', A) = \phi_A(R, A) = M(R, A) = M(R', A)$ for any $R' \in R(U)^n$ with restriction $R \in R(A)^n$. □

Theorem 1.35. No anonymous, neutral and positive responsive SCF is rationalizable, whenever there are at least three alternatives and voters.

Proof. It is enough to prove this for three alternatives $U = \{a, b, c\}$ and three voters because we may add additional completely indifferent voters and bottom-ranked alternatives (using that the SCF is assumed to be neutral, thus satisfies IIA).
Let ϕ be a SCF that is anonymous, neutral and positive responsive. By Lemma 1.9, the only relation that could potentially rationalize ϕ is its base relation and by Corollary 1.34, this base relation chooses by majority. However, the preference profile

1	1	1
a	b	c
b	c	a
c	a	b

gives an acyclic base relation, showing that ϕ is not rationalizable. □

Any preference profile R induces a relation $C(R)$ on the set of alternatives U by declaring $a \geq b$ if and only if at least as many voters prefer a to b than the other way around. In other words, the relation chooses by majority and is thus called *majority relation.*
The previous theorem states that the base relation of $\phi(R)$ for $R \in R(U)^n$ is the majority relation.

Definition 1.36. An alternative $x \in U$ is a **Condorcet winner** w.r.t. a preference profile R, if it is a greatest element in U w.r.t. the strict part of $C(R)$.

The previous theorem essentially boils down to the fact that Condorcet winners need not exist. However, if they do, they are unique.

Definition 1.37. A SCF ϕ is **dictatorial**, if there exists a voter $i \in N$, such that for all $A \in F(U)$ and $R \in R(U)^n$, we have $\phi(R, A) = \{x\}$ whenever voter i strictly prefers x to all other alternatives in A.

Of course, anonymity implies non-dictatorship.
Arrow showed the following important result in 1951.

Theorem 1.38 (Arrow's Impossibility Theorem). There exists no SCF that satisfies IIA on two alternatives, Pareto-optimality on two alternatives, non-dictatorship on two alternatives and transitive rationalizibility whenever there are at least three alternatives.

1.4 Social Welfare Functions

Lec 4
2021-11-16

Definition 1.39. A **social welfare function** (**SWF**) is a function $\psi\colon R(U)^n \to R(U)$.

Intuitively, a social welfare function aggregates individual preference relations into a single, collective one.

For any SCF ϕ that is transitively rationalizable we may construct a canonical SWF, namely the one that takes a preference profile and returns the corresponding base relation of ϕ. On the other hand, any SWF ψ induces the transitively rationalizable SCF $\mathrm{Max}(\psi(-), -)$, which assigns $\mathrm{Max}(\psi(R), A)$ to $R \in R(U)^n$ and $A \in F(U)$.
Therefore, SWFs correspond to transitively rationalizable SCFs.

This bijection is used to translate the various notions we defined from SCFs to SWFs. Using those, we can restate Arrow's Impossibility Theorem for SWFs.

Theorem 1.40 (Arrow's Impossibility Theorem for SWFs). Every SWF that satisfies IIA and Pareto-optimality is dictatorial for at least three alternatives.

In order to prove this theorem, we need to introduce some terminology.

Definition 1.41. Let ψ be a SWF. A group of voters $G \subset N$ is called **decisive for** $a \in U$ **against** $b \in U$, if $a > b$ w.r.t. $\psi(R)$ for all $R \in R(U)^n$ with $a >_i b$ for all $i \in G$.
It is called **decisive**, if it is decisive for all $a, b \in U$.
It is called **semidecisive for** $a \in U$ **against** $b \in U$, if $a > b$ w.r.t. $\psi(R)$ for all $R \in R(U)^n$ with $a >_i b$ for all $i \in G$ and $b >_j a$ for all $j \notin G$.

By definition, decisiveness implies semidecisiveness.

Example 1.42. If ψ is dictatorial with dictator i, then precisely the supersets of i are decisive. W.r.t. majority rule, precisely those subsets are decisive, which have a size greater than $|N|2$. If ψ is Pareto-optimal, then the set of all voters N is decisive.

Lemma 1.43 (**Field Expansion Lemma**)**.** Let ψ be a SWF that satisfies IIA and Pareto-optimality and suppose that there are at least three alternatives. If there exists a group of voters $G \subset N$ that is semidecisive for $a \in U$ against $b \in U$ $(a \neq b)$ then G is decisive.

Lemma 1.44 (**Group Contraction Lemma**)**.** Let $G \subset N$ be a decisive group, partitioned into two subsets $G = G_1 \coprod G_2$. Then G_1 or G_2 is another decisive group.

The essence of Arrow's Impossibility Theorem and many related results is that rationalizability is incompatible with collective choices whenever there are at least three alternatives.

2 Domain Restrictions

2.1 Social Choice Functions in Domain

Lec 5
2021-11-23

Arrow's Impossibility Theorem is a very significant negative result in Social Choice Theory, since despite the rather weak assumptions, it asserts the nonexistence of any such SCF or SWF.
However, by imposing additional assumptions, one can obtain many positive results. Our first approach is to restrict the set of possible preference relations of the voters.

Definition 2.1. A subset $D(U) \subset R(U)$ is called a **domain**.
A SCF ϕ **satisfies some property in domain** $D(U)$, if the restriction $\phi|_{D(U)^n}$ satisfies this property.
These restrictions are functions of the form

$$\phi\colon D(U)^n \to \{\text{choice function } S\colon F(U) \to F(U)\},$$

with $\phi(D, A) \subset A$ for all $D \in D(U)$ and $A \in F(U)$ and are called **social choice function** (**SCF**) **in domain** $D(U)$.

Example 2.2. The domain of *strict* (or *linear*) *preferences* is

$$D(U) := \{R \in R(U) : \forall x, y \in U : x > y \vee y > x\};$$

there must not be any ties in the preferences of a voter.

One can show Arrow's Impossibility Theorem for the domain of strict preferences and in fact that version is independent from the other version (being a dictator in domain does not imply being a dictator in the non-restricted case).

The next property is similar to that of strategyproofness (Definition 1.28).

Definition 2.3. Let $n \in \mathbb{N}$. A collection of SCF

$$\{\phi_j \colon R(U)^j \to \{\text{choice function } S\colon F(U) \to F(U)\} : j \in \{2, \dots, n\}\}$$

can be **manipulated by strategic abstention**, if there exists $j \in \{3, \dots, n\}$, $R \in R(U)^j$ and $A \in F(U)$, such that voter i strictly prefers $\phi_{j-1}(R_{-j}, A)$ to $\phi_j(R, A)$ w.r.t. the power set relation $\mathcal{P}(P_j)$. Here $R_{-j} \in R(U)^{j-1}$ denotes the preference profile R, but with the j-th component removed.
Otherwise it satisfies **participation**.

For any preference profile $R \in R(U)^n$ (and any $n \in \mathbb{N}$, $n \geq 2$), there is the natural relation R_M induced by the majority SCF M, where

$$x \geq_{R_M} y \;:\Longleftrightarrow\; x \in M(R, \{x, y\})) \iff \text{In a runoff, } x \text{ scores at least as many points as } y.$$

By Lemma 1.7, $\mathrm{Max}(R_M, -)$ constitutes a SCF in domain D if and only if R_M is acyclic in D.

Theorem 2.4. If R_M is quasi-transitive in domain D, then $\mathrm{Max}(R_M, -)$ is a SCF in domain D, which is strategyproof and satisfies participation in D.

Proof. This follows easily from the oberservation that by quasi-transitivity of R_M, we have

$$\{a\} = \mathrm{Max}(R_M, A) \iff a \text{ is a Condorcet winner in } A.$$

□

2.2 Dichotomous Preferences

Definition 2.5. The domain of **dichotomous preferences** D_{DI} is

$$D_{DI} \coloneqq \{R \in R(U) : \forall\, x, y, z \in U, x >_R y : z \sim_R x \vee z \sim_R y\}.$$

For $D \in D_{DI}(U)^n$, we say that a voter i **likes** a candidate x, if x is top-ranked in D_i; i.e. if x is a largest element in D_i. Otherwise, x is a smallest element in D_i and we say that i **dislikes** x.

Intuitively, dichotomous preferences are precisely those where a voter may only agree or disagree with an alternative and are not allowed to have more sophisticated opinions.

In the domain D_{DI}, $\mathrm{Max}(R_M, -)$ is known as **approval voting**. Since

$$x \geq_{R_M} y \iff |\{i \in N : i \text{ likes } x\}| \geq |\{i \in N : i \text{ likes } y\}|,$$

the winners of this SCF are exactly the alternatives with the highest number of likes.

Since $\geq$ is transitive on $\mathbb{N}$, we immediately obtain the following result, first proven by Inada in 1964.

Theorem 2.6. The majority relation R_M is transitive in the domain of dichotomous preferences D_{DI}.

By Theorem 2.4, R_M is strategyproof and satisfies participation in D_{DI}.

2.3 General Preference Restrictions

From now on, we only consider SCFs in the domain of total orders (without noting this explicitly); i.e. we assume that the voters only have strict preferences.
Therefore, let $R(U)$ denote the set of all anti-symmetric, transitive and complete relations. Note that the asymmetric part P of a relation $R \in R(U)$ is just $P = R \setminus \{(x,x) : x \in R\}$. In many cases, we will additionally assume that the number of voters is odd, since this ensures that the majority relation R_M is strict.

2.4 Single-Peaked Preferences

In many cases, the space of alternatives comes equipped with a total order. For example, when the question is "How high should the tax rate be?", the set of alternatives $\{0\%, 1\%, \ldots, 100\%\}$ has the usual total order $\geq$. Then a "rational" voter would be assumed to have a favorite tax rate r (the peak when viewed as a graph) and to always prefer tax rates that are close to r to those which are further away. This idea of rationality is formalized by *single-peaked preferences.*

Definition 2.7. A preference profile $R \in R(U)$ is called **single-peaked** with respect to a total order $\geq$ over U, if for all $i \in N$ and $x, y, z \in U$:

- If $x > y > z$ and $x >_i y$, then $y >_i z$.
- If $z > y > x$ and $x >_i y$, then $y >_i z$.

We denote the domain of single-peaked preferences by D_{SP}.

Theorem 2.8. The majority relation R_M is transitive in the domain of single-peaked preferences D_{SP}.

Proof. Assume that xP_My and yP_Mz. In order to show xR_Mz, we differentiate three cases.
Case 1: $x > y > z$ (or $z > y > x$). Any voter that prefers x to y must also prefer y to z, and by transitivity, this means that any such voter prefers x to z. It follows xP_Mz.
Case 2: $z > x > y$ (or $y > x > z$). Any voter that prefers z to x must also prefer to x to y and by transitivity, this means that any such voter prefers z to y. Since yP_Mz, this means that less than half of the voters prefer z to x, implying xP_Mz.
Case 3: $y > z > x$ (or $x > z > y$). Any voter that prefers y to z must also prefer z to x and by transitivity, this means that any such voter prefers y to x. Therefore, more than half the voters prefer y to x, contradicting xP_My, thus showing that this case cannot occur. □

Since we assume that the number of voters is odd, the previous theorem implies that the SCF $\text{Max}(R_M, -)$ is resolute in D_{SP} and the unique winner is precisely the Condorcet winner.

For voter $i \in N$, let t_i be the most-prefered alternative ("peak") of voter i; that is, $\{t_i\} = \text{Max}(R_i, U)$. Let $x, y \in U$ with $x < y$ consecutive (i.e. there is no $z \in U$ with $x < z < y$). Then

$$xP_My \iff |\{i \in N : t_i \leq x\}| > |\{i \in N : t_i \geq y\}|.$$

By transitivity of P_M, it follows

$$x \text{ is Condorcet winner.} \iff |\{i \in N : t_i \leq x\}| > \frac{n}{2} \land |\{i \in N : t_i \geq x\}| > \frac{n}{2}.$$

In particular, there must be at least one voter whose top choice is the Condorcet winner. We can visualize the alternatives on a line ordered by $\leq$. Then we may find the unique Condorcet winner by determining a *median voter*; i.e. a voter whose top choice has more than half of the other top choices "left" to it (including the same choice) and the same must hold for the "right". So by sorting the voters according to their top choices, the $\frac{n+1}{2}$-th voter is a median voter and their top choice is the Condorcet winner.
In the domain D_{SP}, $\text{Max}(R_M, -)$ is known as **median voting** and it satisfies strategyproofness and participation by Theorem 2.4.

It is possible in polynomial time to determine if for a given preference profile, there exists some linear order $\geq$ on U, such that the preference profile is single-peaked with respect to $\geq$.

2.5 Value Restriction

Definition 2.9. A domain D is **value-restricted**, if for each $x, y, z \in U$, there is some alternative, say x, such that one of the following conditions is true:

- x is never the worst alternative ($\forall\, R \in D$: $x > y$ or $x > z$).
- x is never the best alternative ($\forall\, R \in D$: $y > x$ or $z > x$).
- x is never the middle alternative ($\forall\, R \in D$: $(x > y \land x > z)$ or $(y > x \land z > x)$).

Note that the first condition corresponds to $U' := \{x, y, z\}$ being single-peaked with linear order $y > x > z$ or $z > x > y$.
In particular, a value-restricted domain can never contain a Condorcet cycle. The domain of single-peaked preferences D_{SP} is value restricted for the corresponding linear order.
Checking whether a domain is value-restricted can be done in polynomial time by simply checking all triples $x, y, z \in U$.

Theorem 2.10. R_M is transitive in domain D if and only if D is value-restricted.

Therefore, if a domain is value-restricted, $\text{Max}(R_M, -)$ constitutes a SCF satisfying all of Arrow's axioms, strategyproofness and participation. Lec 6 2021-11-30

3 Ignoring Consistency

In this section, we try to resolve Arrow's impossibility theorem (1.38) by ignoring consistency. We do this by essentially fixing the set of alternatives and instead varying the number of voters. Throughout this section, we do not assume that the number of voters is odd.

Definition 3.1. For $k \leq |U|$, a **score vector of dimension** k is just a vector $s \in \mathbb{R}^k$. For $A \in F(U)$, any $R \in R(U)|_A$ corresponds precisely to a permutation on A, so the action of $R|_A$ on s defines a function

$$R(U) \times \{A \in F(U) : |A| = k\} \to \mathbb{R}^k,$$

which by definition assigns the first-ranked candidate s_1 points, the second-ranked candidate s_2 points and so on. By summing the results, we obtain a function

$$\phi_s \colon R(U)^n \times \{A \in F(U) : |A| = k\} \to \mathbb{R}^k.$$

Furthermore, for any $A \in F(U)$ with $|A| = k$, a vector $\mathbb{R}^k$ induces a complete and transitive order on A, which in turn gives rise to the subset of maximizers in A w.r.t. this order. A collection $\{s_i \in \mathbb{R}^i : i \in \{1, \dots, |U|\}\}$ of score vectors thus defines a SCF

$$R(U)^n \to \{\text{choice function } S \colon F(U) \to F(U)\},$$

which is called **scoring rule**.

Since the scoring vectors belonging to different dimensions do not have to be related in any way, scoring rules are generally not rationalizable. After all, the goal of this section is to ignore consistency.

Example 3.2. (a) The *(narrowed) Borda's rule* is the scoring rule induced by the collection of score vectors $s_i = (i-1, \dots, 0)$.

(b) The *plurality rule* is the scoring rule induced by the collection of score vectors $s_i = (1, 0, \dots, 0)$.

(c) The *veto rule* (*anti-plurality*) is the scoring rule induced by the collection of score vectors $s_i = (1, \dots, 1, 0)$, which establishes precisely those alternatives that are last-ranked the fewed number of times as the winners.

Note that scoring rules are invariant under positive affine transformations of the score vectors.

Proposition 3.3. A scoring rule is monotonic if and only if every score vector s is monotonically decreasing (i.e. $s_{i-1} \geq s_i$).

Definition 3.4. A monotonic scoring rule is called **non-trivial**, if not all components are the same; i.e. if there exists $k \leq |U|$, such that $s_{k,1} > s_{k,k}$.

Definition 3.5. A SCF ϕ is a **Condorcet extension**, if $\phi(R, A) = \{x\}$, whenever x is the Condorcet winnner in A w.r.t. R.

Definition 3.6. For a given preference relation $R \in R(U)$, we write

$$n_{x,y}(R) := |\{i \in N : x \geq_i y\}|.$$

Note that many SCFs actually do not depend on the whole preference relation $R \in R(U)^n$, but only on the majority relation (*C1*). By our oddness assumption, the corresponding directed graph is a tournament graph and each edge has the nonnegative weight $|n_{x,y}|$ assigned to it. Then a SCF could also depend on the weighted graph, instead of just the unweighted version (*C2*). This terminology is made precise in the following definition.

Definition 3.7. A SCF ϕ is called

- **C1**, if it only depends on the majority relation; i.e. if it descends to a function

$$\{M : M \text{ is majority relation induced by some } R \in R(U)^n\} \times F(U) \to F(U).$$

- **C2**, if it if it decends to a function

$$\{(x, y, n_{x,y}(R)) : x, y \in A, R \in R(U)^n\} \times F(U) \to F(U)$$

  and is not C1. The set of triples can also be interpreted as a matrix (the adjacency matrix of the weighted graph).

- **C3**, if it is neither C1 nor C2.

Note that a SCF that is C1 or C2 induces a collection of SCFs (one for each number of voters), because the weighted majority graph is defined for any number of voters.

Example 3.8.

(a) **Copeland's rule** ϕ is the C1 SCF defined via

$$\phi(R, A) \coloneqq \arg\max_{x \in A} |\{y \in A : xP_M y\}|.$$

In words, it chooses the vertices with maximal degree in the tournament graph.

(b) The (narrowed) Borda's rule is C_2, as the score of a candidate x is

$$s(x) = \sum_{i \in N} |\{y \in A : x >_i y\}| = \sum_{y \in A \setminus \{x\}} n_{x,y}.$$

(c) The **Maximin rule** ϕ is the C2 SCF defined via

$$\phi(R, A) \coloneqq \arg\max_{x \in A} \min_{y \in A \setminus \{x\}} n_{x,y}.$$

Roughly, this function chooses those candidates which do not loose "too significantly" to any other candidate.

Proposition 3.9. Borda's rule is not a Condorcet extension whenever there are at least three alternatives.

Proof. Consider e.g.

2	1	1	1
a	b	b	c
b	c	a	a
c	a	c	b

or

3	2
a	b
b	c
c	a

□

In fact, more is true.

Theorem 3.10. No scoring rule is a Condorcet extension whenever there are at least three alternatives.

Smith showed the following theorem in 1973.

Theorem 3.11. A Condorcet winner is never the alternative with the lowest Borda score and a Condorcet loser is never the alternative with the highest Borda score. Borda's rule is the only scoring rule that never ranks a Condorcet winner last or a Condorcet loser first.

Example 3.12. (a) **Black's rule** selects the Condorcet winner if one exists and the Borda winners otherwise.

(b) **Baldwin's rule** removes the alternatives with the lowest Borda scores, recomputes Borda scores and then iterates this procedure until no more deletions are possible. Note that Baldwin's rule is not monotonic though.

Definition 3.13. A collection of SCF $(\phi_n)_{n\in\mathbb{N}}$ (we will just write ϕ for ϕ_n) satisfies **reinforcement**, if for all $A \in F(U)$, disjoint subsets $N', N'' \subset N$ containing at least two voters and all $R' \in R(U)^{|N'|}$, $R'' \in R(U)^{|N''|}$ we have

$$\phi(R', A) \cap \phi(R'', A) \neq \emptyset \quad \Longrightarrow \quad \phi(R', A) \cap \phi(R'', A) = \phi(R' \cup R'', A).$$

This property is somewhat analogous to rationalizibility.

Definition 3.14. A SCF ϕ is a **composed scoring rule** if it is a composition of finitely many (and at least one) scoring rules f_i; i.e.

$$\phi(R) = f_1(R) \circ f_2(R) \circ \cdots \circ f_k(R).$$

In words, a composed scoring rule applies a scoring rule to narrow down the set of winners and may perform this procedure iteratively.

The significance of composed scoring rules is underlined by the following theorem.

Theorem 3.15. A neutral and anonymous SCF is a composed scoring rule if and only if it satisfies reinforcement.

This means that reinforcement is the defining property of scoring rules.

Definition 3.16. A SCF ϕ satisfies **cancellation**, if for all preference profiles $R \in R(U)^n$ whose majority relation is the all-relation, we have $\phi(R, -) = \mathrm{id}_{F(U)}$.

Young showed in 1974.

Theorem 3.17. Borda's rule is the only SCF satisfying neutrality, Pareto-optimality (it is even enough to assume Pareto-optimality for profiles with only one voter), reinforcement and cancellation.

The following theorem shows that the concepts of Condorcet and Borda are fundamentally incompatible.

Theorem 3.18. No Condorcet extension satisfies reinforcement for at least three alternatives.

4 Kemeny's Rule

Lec 7 2021-12-09

We now assume that there exists a "true" cumulative preference ranking over the alternatives (e.g. in a jury decision there is only "guilty" or "not guilty" and one of those must be true). Then we can see the preferences of the voters as estimations of the true preference ranking. Inspired by the quote

> "Democracy is the recurrent suspicion that more than half the people are right more than half the time."

from E.B. White, it might be reasonable to assume that voters are right more often than wrong; i.e. that when deciding whether $a > b$ or $b < a$ a voter has a probability $\frac{1}{2} < p \leq 1$ to select the "correct" preference.
This idea is made precise in the following.

Definition 4.1. A SCF is a **maximum likelihood SCF** for a given $p \in (\frac{1}{2}, 1]$, if it yields all alternatives that are most likely to be top-ranked in the "true" ranking.

Condorcet showed the following result in 1785.

Theorem 4.2. Majority rule is the maximum likelihood SCF for two alternatives and any $p \in (\frac{1}{2}, 1]$.

It can be shown that Borda's rule is the maximum likelihood SCF for some $p \in (\frac{1}{2}, 1]$ that is sufficiently close to $\frac{1}{2}$.

Definition 4.3. A **social preference function (SPF)** is a function $\phi\colon R(U)^n \to F(R(U))$.

A SPF assigns to a preference profile a set of preference relations on U, so it is a set-valued version of a SWF. The notions defined for SCFs directly translate to SPFs. Note that if ϕ is anonymous, then $R(U)^n$ can be identified with $F(R(U))$, so a SPF is an endomorphism in that case.

Definition 4.4. Kemeny's rule is the SPF

$$\phi(R) \coloneqq \arg \max_{S \in R(U)} \sum_{i \in N} |S \cap R_i|.$$

Kemeny's rule yields all rankings that maximize pairwise agreement. For two relations R_1, R_2 on U, the absolute value of the symmetric difference $|R_1 \Delta R_2|$ is called *Kenoall-Tau difference* of R_1 and R_2.

Theorem 4.5. Kemeny's rule is the maximum-likelihood SPF for any $p \in (\frac{1}{2}, 1]$.

If we consider the complete, weighted majority graph induced by a preference profile, then a Kemeny ranking is just an acyclic subgraph with maximum (accumulated) weight. Note that if we allowed Kemeny's rule to return non-transitive relations, then the majority relation R_M would be the unique ranking that maximizes the Kemeny "score".
Therefore, Kemeny's rule can also be computed from the weighted majority graph by finding a collection of edges with minimal accumulated weight, such that when they are "inverted", they make the graph acyclic.

Lemma 4.6. Let $G = (V, E)$ be a directed graph and $E' \subset E$. G can be made acyclic by inverting a subset of edges in E' if and only if $(V, E \setminus E')$ is acyclic.

The "subset of edges" part is needed, since inverting all edges in a 3-cyclic graph is not acyclic. The lemma tells us that instead of inverting edges, we can just remove them. Lec 8 2021-12-14

Theorem 4.7 (McGarvey's Theorem). For every majority graph $G = (V, E)$ with weight 1 on every edge, there exists a preference profile R with an odd number of voters, such that the strict majority relation P_M equals E.

Proof. Define the preference relation R_1 of the first voter arbitrarily. For every edge $(x, y) \in E$ with $(x, y) \notin R_1$, call the other alternatives $a_1, \ldots, a_l$ and add two voters

1	1
x	a_l
y	$\ldots$
a_1	a_1
$\ldots$	x
a_l	y

which by construction enforce this edge while not changing any of the other edges. The resulting preference profile R has the desired properties. □

5 Expansion Consistency

Our next approach is to only require expansion consistency, which is motivated by the following theorem.

Theorem 5.1. There is no SCF satisfying anonymity on two elements, neutrality on two elements, positive responsiveness on two elements and α.

Proof. In a Condorcet cycle with three alternatives $\{a, b, c\}$, some element, say a, has to be in the chosen set. But by α, it also has to be contained in the winner sets of $\{a, b\}$ and $\{a, c\}$, which contradicts May's theorem (Theorem 1.33). □

5.1 The Top Cycle

Definition 5.2. A **dominant set of alternatives** is a set A such that any element in A is greater than any element that is not in A w.r.t. the strict majority relation.
For $A \in F(U)$ and the majority relation R_M, we write $\text{Dom}(A, R_M)$ for the set of dominant sets.

Note that $\text{Dom}(A, R_M)$ is totally ordered by set inclusion and thus contains a smallest element. This means that every tournament contains a unique minimal dominant set of alternatives.

Definition 5.3. The minimal dominant set of alternatives is called the **top cycle**.

Clearly, the top cycle is a Condorcet extension.
The notion of top cycles induces a SCF with some nice properties.

Definition 5.4. A SCF ϕ is **finer** than another SCF ϕ', if $\phi \subset \phi'$.

Definition 5.5. A SCF ϕ is called **binary**, if the restriction of ϕ onto two element sets uniquely determines the function; i.e. if $\phi(R)|_T = \phi(R')|_T$ for all two element sets T, then $\phi(R) = \phi(R')$.

Intuitively, for a binary SCF, the choices from larger sets only depend on the choices from the two element sets. In particular, any rationalizable SCF is binary.

Definition 5.6. A SCF is called **majoritarian** if it is

- anonymous,

- neutral,
- positive responsive on two element sets,
- binary.

This implies that a majoritarian SCF only depends on the base relation, which must be the majority relation (Theorem 1.33). On two element sets it chooses according to majority. In particular, any majoritarian SCF is C1. As any C1 SCF, we may view a majoritarian SCF $\phi(-, A)$ for $A \in F(U)$ as a function assigning a set of winners to an oriented tournament graph.
The main difference is that a C1 SCF need not be anonymous or neutral and can choose "arbitrarily" on two element sets. A majoritarian SCF satisfies three of Arrow's conditions, namely non-dictatorship, IIA and Pareto-optimality on two element sets.

Theorem 5.7. The top cycle SCF is the finest majoritarian SCF that satisfies β^+.

In particular, on two element sets, the top cycle SCF is precisely majority rule.

Definition 5.8. For an asymmetric complete relation on a set X (and thus also for an oriented tournament graph), the upper set $D(x) :=\uparrow x$ is called **dominion** (or **set of successors**) and the lower set $\overline{D}(x) :=\downarrow x$ is the **set of dominators** (or **predecessors**) of $x \in X$.
We write $D^k(x)$ ($D^0(x) = \{x\}$) for the set of k-times iterated dominion and $\overline{D}(x)$ ($\overline{D}^0(x) = \{x\}$) for the set of k-times iterated dominators of x. Furthermore, we define

$$D^*(x) := \bigcup_{k=0}^{\infty} D^k(x) \quad \text{and} \quad \overline{D}^*(x) := \bigcup_{k=0}^{\infty} \overline{D}^k(x).$$

In the corresponding tournament graph, the set $D^k(x)$ contains precisely those vertices which can be reached from x by using at most k edges and $D^*(x)$ is the connected component of x. Analogously for $\overline{D}^k$ and $\overline{D}^*$.
Note that we obtain a partition $X = D(x) \cup \{x\} \cup \overline{D}(x)$. Lec 9

Starting with an alternative $a \in A$, any dominant set containing a must also contain $\overline{D}^*(a)$ and this is the smallest dominant set containing a. Because $\mathrm{Dom}(A, R_M)$ is totally ordered by inclusion, this implies 2021-12-21

$$\mathrm{Dom}(A, R_M) = \left\{\overline{D}^*(x) : x \in A\right\}.$$

This gives rise to the following simple algorithm: Start with the "working set" $B := \{a\}$ and iteratively add all alternatives that dominate some alternative in B until no more such alternatives exist. By applying this algorithm to every $a \in A$ once, we can compute $\mathrm{Dom}(A, R_M)$ in $\mathcal{O}(|A|^3)$.

The Copeland winners (i.e. those alternatives with a maximal number of outgoing edges) are always contained in the top cycle. This is true, because for any alternative $a \in A$ not contained in the top cycle, any element of its domain $D(a)$ is also dominated by any element of the top cycle. This observation allows us to apply the above algorithm to a Copeland winner and gives a runtime of $\mathcal{O}(|A|^2)$.

Another idea to solve the case when the base relation is not transitive is to consider the transitive closure R_M^* of the majority relation R_M. In our notation, this means that $x > y$ w.r.t. R_M^* if and only if $y \in D^*(x)$, which in turn is equivalent to $x \in \overline{D}^*(y)$. But this idea just leads to the top cycle, as the following theorem shows.

Theorem 5.9. The SCF induced by the transitive closure R_M^* of the majority relation R_M is the top cycle SCF.

Proof. This follows by observing that

$$\begin{aligned} x \in \mathrm{Max}(R_M^*, A) &\iff \forall\, y \in A\colon x \geq_{R_M^*} y \iff D^*(x) = A \\ &\iff \forall\, y \in A\colon x \in \overline{D}^*(y) \iff \forall\, B \in \mathrm{Dom}(A, R_M)\colon x \in B. \end{aligned}$$

□

This gives rise to an alternative algorithm to compute the top cycle by first determining the strong components of the strict majority graph and then choosing the maximal one.
In general, the top cycle is rather large. A bad consequence of this is that the alternatives returned are not necessarily Pareto-optimal. In the following example, the top cycle consists of all alternatives, but the alternative c is Pareto-dominated by b.

1	1	1
a	b	d
b	c	a
c	d	b
d	a	c

5.2 The Uncovered Set

Definition 5.10. Let R be a asymmetric complete relation (an oriented tournament graph). The **cover relation** of R is defined as follows:

$$x > y \;:\iff D(y) \subset D(x) \iff \overline{D}(x) \subset \overline{D}(y).$$

and in that case, we say that x **covers** y.

Any cover relation is a partial order and its asymmetric part is a subrelation of the original relation.
Applying the cover relation to the strict majority relation gives rise to another SCF. This approach is similar to the idea using the transitive closure in the previous section, but instead of enlarging the majority relation, we shrink it so that it becomes transitive.

Definition 5.11. The **uncovered set** consists of all alternatives that are maximal w.r.t. the cover relation C of the strict majority relation. The corresponding SCF $\mathrm{UC}(A, P_M) :=$ $\mathrm{Max}(C, A)$ is the **uncovered set SCF**.

This is a Condorcet extension.
It is helpful to observe that

$$a \in \mathrm{UC}(A) \iff \nexists x \in A\colon x >_C a \iff \forall\, x \in \overline{D}(a)\; \exists\, y \in D(a)\colon y >_{P_M} x.$$

In other words, for $a \in \mathrm{UC}(A)$, every edge $b \to a$ (i.e. b dominates a) has a corresponding cycle $b \to a \to c \to b$ for some $c \in A$. In particular, this implies the following equivalent characterization of the uncovered set.

Theorem 5.12. The uncovered set w.r.t. the strict majority relation on a set of alternatives consists precisely of those alternatives that reach every other alternative in at most two steps.

This gives rise to an algorithm that computes $A := M^2 + M + M^0$ for the adjacency matrix M and only returns those indices for which the corresponding row in A has only nonzero entries.

Moulin proved the following characterization of the uncovered set in 1986.

Theorem 5.13. The uncovered set SCF is the finest majoritarian SCF satisfying γ.

Because β^+ implies γ, Theorem 5.7 shows that the uncovered set is always contained in the top cycle.

In contrast to the uncovered set SCF, the top cycle SCF is Pareto-optimal.

Proposition 5.14. The uncovered set SCF is Pareto-optimal.

One can show that the uncovered set SCF is also the largest majoritarian SCF satisfying Pareto-optimality, so it is the only majoritarian SCF that satisfies Pareto-optimality and γ.

5.3 The Banks Set

Lec 10
2022-01-11

Definition 5.15. Let R be a relation on A. A **transitive subset** of A is a subset $B \subset A$, such that R is transitive on B.

Note that a tournament (A, P_M) is transitive if every subtournament has a Condorcet winner.

Definition 5.16. The **Banks set** is the set consisting of the maximal elements of all inclusion-maximal transitive subsets of the alternatives A w.r.t. the strict majority relation P_M. This induces the **Banks SCF**.

Note that x is in the Banks set of A if and only if it "cannot be extended from above"; i.e. if and only if there exists a transitive subset $B \subset A$, such that $x \in \mathrm{Max}(R, B)$ and there is no $a \in A$ that dominates all $b \in B$.

Definition 5.17. A choice function $S\colon F(U) \to F(U)$ satisfies **strong retentiveness** (ρ^+), if for all $A \in F(U)$ and $x \in A$, we have $S(\overline{D}(x) \cap A) \subset S(A)$.

Intuitvely, a choice function satisfies ρ^+, if the best elements from all dominator sets have to be chosen.

Lemma 5.18. A majoritarian SCF that satisfies γ also satisfies ρ^+.

Theorem 5.19. The Banks SCF is the finest majoritarian SCF satisfying ρ^+.

Combining Theorem 5.13, Lemma 5.18 and Theorem 5.19, it follows that the Banks set is always contained in the uncovered set. With Proposition 5.14, this implies that the Banks SCF is Pareto-optimal.

Furthermore, the Banks SCF is a Condorcet extension. It is even a *strong Condorcet extension*, meaning that it selects a unique winner if and only if that winner is a Condorcet winner. Because the Banks set is included in the uncovered set, which is included in the top cycle, it follows that all three of those SCFs are strong Condorcet extensions.

By starting with a “working set” B consisting of an arbitrary alternative and iteratively adding alternatives that dominate all elements in B, we can compute an element of the Banks set in $\mathcal{O}(|A|^2)$. However, deciding whether a given alternative is contained in the Banks set is NP-complete.

Definition 5.20. A choice function $S\colon F(U) \to F(U)$ satisfies **retentiveness** (ρ), if for all $A \in F(U)$ and $x \in S(A)$, we have $S(\overline{D}(x) \cap A) \subset S(A)$.

By definition, ρ is a weakening of ρ^+.

5.4 The Tournament Equilibrium Set

Definition 5.21. Let $S\colon F(U) \to F(U)$ be an arbitrary choice function.
A set $B \in F(U)$ is S**-retentive**, if for all $x \in B$, we have $S(\overline{D}(x)) \subset B$.

The idea behind the notion of S-retentive sets is that no alternative in the set should be “properly” dominated by an outside alternative.

Definition 5.22. For a choice function $S\colon F(U) \to F(U)$, we construct the new choice function $\tilde{S}\colon F(U) \to F(U)$ that assigns to $A \in F(U)$ the union of all inclusion-minimal S-retentive subsets of A.

Proposition 5.23. The $\tilde{}$ operator only returns choice functions that satisfy ρ.

Example 5.24. For the identity choice function id, $\widetilde{\mathrm{id}}$ is the top cycle.

Definition 5.25. The **tournament equilibrium set** of a tournament TEQ is the unique fixed point of the $\tilde{}$ operator; that is, $TEQ = \widetilde{TEQ}$.

It can be shown that the tournament equilibrium set is always contained in the Banks set.

6 Consistency on Sets

6.1 Set-rationalizability

Lec 11 2022-01-18

We redefine notions like α and γ by only refering to the set of chosen alternatives, instead of the individual elements themselves.

Definition 6.1. A choice function $S\colon F(U) \to F(U)$ is called **set-rationalizable**, if there is a relation $R \subset F(U) \times F(U)$ on $F(U)$, such that there is a unique maximum $\mathrm{Max}(R, A)$ for every $A \in F(U)$ and that $S = \mathrm{Max}(R, -) \circ F$.

The base relation (on U) of $S\colon F(U) \to F(U)$ can be extended to $F(U)$ by defining

$$X > Y \;:\Longleftrightarrow\; X = S(X \cup Y).$$

We recall that a choice function $S\colon F(U) \to F(U)$ satisfies

- α, if

$$\forall A, B \in F(U),\ x \in A \cap B\colon \quad x \in S(A \cup B) \Rightarrow x \in S(A) \wedge x \in S(B).$$

- γ, if

$$\forall A, B \in F(U),\ x \in A \cap B\colon \quad x \in S(A \cup B) \Leftarrow x \in S(A) \wedge x \in S(B).$$

We transfer these concepts to our new setting.

Definition 6.2. A choice function $S\colon F(U) \to F(U)$ satisfies

- $\hat{\alpha}$, if

$$\forall A, B \in F(U),\ X \subset A \cap B\colon \quad X = S(A \cup B) \Rightarrow X = S(A) \wedge X = S(B).$$

- $\hat{\gamma}$, if

$$\forall A, B \in F(U),\ X \subset A \cap B\colon \quad X = S(A \cup B) \Leftarrow X = S(A) \wedge X = S(B).$$

Despite the formal similarities, α and $\hat{\alpha}$ are not logically related and the same holds true for γ and $\hat{\gamma}$.

Lemma 6.3. A choice function $S\colon F(U) \to F(U)$ satisfies $\hat{\alpha}$, if and only if for all $V, W \in F(U)$ with $S(V) \subset W \subset V$, we have $S(V) = S(W)$.

Intuitively, the previous lemma states that a choice function satisfying $\hat{\alpha}$ is invariant under removing non-chosen alternatives.
In particular, any SCF satisfying $\hat{\alpha}$ must be idemptotent (in its second argument).

Theorem 6.4. A choice function $S\colon F(U) \to F(U)$ is set-rationalizable if and only if it satisfies $\hat{\alpha}$.

Definition 6.5. A choice function $S\colon F(U) \to F(U)$ is called **stable**, if it satisfies $\hat{\alpha}$ and $\hat{\gamma}$.

Theorem 6.6. A choice function $S\colon F(U) \to F(U)$ is quasi-transitively rationalizable if and only if it satisfies α, $\hat{\alpha}$ and $\hat{\gamma}$.

Proposition 6.7. Every non-trivial monotonic scoring rule violates $\hat{\alpha}$.

The choice functions induced by most social choice functions we discussed so far violate $\hat{\alpha}$, but it is not hard to see that e.g. the top cycle is even stable.

Definition 6.8. Let (A, R) be a complete relation and $p\colon A \to [0,1]$ a probability distribution. p is called **optimal**, if the *utility function*

$$u_p\colon A \to \mathbb{R},\ x \mapsto \sum_{y \in \overline{D}(x)} p(y) - \sum_{y \in D(x)} p(y)$$

maps to the nonnegative real numbers.

Intuitively, an optimal probability distribution puts "more probability" on the alternatives that dominate x than those that are dominated by it for each $x \in A$.

Theorem 6.9. Every complete relation (A, R) admits a unique optimal probability distribution.

Proof. We first note that for any probability distribution p

$$\sum_{x\in A} p(x)u_p(x) = \sum_{x\in A} p(x)\left(\sum_{y\in \overline{D}(x)} p(y) - \sum_{y\in D(x)} p(y)\right)$$
$$= \sum_{x,y\in A,\ x\neq y} p(x)p(y)(2\cdot [\![yP_M x]\!] - 1) = 0.$$

As a consequence, if p is optimal, then $u_p(x) = 0$ for all $x \in A$ with $p(x) > 0$.
We now prove the existence of an optimal probability distribution by induction on $|A|$. The base case $|A| = 1$ is clear, so suppose $|A| > 1$. Aiming for contradiction, assume that for any probability distribution p, $\nu := \min_{x\in A} u_p(x) < 0$. We may assume that p is chosen such that ν is maximal. Then there exists $z \in A$ with $u_p(z) \geq 0$. By the induction hypothesis, there is a probability distribution q with $q(z) = 0$ and $u_q(z) \geq 0$ for all $x \in A \setminus \{z\}$. But then for $\epsilon \in (0, 1)$, $r := (1-\epsilon)p + \epsilon q$ satisfies

$$u_r(x) \geq (1-\epsilon)\nu + \epsilon\cdot 0 > \nu \quad \forall\, x \in A \setminus \{z\}$$

and if ϵ is small enough ($\epsilon < -\frac{\nu}{|u_q(z)|}$) also

$$u_r(z) = (1-\epsilon)u_p(z) + \epsilon u_q(z) > \nu.$$

But this contradicts the maximality of p, which proves the claim.
We now show the uniqueness of optimal probability distributions. Aiming for contradictions, we assume that p and q are optimal for (A, R) and $p \neq q$. By considering a convex combination $\lambda p + (1-\lambda)q$ and $\mu p + (1-\mu)q$ for $\lambda, \mu \in (0, 1)$, $\lambda \neq \mu$, we may assume that p and q have the same support B. Let $r(x) := p(x) - q(x)$. Then

$$\sum_{y\in \overline{D}(x)} r(y) - \sum_{y\in D(x)} r(y) = \sum_{y\in \overline{D}(x)} (p(y) - q(y)) - \sum_{y\in D(x)} (p(y) - q(y))$$
$$= u_p(x) - u_q(x) = 0 - 0 = 0 \quad \forall\, x \in B.$$

and

$$\sum_{x\in B} r(x) = 0$$

hold and constitute a finite homogeneous system of linear equations for $r \in \mathbb{R}^{|B|}$ with integer coefficients. It is known that such a system has a non-zero integer solution $r^* \in \mathbb{Z}^{|B|}$. By repeatedly dividing r^* by 2, we may assume that $r^*(b)$ is odd for some $b \in B$. We extend r^* to all of A by setting it to zero for elements not in B. For $x \in B$, we calculate

$$0 = \sum_{x\in B} r^*(x) = r^*(x) + \sum_{y\in \overline{D}(x)} r^*(y) + \sum_{y\in D(x)} r^*(y) = r^*(x) + 2\cdot \sum_{y\in \overline{D}(x)} r^*(y),$$

so for all $x \in B$, $r^*(x)$ is even, contradicting our assumption. □

Alternatively, the existence can also be shown using the *minimax theorem*. The optimal distribution corresponds to the unique Nash equilibrium (or maximin strategy) of the zero-sum game given by the skew-adjacency matrix of the tournament.

Definition 6.10. The **bipartisan set** consists precisely of those alternatives that are in the support of the optimal probability distribution on (A, R_M). This induces the **bipartisan set SCF**.

It can be shown that the bipartisan set always contains an odd number of alternatives and is contained in the uncovered set. In particular, the bipartisan SCF is a Condorcet extension.

Theorem 6.11. The bipartisan set SCF is stable and satisfies strong monotonicity (i.e. it is invariant under the weakening of unchosen alternatives).

6.2 Gibbard-Satterthwaite Impossibility Theorem

Lec 12 2022-01-25

Definition 6.12. A SCF ϕ is called **strongly monotonic** if for all $R, R' \in R(U)^n$, $A \in F(U)$ with $R = R'$ except for some $i \in N$, $x \in A \setminus \phi(R, A)$, $y \in A$, where $(x, y) \in R$ but $(y, x) \in R'$, we have $\phi(R, A) = \phi(R', A)$.

Intuitively, a strongly monotonic SCF is one that is invariant under weakening unchosen alternatives. For resolute SCFs, the property allows us to apply any changes to the preferences of the voters without changing the outcome, as long as the set of candidates that are less prefered than the winner is kept or enlarged.
In particular with a strongly monotonic SCF one can arbitrarily change the preferences of the voters over the unchosen alternatives.

Lemma 6.13. Any strongly monotonic SCF is monotonic.

Muller and Satterthwaite showed the following theorem in 1977.

Theorem 6.14. A resolute SCF is strategyproof if and only if strongly monotonic.

Proof. Fix a set $A \in F(U)$ and for simplicity we omit it from the notation.
Let ϕ be a resolute SCF and $R, R' \in R(U)^n$ with $R|_{N\setminus\{i\}} = R'|_{N\setminus\{i\}}$. Note that strategyproofness is equivalent to

$$\forall\, R_i, R_i' \in R(U)\colon \phi(R) = \{x\} \neq \{y\} = \phi(R') \Rightarrow xP_iy \wedge yP_i'x. \qquad (*)$$

If ϕ satisfies $(*)$, $\phi(R) = \{x\}$ and for all $z \in A$ with xP_iz we have $xP_i'z$, then $\phi(R') = \{x\}$, because the opposite would contradict $(*)$.
For the other direction, we suppose that $\phi(R) = \{x\} \neq \{y\} = \phi(R')$ and yP_ix or $xP_i'x$. By potentially interchanging R and R', we may assume that yP_ix. Define R'' to be

$$R_i'' = R_i \setminus \{(z, y) : z \neq y\} \cup \{(y, z) : z \neq y\},$$

which amounts to moving y to the top of the preference profile. Now we can conclude that ϕ cannot be strongly monotonic, since otherwise $\phi(R'') = \{x\}$ and $\phi(R'') = \{y\}$. $\square$

Oftentimes it is easier to work with strong monotonicity than strategyproofness.

Of course, for two alternatives monotonicity and strong monotoncity are equivalent, so we obtain the first part of Theorem 1.32 as a corollary.

Theorem 6.15. No resolute Condorcet extension satisfies strong monotonicity for at least three voters and alternatives.

Proof. Let ϕ a resolute function that satisfies strong monotonicity and let R be a Condorcet cycle with three voters and alternatives. Then by resoluteness, ϕ produces a unique winner for R. Using strong monotoncity we can move one of the unchosen alternatives down in such a way that a different alternative emerges as the Condorcet winner without changing the result. This shows that a SCF satisfying strong monotoncity cannot be a Condorcet extension on three voters and alternatives.
For the general case, just add alternatives to the bottom of the preference profiles and for each pair of additional voters (we are assuming that the number of voters is odd, even though the statement can also be proven for an even number of voters) we may add an arbitrary preference profile and its inverse (reversed version). □

Definition 6.16. A SCF ϕ is called **non-imposing** if for each $A \in F(U)$, $\phi(-, A): R(U)^n \to F(U)$ satisfies $\{\{x\} : x \in A\} \subset \text{im}(\phi(-, A))$.

This is a very weak condition. For example, Pareto-optimality implies the non-imposing condition. The following lemma shows that for strongly monotonic SCFs, the two notions are actually equivalent.

Lemma 6.17. A non-imposing and strongly monotonic SCF is Pareto-optimal.

We can now prove another impossibility theorem, which is just as important as Arrow's impossibility theorem (Theorem 1.38).

Theorem 6.18 (Gibbard-Satterthwaite Impossibility Theorem). Every non-imposing, strategyproof and resolute SCF over at least three alternatives is dictatorial.

For two alternatives, majority rule satisfies the conditions in the theorem without being dictatorial.

In particular, the theorem directly implies Theorem 6.15.

Theorem 6.19. No resolute Condorcet extension satisfies participation when there are at least 12 voters and 4 alternatives.

6.3 Extending Preference Relations

Lec 13 2022-02-02

Because resolute SCFs are never neutral, we usually prefer to work with non-resolute SCFs. However, in that case, it is hard to decide whether an outcome (which is a set of alternatives) is prefered by some voter to another outcome. Therefore, we need to extend the relations of the voters on U to a relation on $F(U)$. Our first solution to this problem was to consider the power set relation (see Definition 1.28). We can now generalize this notion and introduce another way to lift preferences on alternatives to preferences on sets of alternatives.

Definition 6.20. We call a map $\psi: R(U) \to R(F(U))$ a **preference extension**.
A SCF ϕ is called **ψ-manipulable** by voter $i \in N$, if there exists $A \in F(U)$ and two preference profiles $R, R' \in R(U)^n$ whose j-th components agree for all $j \neq i$, such that i strictly prefers $\phi(R', A)$ to $\phi(R, A)$ w.r.t. $\psi(P_i)$.
A SCF is called **ψ-strategyproof**, if it is not ψ-manipulable by any voter.

Thus Definition 1.28 is a special case, where ψ assigns to $R \in R(U)$ the power set relation $\mathcal{P}(R)$.

Definition 6.21. Kelly's preference relation R^k on $F(U)$ induced by a preference relation R on U is defined by

$$X \geq_{R^k} Y \;:\Longleftrightarrow\; \forall\, x \in X, y \in Y : x \geq_R y.$$

We denote the strict part of R^k by P^k.
The corresponding preference extension is called **Kelly's preference extension**.

By definition, we have

$$X >_{P^k} Y \;:\Longleftrightarrow\; (\forall\, x \in X, y \in Y : x \geq_R y) \wedge (\exists\, x \in X, y \in Y : x >_P y).$$

If the given preference relation P is strict (e.g. the preference relation of a voter) with symmetric completion R, then two comparable sets w.r.t. P^k can have at most one element in their intersection.

Clearly, a SCF that is manipulable is also R^k-manipulable.

Theorem 6.22. Every strongly monotonic SCF is R^k-strategyproof.

Theorem 6.23. Every SCF satisfying monotonicity, $\hat{\alpha}$ and IIA is also strongly monotonic.

Proof. Let ϕ be a SCF with the above properties, $R, R' \in R(U)^n$ identical preference profiles except that xP_iy instead of yP'_ix and $x \notin \phi(R)$. Then $x \notin \phi(R')$ by monotonicity, so IIA implies for all $A \in F(U)$

$$\phi(R, A) \overset{\hat{\alpha}}{=} \phi(R, A \setminus \{x\}) = \phi(R', A \setminus \{x\}) \overset{\hat{\alpha}}{=} \phi(R', A). \qquad \square$$

For example, the top cycle SCF and bipartisan set SCF are strongly monotonic by the above theorem and thus in particular R^k-strategyproof by Theorem 6.22. Furthermore, it can be shown that the uncovered set SCF is R^k-strategyproof.

We now generalize Definition 2.3.

Definition 6.24. Let $\psi\colon R(U) \to R(F(U))$ be a preference extension and $n \in \mathbb{N} \cup \{\infty\}$. A collection of SCF

$$\left\{\phi_j\colon R(U)^j \to \{\text{choice function } S\colon F(U) \to F(U)\} : j \in \{2, \ldots, n\}\right\}$$

can be **ψ-manipulated by strategic abstention**, if there exists $j \in \{3, \ldots, n\}$, $R \in R(U)^j$ and $A \in F(U)$, such that $\phi_{j-1}(R_{-j}, A) > \phi_j(R, A)$ w.r.t. $\psi(P_j)$. Here $R_{-j} \in R(U)^{j-1}$ denotes the preference profile R, but with the j-th component removed.
Otherwise it satisfies **ψ-participation**.

Definition 6.25. We call a collection

$$\left\{\phi_j\colon R(U)^j \to \{\text{choice function } S\colon F(U) \to F(U)\} : j \in \mathbb{N}, j \geq 2\right\}$$

of C1 (or C2) SCF ϕ_j **compatible**, if $\phi_i(R) = \phi_j(R')$ whenever $R \in R(U)^i$ induces the same majority graph (or weighted majority graph) as $R' \in R(U)^j$.

Theorem 6.26. Let ψ be a preference extension. A compatible collection

$$\big\{\phi_j \colon R(U)^j \to \{\text{choice function } S \colon F(U) \to F(U)\} : j \in \mathbb{N}, j \geq 2\big\}$$

of majoritarian ψ-strategyproof SCFs ϕ_j also satisfies ψ-participation.

Proof. By contraposition, we assume that there exist $n \in \mathbb{N}$, $R \in R(U)^n$ and $i \in N$ such that $\phi_{n-1}(R_{-i}) > \phi_n(R)$ w.r.t. $\psi(P_i)$. We want to show that at least one of the ϕ_j is not ψ-strategyproof.
Consider the "doubled" preference profile $R' \in R(U)^{2n}$ in $2n$ voters, in which the first j voters have the same preference as they have in R and the $(n+j)$-th voter has the same preferences as voter j for $j \in \{1, \dots n\}$. Additionally, we define the preference profile R'' like R', but invert the preferences of the i-th voter. This means that voter i and $n+i$ "cancel each other out" in the resulting majority graph.
Because the collection is compatible, all SCFs are majoritarian and by assumption, we have

$$\phi_{2n}(R'') = \phi_{n-1}(R_{-i}) > \phi_n(R) = \phi_{2n}(R') \quad \text{w.r.t. } \psi(P_i),$$

which shows that ϕ_{2n} is ψ-manipulable by voter i. □

In particular, the top cycle, the uncovered set and the bipartesian set SCFs satisfy R^k-participation.

We now define another preference extension, which is stronger than Kelly's preference extension in the sense that more sets are comparable.

Definition 6.27. Fishburn's preference relation R^f on $F(U)$ induced by a preference relation R on U is defined by

$$X \geq_{R^f} Y \;:\Longleftrightarrow\; (\forall\, x \in X \setminus Y, y \in Y \colon x \geq_R y) \wedge (\forall\, x \in X, y \in Y \setminus X \colon x \geq_R y).$$

The corresponding preference extension is called **Fishburn's preference extension**.

As mentioned, Fishburn's preference relation extends that of Kelly; that is, $R^k \subset R^f$.

We adapt the notion of non-imposing SCFs (Definition 6.16) to sets.

Definition 6.28. A SCF is **set non-imposing**, if for every $A \in F(U)$ and $X \in F(A)$, there exists some $R \in R(U)^n$ such that $f(R, A) = X$.

Theorem 6.29. The top cycle SCF is the only majoritarian SCF that is R^f-strategyproof and set non-imposing.

7 Probabilistic Social Choice

Our approach using SCFs has lead to many impossibility results, which is rather unsatisfying. It thus seems natural to consider another notion of "decision functions"; namely those which are allowed to return a probability distribution over the alternatives that can be used to decide the winner.
By fixing the set of feasible alternatives $A \subset U$, any SCF ϕ gives rise to a map $R(U)^n \to F(A)$. Furthermore, if ϕ satisfies IIA, this is really a map $R(A)^n \to F(A)$ so by replacing U with A we may assume that $U = A$.
We now generalize this type of function by allowing the returned value to be a probability vector whose i-th component describes the probability that the i-th alternative should win.

Definition 7.1. A **social decision scheme** (**SDS**) is a function $\phi\colon R(U)^n \to P$, where P denotes the simplex

$$P := \left\{ p \in \mathbb{R}^m, p_i \geq 0 \;\forall\, i \in \{1, \ldots, m\}, \sum_{j=1}^{m} p_j = 1 \right\} \subset \mathbb{R}^m$$

of probability distributions on the m candidates.

In order to decide whether a voter prefers a given probability distribution to another, we define the following concept.

Definition 7.2. A function $u\colon U \to \mathbb{R}$ is a **utility function** w.r.t. a given relation R on U if it is strictly monotonic.

If the relation on U is complete (e.g. the preference relation of a voter), this just means that

$$u(x) \geq u(y) \iff x \geq_R y \quad \forall\, x, y \in U.$$

For a given preference profile $R \in R(U)^n$, any SDS yields a probability distribution p on the universe of alternatives U, so we obtain a probability space $(U, \mathcal{P}(U), p)$. Then any function $u\colon U \to \mathbb{R}$ is a random variable. If u is a utility function w.r.t. the preferences of some voter i, then the expected value

$$\mathbb{E}[u] = \sum_{a \in U} p_a u(a)$$

of u can be used to quantify how much the voter likes this particular outcome of the election (namely drawing a random candidate w.r.t. p in order to determine the winner). This is made precise in the following definition.

Definition 7.3. Let U be the set of alternatives, p a probability distribution on U, $i \in N$ a voter, $R \in R(U)^n$ a preference profile and $u : U \to \mathbb{R}$ a utility function w.r.t. R_i. The **expected utility** $u(i, p)$ of voter i is the expected value $\mathbb{E}[u]$ of u w.r.t. p.

We can now adapt Definition 1.28 to SDSs.

Definition 7.4. A SDS ϕ is **manipulable** if there exist $R, R' \in R(U)^n$ with $R|_{N\setminus\{i\}} = R'|_{N\setminus\{i\}}$, $i \in N$ and a utility function $u\colon U \to \mathbb{R}$ w.r.t. R_i, such that $u(i, \phi(R')) > u(i, \phi(R))$.
ϕ is **strategyproof**, if it is not manipulable by any voter.

Just like before, the intuition is that a manipulable SDS is one where a voter might be able to submit a "fake" version of their preferences in order to obtain an outcome they prefer, provided they know how the other voters vote.

Theorem 7.5. Every SDS that puts probability 1 on a Condorcet winner can be manipulated if there are at least three voters and alternatives.

Proof. We only prove the special case of three voters and alternatives.
Let ϕ be a SDS that puts probability 1 on a Condorcet winner and R' the preference profile given by a Condorcet cycle over three voters and alternatives:

1	1	1
a	b	c
b	c	a
c	a	b

Without loss of generalize we may assume that $\phi(R')(a) > 0$. Let R be the preference profile obtained from R' by interchanging b and c in the first column. Then $\phi(R)(c) = 1$, because c is a Condorcet winner in R. Let $u\colon U \to \mathbb{R}$ be a utility function such that

$$u(a) = 1, \quad u(c) \in (0, \phi(R')(a)), \quad u(b) = 0.$$

Then this constitutes a utility function w.r.t. R_1 and by construction

$$u(1, \phi(R')) \geq \phi(R')(a) > u(c) = u(1, \phi(R)),$$

so ϕ is manipulable by voter 1. □

Definition 7.6. Let p be a probability distribution on the voters. The SDS that picks a voter at random w.r.t. p and then returns that voter's favorite alternative is called **random dictatorship SDS**.

Next we adapt Definition 6.16 to this probabilistic setting.

Definition 7.7. A SDS is called **non-imposing**, if its image contains all degenerate probability distributions (i.e. those which assign probability 1 to some alternative).

We can now state an impossibility result similar to that of Gibbard-Satterthwaite (Theorem 6.18).

Theorem 7.8. Every non-imposing, non-manipulable SDS is a random dictatorship for at least three voters.

We can obtain a SDS from any scoring rule (assuming all scores are non-negative and some alternative always has a positive score) by assigning an alternative its score divided by the sum of the scores of all alternatives.
For example, the SDS obtained from Borda's rule is not non-imposing.

Definition 7.9. For a preference profile $R \in R(U)^n$, let M be the majority margin matrix $M_{i,j} = n_{x,y} - n_{y,x}$. A probability distribution p is called **maximal** if $M^T \cdot p \geq 0$.

This can be thought of as a randomized Condorcet winner or a C2 version of optimal distributions (Definition 6.8). It can be shown that a maximal lottery always exists, is unique and can be computed efficiently using linear programming.