Collective Choices
Lecture 3: Ranking methods

René van den Brink
VU Amsterdam and Tinbergen Institute

May 2016

Introduction I

In Lectures 1 and 2 we considered preference aggregation and discussed several social choice functions and social welfare fuctions.

We saw that using the Borda rule we can assign to every preference profile a unique weak (social) preference relation by ordering the alternatives by their Borda score.
The same can be done for any scoring rule.

We also considered the Condorcet rule that uses the majority relation.

This majority relation can also be considered as a social preference relation, but it need not be transitive.

Question: How would you choose when the majority relation is not transitive?

Introduction II

In this lecture we apply score functions for directed graphs (digraphs) which assign real numbers to every node in a digraph.

Using these score functions we define social choice functions and social welfare functions by simply ranking the nodes according to their score in an associated digraph.

Score functions for digraphs have many applications:

- ranking alternatives in a preference relation (our main application in this course)
- ranking teams in a sports competition (based on the results of the matches)
- ranking web pages (based on their links)
- ranking positions in a network by their importance, centrality, ...

Introduction III

To stress the general use, we discuss score functions and ranking methods for digraphs.

Introduction IV

Contents

- Directed graphs
- Score functions
- Properties of score functions
- Eigenvector scores
- Application to social choice

Directed graphs I

1. Directed graphs

Let $A = \{a_1, a_2, \ldots a_m\}$ be a fixed finite set of alternatives.

A **directed graph** or **digraph** on the set of alternatives A is a collection of ordered pairs $D \subseteq A \times A$, where $(a, b) \in D$ can be interpreted as 'a weakly defeats b'.

The ordered pairs $(a, b) \in D$ are called **arcs**.

Remark: If $(a, b) \in D$ and $(b, a) \notin D$ then we say that 'a (strictly) defeats b'.

Directed graphs II

Remark: Usually, the set A is called a set of *nodes*, but since we will mainly apply this to social choice situations, we refer to the set A as a set of alternatives.
But notice that A also can be a set of 'teams in a sports competition', 'web pages on the www', 'positions in a network', etc.

Remark: Since in this lecture we take the set of alternatives A fixed, we represent a digraph (A, D) just by its binary relation, and speak about digraph D.

Directed graphs III

Applications:

- **Individual or Social choice:**
  For two alternatives $a, b \in A$, $(a, b) \in D$ means that a is weakly preferred to b.
- **Sports competition**:
  For two teams $a, b \in A$, $(a, b) \in D$ and $(b, a) \notin D$ means that team a has won the match it played against team b.
  For two teams $a, b \in A$, $[(a, b) \in D$ and $(b, a) \in D]$ means that teams a and b played a draw.
- **Web page ranking**:
  For two web pages $a, b \in A$, $(a, b) \in D$ means that there is a link from webpage a to webpage b.

Directed graphs IV

- **Hierarchical networks**
  In a hierarchical network (a, b) means that a is dominant over b. For example, if the hierarchy is a firm structure, then $(a, b) \in D$ can be that manager a is the direct boss of employee b.

Assumption We assume the digraph D to be reflexive, i.e. $(a, a) \in D$ for all $a \in A$.

Score functions I

2. Score functions

Definition

A **score function** on a set of alternatives A is a function σ that assigns a real number $\sigma_a(D)$ to every alternative a in any digraph D on A.

So, $\sigma(D) = (\sigma_1(D), \sigma_2(D), \ldots, \sigma_m(D)) \in \mathbb{R}^m$ where $\sigma_a(D)$ is a measure of the 'power' or 'strength' of alternative $a \in A$ in digraph D.

For preference relations it can be a measure of 'desirability'.

Score functions II

For digraph D on A and alternative $a \in A$, the alternatives in the set

$$Succ_a(D) = \{b \in A \setminus \{a\} \mid (a, b) \in D\}$$

are called the **successors** of a in D. These are the alternatives that 'are weakly defeated' by a.

The alternatives in the set

$$Pred_a(D) = \{b \in A \setminus \{a\} \mid (b, a) \in D\}$$

are called the **predecessors** of a in D. These are the alternatives that 'weakly defeat' a.

Question: Suppose that you know $Succ_a(D)$ for all $a \in A$. Do you know $Pred_a(D)$ for all $a \in D$? And do you know D?

Score functions III

The outdegree and β-scores
The outdegree of alternative a is the number of other alternatives that are weakly defeated by a.

Definition
The **outdegree** of alternative $a \in A$ in digraph D is the number of successors of a in D:

$$out_a(D) = \#Succ_a(D)$$

Note that the outdegree does not take account of 'who' are the successors of a. Only the number of successors matters.
You can say that an alternative gets '1 point' for every alternative it weakly defeats.

Score functions IV

A disadvantage of the outdegree is that, in determining the score of an alternative, it does not take account of the 'strength' of the alternatives it defeats.
We can take account of this by assigning for every alternative b that is weakly defeated by a, $\frac{1}{\#Pred_b(D)}$ point to a.

This yields the following score function.

Definition
The β-**score** of alternative $a \in A$ in digraph D is given by

$$\beta_a(D) = \sum_{b \in Succ_a(D)} \frac{1}{\#Pred_b(D)}$$

Score functions V

The outdegree and β-score can give very different outcomes and rankings of the alternatives.

One way to understand score functions, or understand the difference between different score functions, is to find axiomatizations.
That means finding properties (axioms) that are satisfied by a score function, and only by this score function.

3. Properties of score functions

Dummy property For every digraph D on A and alternative $a \in A$ with $Succ_a(D) = \emptyset$, it holds that $\sigma_a(D) = 0$.

Interpretation: If an alternative has no successors then its score is zero.

Symmetry For every digraph D on A and alternatives $a, b \in A$ such that $Succ_a(D) = Succ_b(D)$ and $Pred_a(D) = Pred_b(D)$, it holds that $\sigma_a(D) = \sigma_b(D)$.

Interpretation: If two alternatives have the same successors and predecessors, then they have the same score.

Question: Suppose D represents a preference relation. Can you consider $\sigma(D)$ as a utility function? If yes, how do you interpret the two properties above?

Properties of score functions II

For digraph D on set of alternatives $A = \{a_1, \ldots, a_m\}$, and alternative $a_k \in A$, the **loss graph** of alternative a_k is the digraph

$$D_k = \{(b, a) \in D \mid a = a_k\}.$$

D_k is the digraph that consists of all arcs where alternative a_k is the successor.

Property

Additivity over loss graphs For every digraph D on $A = \{a_1, \ldots, a_m\}$, it holds that

$$\sigma(D) = \sum_{k=1}^{m} \sigma(D_k).$$

Properties of score functions III

Remark: Remember that $\sigma(D)$, and thus also all $\sigma(D_k),\ k \in \{1, \ldots, m\}$, are m-dimensional vectors.

Proposition

The outdegree and β-score satisfy the dummy property, symmetry and additvity over loss graphs.

These two scores satisfy a different normalization.

Properties of score functions IV

Score normalization For every digraph D on A, it holds that

$$\sum_{a \in A} \sigma_a(D) = \#D.$$

According to score normalization the total number of points to be allocated is the number of 'pairs' ('matches') in D.

Dominance normalization For every digraph D on A, it holds that

$$\sum_{a \in A} \sigma_a(D) = \#\{b \in A \mid Pred_b(D) \neq \emptyset\}.$$

According to dominance normalization the total number of points to be allocated is the number of weakly defeated alternatives.

Properties of score functions V

Proposition

The outdegree satisfies score normalization

The β-score function satisfies dominance normalization.

Properties of score functions VI

Theorem

(i) The outdegree is the unique score function that satisfies the dummy property, symmetry, additvity over loss graphs and *score normalization*.

(ii) The β-score is the unique score function that satisfies the dummy property, symmetry, additvity over loss graphs and *dominance normalization*.

Remark: The difference between the outdegree and β-score is only in the normalization, i.e. the total number of points that is allocated over the alternatives.

Just deciding how many ‘points‘ to allocate is not as innocent as it seems in ranking alternatives (sport teams, web pages, etc.).

It can lead to different rankings, even to a different ‘winner‘.

Properties of score functions VII

As we have seen, a disadvantage of the outdegree of an alternative is that it does not take account of the 'strength' of the alternatives it defeats. The β-score takes this into account.

However, a disadvantage of the β-score is that alternatives can get higher in the ranking by 'losing' instead of 'winning' a pairwise comparison (match).

Question: Do you see why this might happen?

This disadvantage of the β-score can be 'repaired' by letting every alternative also weakly defeat itself, yielding the following score function.

Properties of score functions VIII

Definition

The *modified* β-score of alternative $a \in A$ in digraph D is given by

$$\beta_a^{mod}(D) = \sum_{b \in Succ_a(D) \cup \{a\}} \frac{1}{\#Pred_b(D) + 1}$$

Properties of score functions IX

Proposition

Consider digraph D and alternatives $a, b \in A$ such that $\beta_a^{mod}(D) \geq \beta_b^{mod}(D)$, and alternative $c \in A \setminus \{a\}$.

- If $D' = D \cup \{(a, c)\}$ then $\beta_a^{mod}(D') \geq \beta_b^{mod}(D')$.
- If $D' = D \setminus \{(c, a)\}$ then $\beta_a^{mod}(D') \geq \beta_b^{mod}(D')$.

Interpretation: If alternative a wins one more match (instead of losing it) then it will not do worse in the ranking of the alternatives.

Eigenvector scores I

4. Eigenvector scores

The β- and modified β-score of an alternative a depend on the number of predecessors of its successors.

In this way, we take account of the 'strength' of alternative b that is weakly defeated by a, in determining the score of alternative a.

But if we want to take account of the strength of a's successor b, then the score of b should appear in the score of a.

This can be done using **eigenvectors**.

Eigenvector scores II

Definition

Let A be a set of alternatives. The **transition matrix** of digraph D on A, is the $\#A \times \#A$ matrix Π^D with entries given by

$$\pi^D_{ab} = \begin{cases} \frac{1}{\#Pred_b(D)+1} & \text{if } (a,b) \in D \text{ or } a = b \\ 0 & \text{otherwise.} \end{cases}$$

Then $\beta^{mod}(D) = \Pi^D \mathbf{1}_A$, with $\mathbf{1}_A \in \mathbb{R}^m$ the unit vector given by $(\mathbf{1}_A)_a = 1$ for all $a \in A$.

Question: What are the entries in the column corresponding to alternative a in matrix π^D? What are the entries in the row corresponding to alternative a?

Eigenvector scores III

Proposition

Let D be a digraph on A. Then the matrix Π^D has eigenvalue 1.

Question: How would you use this proposition to define a score function?

Eigenvector scores IV

Let $\lambda \in \mathbb{R}^m$ be an eigenvector of digraph D corresponding to eigenvalue 1, i.e.

$$\Pi^D \lambda = \lambda.$$

Written differently, for every alternative $a \in A$,

$$\lambda_a = \sum_{b \in Succ_a(D) \cup \{a\}} \frac{\lambda_b}{\#Pred_b(D) + 1}.$$

Question: How do you interpret such a vector λ?

Eigenvector scores V

Unfortunately, the eigenvector λ need not be unique. (It is unique upto normalization if the digraph has only one top cycle, see later in Application to social choice functions).

We can single out one of these eigenvectors by the following iterative procedure.

Start with the modified β-score β^{mod}:

$$\beta^1(D) = \beta^{mod}(D)$$

Then, for $t \in \{2, 3, \ldots\}$, we can define the t^{th}**-order** β**-score** β_a^t of alternative $a \in A$ in digraph D on A, iteratively, as

$$\beta_a^t(D) = \sum_{b \in Succ_a(D) \cup \{a\}} \frac{\beta_b^{t-1}(D)}{\# Pred_b(D) + 1}.$$

Eigenvector scores VI

Proposition

For every digraph D on set of alternatives A, the limit

$$\lim_{t \to \infty} \beta^t(D)$$

exists and is unique.
Moreover, it is an eigenvector corresponding to eigenvalue 1.

This gives the following definition.

Definition

The λ-**score function** of digraph D on A is given by

$$\lambda(D) = \lim_{t \to \infty} \beta^t(D).$$

Eigenvector scores VII

Remark: This iterative process can be seen as a Markov process.

Remark: Also the famous Google Page Rank method to rank web pages is based on a limit and eigenvector approach.

5. Application to Social Choice

Consider a finite set of alternatives A, and finite set of agents N.

Short summary of previous lectures 1

A **preference relation** $\succsim_i$ on A represents preferences of agent $i \in N$ over set of alternatives A.

$a \succsim_i b$ means that agent i considers alternative a 'at least as good' as alternative b.

Application to Social Choice II

Preference relation $\succsim_i$ on A is

- **complete** if $a \succsim_i b$ or $b \succsim_i a$, $a \neq b$
- **transitive** if $a \succsim_i b$ and $b \succsim_i c$ implies that $a \succsim_i c$.
- **anti-symmetric** if $[a \succsim_i b$ and $a \neq b]$ implies that $b \not\succsim_i a$,

Application to Social Choice III

A **preference profile** is a tuple $p = (\succsim_i)_{i \in N}$ of individual preference relations on A.

A triple (N, A, p) is a **social choice situation**.

Since we take the set of agents and set of alternatives fixed, we represent a social choice situation just by its **preference profile** p.

A **social choice function** C assigns to every preference profile p a non-empty subset of alternatives: $C(p) \subseteq A$.

Given a social choice function C and a preference profile p, we call $C(p)$ the corresponding **social choice set**.

Application to Social Choice IV

Let $n^p(a, b) = \#\{i \in N \mid a \succ_i b\}$ be the number of agents that consider a to be better than b in profile p.

The **majority relation** of preference profile p is the preference relation $\succsim^p$ given by

$$a \succsim^p b \Leftrightarrow n^p(a, b) \geq n^p(b, a).$$

Application to Social Choice V

A social choice function is **majoritarian** if the social choice set assigned to each preference profile p only depends on the majority relation $\succsim^p$.

A **Condorcet winner** in preference profile p is an alternative $a \in A$ such that $a \succsim^p b$ for all $b \in A \setminus \{a\}$.

A Condorcet winner is a best element in the majority relation.

A social choice function is a **Condorcet social choice function** if it chooses the Condorcet winner in any preference profile that has a Condorcet winner.

End of Summary 1

Application to Social Choice VI

Majoritarian social choice functions based on score functions

Assumption: All individual preference relations $\succsim_i$ are complete and transitive.

For every preference profile $p = (\succsim_i)_{i \in N}$, we define the **majority digraph** D^p by

$$(a, b) \in D^p \text{ if and only if } a \succsim^p b.$$

Given a score function σ, we define the corresponding social choice function C^σ as the social choice function that chooses the alternatives with the highest scores according to σ:

$$C^\sigma(p) = \{a \in A \mid \sigma_a(D^p) \geq \sigma_b(D^p) \text{ for all } b \in A\}.$$

Application to Social Choice VII

Two special cases

The **(modified)** β**-rule** is the majoritarian social choice function $C^{\beta^{mod}}$ given by

$$C^{\beta^{mod}}(p) = \{a \in A \mid \beta_a^{mod}(D^p) \geq \beta_b^{mod}(D^p) \text{ for all } b \in A\}.$$

The λ**-rule** is the majoritarian social choice function C^{λ} given by

$$C^{\lambda}(p) = \{a \in A \mid \lambda_a(D^p) \geq \lambda_b(D^p) \text{ for all } b \in A\}.$$

Application to Social Choice VIII

Short summary of previous lectures 2

Let $\overline{\succsim}$ be the **transitive closure** of preference relation $\succsim$, i.e., $a \overline{\succsim} b$ if and only if there exist a sequence $a_1, \ldots, a_t \in A$ such that

- $a_1 = a$,
- $a_k \succsim a_{k+1}$ for all $k \in \{1, \ldots, t-1\}$,
- $a_t = b$.

Application to Social Choice IX

Definition

A subset of alternatives $T \subset A$ is a **Top cycle** in preference profile p if

- $a, b \in T, a \neq b \Rightarrow a \overline{\succsim^p} b$, and
- $a \notin T, b \in T \Rightarrow a \overline{\not\succsim^p} b$.

For preference profile p, we define the **Top set** $TOP(p)$ as the union of all Top cycles in p.

Remark: If the majority relation $\succsim^p$ is a complete and anti-symmetric relation on A, then p has exactly one Top cycle.

Application to Social Choice X

Definition

An alternative b is **covered** by alternative a in preference profile p if

- $a \succsim^p b$, and
- $b \succsim^p c \Rightarrow a \succsim^p c$ for all $c \in A$.

The **uncovered set** $UNC(p)$ is the set of alternatives that are not covered by some other alternative in p.

End of Summary 2

Application to Social Choice XI

Theorem

Consider the preference profile p. Then

- $C^{\beta^{mod}}(p) \subseteq UNC(p)$.
- If $\succsim^p$ is complete and anti-symmetric, then $C^{\beta^{mod}}(p) \subseteq TOP(p)$.

Theorem

For every preference profile p it holds that $C^{\lambda}(p) \subseteq TOP(p)$, and thus $C^{\lambda}(p) \subseteq UNC(p)$.

Application to Social Choice XII

Remarks: $C^{\beta^{mod}}$ is a refinement of the uncovered set for every social choice situation, but not a refinement of the Top Cycle for every social choice situation.

C^{λ} is a refinement of the uncovered set and the Top cycle for every social choice problem.

In this way, C^{λ} seems looks better then $C^{\beta^{mod}}$.

Next, we discuss an advantage of $C^{\beta^{mod}}$ over C^{λ}.

Application to Social Choice XIII

Two properties of social choice functions

Definition

A social choice function satisfies **Pareto optimality** if for every preference profile $p = (\succsim_i)_{i \in N}$ and alternatives $a, b \in A$ such that

- $b \succsim_i a$ for all $i \in N$, and
- there is an $i \in N$ such that $b \succ_i a$,

it holds that $a \notin C(p)$.

Interpretation: If all agents consider b at least as good as a, and at least one agent considers b better than a, then a cannot be in the choice set.

Remark: Notice the difference with Pareto optimality of social welfare functions discussed in Lecture 2.

Application to Social Choice XIV

Definition

A social choice function satisfies **monotonicity** if for every two preference profiles $p = (\succsim_i)_{i \in N}$ and $p' = (\succsim'_i)_{i \in N}$, and alternative $a \in A$ such that for every $i \in N$

- $a \succsim_i b \Rightarrow a \succsim'_i b$ for all $b \in A \setminus \{a\}$, and
- $b \succsim_i c \Leftrightarrow b \succsim'_i c$ for all $b, c \in A \setminus \{a\}, b \neq c$.

it holds that $a \in C(p)$ implies that $a \in C(p')$

Interpretation: If alternative a is in the choice set, and the individual preference relations change only by a getting 'more prefered', then a is still in the social choice set.

Application to Social Choice XV

Theorem

The social choice function $C^{\beta^{mod}}$ is a Condorcet social choice function which satisfies Pareto optimality and monotonicity.

Theorem

The social choice function C^{λ} is a Condorcet social choice function which satisfies Pareto optimality.

Remark: C^{λ} is not monotone.

Remark: Comparing $C^{\beta^{mod}}$ with C^{λ}, the disadvantage of $C^{\beta^{mod}}$ is that it is not a refinement of *Top* for every social choice problem.
A disadvantage of C^{λ} is that it is not monotone.

Social welfare functions I

Social welfare functions

Obviously, score functions can also be used to define **social welfare functions** which assign tot every digraph a complete, transitive relation.

Given a score function σ, we can define a corresponding social welfare function that ranks the alternatives according to σ:

$$a \succsim^{\sigma} b \Leftrightarrow \sigma_a(D^p) \geq \sigma_b(D^p) \text{ for all } a, b \in A.$$

Conclusion I

In this lecture we discussed several ranking methods for digraphs, and applied them to define social choice and welfare functions.
These ranking methods can be used to derive a complete and transitive preference relation from any social or individual preference relation.

Besides social choice, ranking methods are also used in, for example, to ranking teams in sports competitions or ranking WWW web pages, one of the most famous being Google Page Rank.