What’s in an axiom?

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

[joint work with Marie Schmidtlein]

Talk Outline

In our work, we often use *axioms* to describe properties of mechanisms. For most such work, what matters are the *specifics* of concrete axioms.

Today, I instead want to talk about the *nature of axioms* in general:

- What is the formal *meaning* of a given axiom?
- Are there natural *classifications* to put order in the space of axioms?

The Model

We focus on irresolute social choice functions for variable electorates.

Terminology

set of *alternatives* $=$ finite set X

preference $=$ linear order on X $=$ element of $\mathcal{L}(X)$

universe $=$ finite set $N^\star$ of *agents*

electorate $=$ set $N \subseteq N^\star$ of agents reporting a preference

profile $=$ function R from some electorate N to $\mathcal{L}(X)$

outcome $=$ nonempty subset of X (ties are allowed)

Now a *voting rule* (or SCF) is a function mapping any given profile in $\textsc{Prof} := \mathcal{L}(X)^{N \subseteq N^\star}$ to an outcome in $\textsc{Out} := 2^X \setminus \{\emptyset\}$:

$$F : \textsc{Prof} \to \textsc{Out}$$

Remark: Much (all?) of what we'll do also works for other models.

Axioms

An axiom is a *normatively desirable property* of voting rules F.

Examples:

- Anonymity = "treat all agents the same"
- Pareto = "do not select dominated alternatives"
- Strategyproofness = "don't incentivise misreporting of preferences"

Usual: Is axiom A normatively *adequate*? Is it *useful* (for the paper)?

Now: What is the *meaning* of axiom A? How do we *define* it?

Example: Defining the Anonymity Axiom

Start with an *intuitive* expression of the idea:

> *The voting rule we use should treat all agents the same.*

Then turn it into a mathematically *rigorous* definition:

$F(R) = F(\sigma \circ R)$ for all profiles R and permutations $\sigma : N^\star \to N^\star$

And maybe even provide a *formal* definition in a formal language:

$$\bigwedge_{R \in \textsc{Prof}} \ \bigwedge_{\sigma \in S_{N^\star}} \ \bigwedge_{\substack{R' \in \textsc{Prof} \text{ s.t.} \\ R'(i) = R(\sigma(i))}} \ \bigwedge_{x \in X} \ p_{R,x} \to p_{R',x}$$

Or be *explicit* and just point to the set of *all* anonymous rules:

$$\{\, \textsc{Borda}, \textsc{Copeland}, \textsc{Plurality}, \ldots, F_{4711}, \ldots \}$$

Meaning of Axioms

Two ways of fixing the *meaning* of an *axiom* A:

- *intensional* definition: list necessary and sufficient conditions
- *extensional* definition: enumerate voting rules satisfying A

Aside: Distinction goes back to Gottlob Frege (Sinn *vs.* Bedeutung).

The intensional approach is the common one in SCT:

- good for intuitions, close to philosophical starting point
- but methodologically *ad hoc*, no general formalism

So let's try the extensional approach ...

G. Frege. Über Sinn und Bedeutung. *Zeitschrift für Philosophie und Philosophische Kritik*, 100(1):25–50, 1892.

Extensional Semantics of Axioms

The *interpretation* (or *extension*) of an axiom A is a set of voting rules:

$$\mathbb{I}(A) \quad \subseteq \quad (\textsc{Prof} \to \textsc{Out})$$

such that $F \in \mathbb{I}(A)$ iff *F satisfies A*

Permits unambiguous definition of meaning of any conceivable axiom.

Applications

Let's review some applications of $\mathbb{I}(\cdot)$ as a notational tool:

- Example for a *relationship* between axioms:

  $\mathbb{I}(\textsc{Pareto}) \ \subseteq \ \mathbb{I}(\textsc{Faithfulness})$

- Example for a *characterisation* result:

  $\mathbb{I}(\textsc{Ano}) \cap \mathbb{I}(\textsc{Neu}) \cap \mathbb{I}(\textsc{PosRes}) \ = \ \{\textsc{Majority}\}$ for $|X| = 2$

- Example for an *impossibility* result:

  $\mathbb{I}(\textsc{Onto}) \cap \mathbb{I}(\textsc{StratProof}) \cap \mathbb{I}(\textsc{NonDict}) \ = \ \emptyset$ for $|X| \geqslant 3$

Classifying Axioms

We now can classify axioms in terms of their *strength*. Like this:

$$strength(A) \;=\; {}^{1}/_{|\mathbb{I}(A)|} \;\in\; (0,1]$$

Other classification approaches coming up next:

- What (or how many) profiles does an axiom talk about?
- What (or how many) profiles at a time does an axiom constrain?

Axioms Talking about Profiles

An axiom may be "talking" about one profile but not another.
Intuitively clear for intensional definitions. But for extensional ones?

Example 1: *Pareto* = "do not select *dominated* alternatives"

talks about

$a \succ b \succ c$
$b \succ c \succ a$
$b \succ a \succ c$

does not talk about

$a \succ b \succ c$
$b \succ c \succ a$
$c \succ a \succ b$

Example 2: *Anonymity* = "be invariant under permutations of agents"
talks about *all* profiles (yet fixes the outcome for none!)

Can we provide a general definition for this concept?

Axioms Talking about Profiles

For any axiom A, define $\mathbb{P}(A)$ as the intersection of all sets $S \subseteq \textsc{Prof}$ for which there exists a family $\mathcal{F}_S \subsetneq (S \to \textsc{Out})$ such that:

$$\mathbb{I}(A) \ = \ \mathcal{F}_S \otimes \{F : (\textsc{Prof} \setminus S) \to \textsc{Out}\}$$

We obtain the following “theorem”:

Axiom A talks about profile R iff $R \in \mathbb{P}(A)$.

To get the intuition, check these cases:

- $A =$ Pareto $\to$ $\mathbb{P}(A) = \{R \mid$ some x is dominated in $R\}$
- $A =$ Anonymity $\to$ $\mathbb{P}(A) = \textsc{Prof}$

Recall: $\mathbb{I}(A) = \{F : \textsc{Prof} \to \textsc{Out} \mid F$ satisfies $A\}$

Intraprofile and Interprofile Axioms

Fishburn was the first (?) to distinguish *intra-* and *interprofile* axioms:

Pareto	Anonymity
Condorcet	Monotonicity
Resoluteness	Reinforcement
⋮	⋮

Clear enough in practice for concrete axioms.

But what about a general definition?

P.C. Fishburn. *The Theory of Social Choice.* Princeton University Press, 1973.

A Hierarchy of Axioms

Call axiom A a *k-axiom* if k is the smallest integer such that:

$$\mathbb{I}(A) \;=\; \bigcap_{(R_1,\ldots,R_k)\,\in\,\textsc{Prof}^k} \{\, F \mid (F(R_1),\ldots,F(R_k)) \in A(R_1,\ldots,R_k) \,\}$$

where $A(R_1,\ldots,R_k) \;:=\; \{\, (F'(R_1),\ldots,F'(R_k)) \mid F' \in \mathbb{I}(A) \,\}$

So a k-axiom only ever imposes a constraint on k profiles at a time.

Some observations:

- Fishburn's intraprofile axioms $=$ 1-axioms
- Fishburn's interpofile axioms $\approx$ k-axioms with $k > 1$ [more soon]
- Every axiom is a k-axiom for some $k \leqslant |\textsc{Prof}|$.

Recall: $\mathbb{I}(A) = \{F : \textsc{Prof} \to \textsc{Out} \mid F$ satisfies $A\}$

Active and Passive Intraprofile Axioms

Fishburn further divides intraprofile axioms into those that are *active* (that "involve specific conditions on contents") and *passive* axioms:

Pareto	Resoluteness
Condorcet	
⋮	⋮

We know how to formalise this!

Axiom A is passive only if $\mathbb{P}(A) = \textsc{Prof}$.

P.C. Fishburn. *The Theory of Social Choice.* Princeton University Press, 1973.

Universal and Existential Axioms

Fishburn restricts the terms intra- and interprofile to *universal* axioms, and distinguishes those from *existential* axioms such as this:

Nonimposition = "every $x \in X$ should win alone in *some* profile"

Intuitively, this is about the type of quantification over profiles:

> "existential [axioms] are based primarily on existential qualifiers [...] universal [axioms] do not use existential qualifiers in any way, or [...] in a secondary manner"

Even less clear what any of this might mean when there is no language.

But: Typical "existential" axioms are k-axioms for $k = |\textsc{Prof}|$.

P.C. Fishburn. *The Theory of Social Choice.* Princeton University Press, 1973.

Last Slide

I shared a few ruminations about the nature of axioms culminating in language-independent definitions of three fundamental concepts:

- the meaning of an axiom
- the notion of an axiom talking about a profile
- the structural complexity of the constraints an axiom can impose

For full details, see Chapter 2 of Marie Schmidtlein's MSc thesis.

M.C. Schmidtlein. *Voting by Axioms*. MSc thesis, University of Amsterdam, 2022.