The Abstract Mathematics of the Axiomatic Method

Frank M. V. Feys
Independent researcher, Antwerp, Belgium
frank@fmvfeys.com
ORCID: 0009-0000-7667-6742

October 29, 2026

Abstract

We study the axiomatic method itself as a mathematical object. An axiom is a subset of a finite domain, an axiom system is a family of such subsets, and the model class is their intersection. The model-class map forms a Galois connection with the theory map, and adding independent axioms drives a strictly decreasing chain of model classes terminating in one of four regimes: vacuity, possibility, uniqueness, or impossibility. Given a probability measure on the domain, every axiom system defines a cooperative game on its axiom set, and the Shapley value of this game gives a canonical attribution of credit or blame to each axiom. We prove that axiom games are characterized by nonnegative Möbius coefficients of the complement function $1 - v$, that all pairwise axiom interactions exhibit diminishing marginal returns, and that the minimum number of axiom drops needed to restore possibility equals the transversal number of the family of impossibility cores. We introduce the impossibility dimension, prove a duality relating it to the minimum domain size certifying independence, show that the simplicial complex of possible sub-systems has homology governed by the Alexander dual of the core family, and establish a sharp phase transition from possibility to impossibility in random axiom systems. The framework is illustrated on a catalogue of fully derived examples, including Condorcet, Arrow-type, and Nash-type systems on small finite domains, in which every model-class size is obtained by direct enumeration.

Keywords. Axiomatic method; impossibility theorem; characterization theorem; Shapley value; Galois connection; Möbius function; Alexander duality; axiom game; cooperative game.

Mathematics Subject Classification (2020). 91B14; 91A12; 06A15; 05E45.

1 Introduction

The axiomatic method is among the most powerful tools in mathematics and economic theory. One fixes a domain of objects, formulates a list of desiderata, and asks: which objects satisfy all axioms simultaneously? The answer may be an entire class, a unique object, or nothing at all. These three outcomes correspond to three celebrated kinds of results: possibility theorems, uniqueness (characterization) theorems, and impossibility theorems.

A simple example illustrates the pattern. Consider the problem of choosing a "fair" division rule for splitting a dollar between two players. One might require the rule to be *efficient* (the entire dollar is distributed), *symmetric* (identical players receive equal shares), and *individually rational* (no player receives less than zero). These three axioms, applied to the domain of all conceivable division rules, single out the equal split. This is a uniqueness theorem. Adding a fourth axiom, say *dictatorial* (one designated player receives everything), produces a contradiction with efficiency and symmetry. This is an impossibility theorem. In each case, the reasoning has the same abstract structure. The axioms are conditions (subsets of the domain), the model class is their intersection, and the conclusion depends on whether that intersection is large, a singleton, or empty.

Despite the ubiquity of this pattern, its abstract mathematical skeleton is rarely made explicit. Papers in cooperative game theory, social choice, mechanism design, and mathematical finance all apply the method, but typically work within a concrete domain rather than studying the method itself. The present paper fills this gap at three levels.

At the qualitative level, we show that an axiom is a subset of a domain, an axiom system is a set of such subsets, and the model class is their intersection. The model-class map is antitone and forms one leg of a Galois connection with the theory map, whose fixed points are the logically closed systems and the definable model classes. Adding an independent axiom strictly shrinks the model class, and the resulting chain of strict inclusions is bounded in length by the size of the domain. The chain terminates in one of four regimes, depending on whether the final model class is all of $\mathcal{U}$ (vacuity), has at least two elements (possibility), is a singleton (uniqueness), or is empty (impossibility). We prove theorems about each regime and about the transitions between them.

At the quantitative level, we develop several contributions. Every axiom system defines a cooperative game on the axiom set, and the *axiom Shapley value* of this game gives a canonical attribution of credit (or blame) to each axiom. The game is normalized, monotone, and submodular; its submodularity is equivalent to the *Axiom Interaction Theorem*, which establishes that axioms within any axiom system always exhibit diminishing marginal returns. The *Characterization of Axiom Games* shows that a set function is an axiom game if and only if the complement function $1 - v$ has nonnegative Möbius coefficients, a condition strictly stronger than submodularity. The axiom Shapley value is then the centroid of the $n!$ greedy marginal vectors, which are the vertices of the base polytope (with possible repetitions). The *Axiom Dominance Theorem* gives an exact formula for the Shapley-value difference of two axioms and shows that the qualitative ranking is invariant over the class of measures under which one axiom is uniformly more restrictive than the other. This places the classical examples on a firm structural foundation. The *Rescue Theorem* identifies the optimal single-axiom drop via the full-coalition marginal and characterizes the weakening condition for restoring possibility. The *Transversal Theorem* shows that the minimum rescue depth equals the minimum transversal number of the family of impossibility cores, separating the Shapley-based ranking from the structural optimum. We also introduce *axiom entropy*, a measure of how evenly the collapse burden is distributed among axioms.

At the structural level, we introduce the *impossibility dimension* $\dim(\Sigma)$, the minimum size of an impossibility core, and prove a Dimension-Witness Duality showing it equals the minimum number of distinct witness objects needed to certify independence. The *Feasibility Complex Theorem* establishes that the collection of possible sub-systems is a simplicial

complex; when the system is itself an impossibility core, this complex is homeomorphic to a sphere. The *Alexander Duality Theorem* shows that the reduced homology of this complex is completely determined by the overlap pattern of the impossibility cores, via the Alexander dual (the "co-core complex"); contractibility and cone-point acyclicity both follow as corollaries. The transition from possibility to impossibility in random axiom systems is abrupt, passing from almost-sure possibility to almost-sure impossibility over a vanishingly narrow window, the threshold governed by the impossibility dimension.

1.1 Related Work

The axiomatic method has a long history, and several threads of existing work touch on the themes of this paper. At the most foundational level, Hilbert's programme sought to reduce mathematics itself to a system of axioms whose consistency and completeness could be established by finitary means. While Gödel's incompleteness theorems showed this programme cannot fully succeed, the impulse to study axiom systems *as mathematical objects* rather than merely *through* them remains central to mathematical logic and foundations. In a different direction, the programme of *reverse mathematics* asks which axioms are necessary and sufficient to prove a given theorem over a weak base theory, yielding a calibration of logical strength that is conceptually related to our notions of independence and impossibility cores, though the setting and methods are entirely different (we work with finite set-theoretic domains rather than subsystems of second-order arithmetic). In universal algebra, the study of equational theories, variety lattices, and Birkhoff's theorem on equational classes provides a general algebraic framework for the interplay between axioms and models. The Galois connection between properties and classes of structures is a recurring theme in that literature, as it is in ours.

In social choice theory, Arrow (1951) initiated the study of impossibility theorems, and the logical structure of such results has been analyzed by many authors, including Fishburn (1970), Sen (1970), Blau (1972), and Campbell and Kelly (2002). Saari (2008) studied the geometry of voting paradoxes, and Baryshnikov (1993) and Weinberger (2004) applied topological methods to social choice impossibilities. The Shapley value (Shapley, 1953) and its properties in cooperative game theory are surveyed in Winter (2002) and Roth (1988). The interaction index for coalitions was axiomatized by Grabisch and Roubens (1999). Submodular functions and polymatroids were studied by Edmonds (1970), with the general theory developed further by Lovász (1983) and Fujishige (2005). Sharp thresholds for monotone properties were established by Bollobás and Thomason (1987) and sharpened by Friedgut and Kalai (1996).

The Galois connection between axiom systems and model classes developed in Section 5 is an instance of a broader construction that appears throughout lattice theory and order theory. In particular, the framework of *formal concept analysis* (FCA), introduced by Wille (1982) and developed systematically in Ganter and Wille (1999), is built on the same adjunction between a set of "objects" and a set of "attributes," with the closed elements of the resulting closure operators called *formal concepts* and the lattice of such concepts called the *concept lattice*. In our setting, the objects are elements of $\mathcal{U}$ and the attributes are axioms; the formal concepts correspond to the fixed points identified in Theorem 5.1(iv). We include the Galois connection for completeness and because the specific consequences we

draw from it (the four-regime taxonomy, the chain structure, and its interaction with the cooperative game and topological layers developed below) appear not to have been made explicit in either the FCA literature or the social choice literature.

The feasibility complex of Section 11 is closely related to the *independence complex* studied in combinatorial topology (see, e.g., Kozlov, 2008), in which the faces of a simplicial complex encode the independent sets of a combinatorial structure and topological invariants carry information about its constraints. Our feasibility complex specializes this idea to the case where "independence" means joint satisfiability of axioms.

The present paper differs from all of the above in that it studies the axiomatic method itself as a mathematical object, rather than working within any particular axiomatic domain. The cooperative game on the axiom set (Section 8), the Rescue Theorem and transversal characterization (Section 9), the Dimension-Witness Duality (Section 10), the feasibility complex and its topology (Section 11), and the phase transition analysis (Section 14) are, to the best of our knowledge, new.

1.2 Contributions and Novelty

To assist the reader in separating new results from known ones, we briefly indicate the status of each major result; Table 1 provides a complete summary.

Framework (known or elementary). The Monotone Collapse Theorem (Theorem 4.1) and the Galois Connection (Theorem 5.1) are standard results from lattice theory and formal concept analysis (Ganter and Wille, 1999), restated here in the axiom-system context. The Uniqueness, Core Existence, and Independence Witnessing Theorems are elementary consequences of the finite-domain setup, included for self-containment.

New contributions. The axiom game (Σ, v_μ), the Axiom Shapley Value Theorem (Theorem 8.4), and the Axiom Interaction Theorem (Theorem 8.8) are new, as is the Characterization of Axiom Games (Theorem 8.17), which identifies axiom games via non-negative Möbius coefficients of $1 - v$. The Axiom Dominance Theorem (Theorem 8.12) gives an exact formula for pairwise Shapley-value differences and certifies the qualitative ranking over entire classes of measures. The Rescue Theorem (Theorem 9.4), the Transversal Theorem (Theorem 9.8), the Dimension-Witness Duality (Theorem 10.6), the Feasibility Complex Theorem (Theorem 11.3), the Alexander Duality Theorem (Theorem 11.11), the Contractibility Theorem (Theorem 11.7), and the Phase Transition Theorem (Theorem 14.3) are all new.

1.3 Outline and Summary of Main Results

Figure 1 provides a synoptic view of the paper's architecture, showing how the main objects and theorems are organized across the qualitative, quantitative, and structural/probabilistic layers. Table 1 summarizes each main result, its status, and its location in the paper.

Result	Type	Status	Ref.
Monotone Collapse	Thm	Known/elementary	4.1
Galois Connection	Thm	Known (FCA)	5.1
Uniqueness	Thm	Elementary	6.1
Core Existence	Thm	Elementary	7.1
Independence Witnessing	Thm	Elementary	7.3
Properties of the Axiom Game	Prop	New	8.2
Axiom Shapley Value	Thm	New	8.4
Axiom Interaction	Thm	New	8.8
Characterization of Axiom Games	Thm	New	8.17
Shapley Value as Centroid	Prop	New	8.14
Axiom Dominance	Thm	New	8.12
Axiom Entropy Bounds	Thm	New	8.23
Rescue	Thm	New	9.4
Impossibility Depth = Transversal Number	Thm	New	9.8
Dimension-Witness Duality	Thm	New	10.6
Topology of the Feasibility Complex	Thm	New	11.3
Contractibility	Thm	New	11.7
Alexander Duality	Thm	New	11.11
Phase Transition	Thm	New	14.3

Table 1: Summary of main results. "Known/elementary" indicates results that are standard or immediate from the finite-domain setup, included for completeness. "New" indicates results that appear to be new to this paper.

2 Setup and Basic Definitions

Throughout, we fix a nonempty *finite* set $\mathcal{U}$, called the *domain* or *universe*, and we write $N = |\mathcal{U}|$ for its cardinality. Elements of $\mathcal{U}$ are the objects under study. Unless stated otherwise, we assume $N \geq 2$; the degenerate case $N = 1$ is discussed in Remark 4.3 below.

Definition 2.1 (Axiom, Axiom System, Model Class)**.** An *axiom* is a nonempty subset $A \subseteq \mathcal{U}$. An *axiom system* is a finite set $\Sigma \subseteq \mathcal{P}(\mathcal{U}) \setminus \{\emptyset\}$ of axioms. The *model class* of Σ is

$$\mathcal{M}(\Sigma) = \bigcap_{A \in \Sigma} A,$$

with the convention that $\mathcal{M}(\emptyset) = \mathcal{U}$. An object $f \in \mathcal{U}$ is a *model* of Σ if $f \in \mathcal{M}(\Sigma)$.

Remark 2.2. The requirement $A \neq \emptyset$ excludes trivially impossible axioms (those satisfied by no object) from the theory. This normalization is harmless, and ensures individual satisfiability. An axiom $A = \emptyset$ makes any system containing it impossible regardless of the remaining axioms, and contributes nothing to the interplay among axioms that is our primary concern. All results below concern axioms that are individually satisfiable; the interesting phenomenon is joint unsatisfiability.

Remark 2.3 (Extensional Sense of "Axiom")**.** In mathematical logic, an axiom is a syntactic object, namely a formula in some formal language, and one speaks of its set of models as its extension. Throughout this paper we use the word *axiom* in the extensional sense, i.e., we

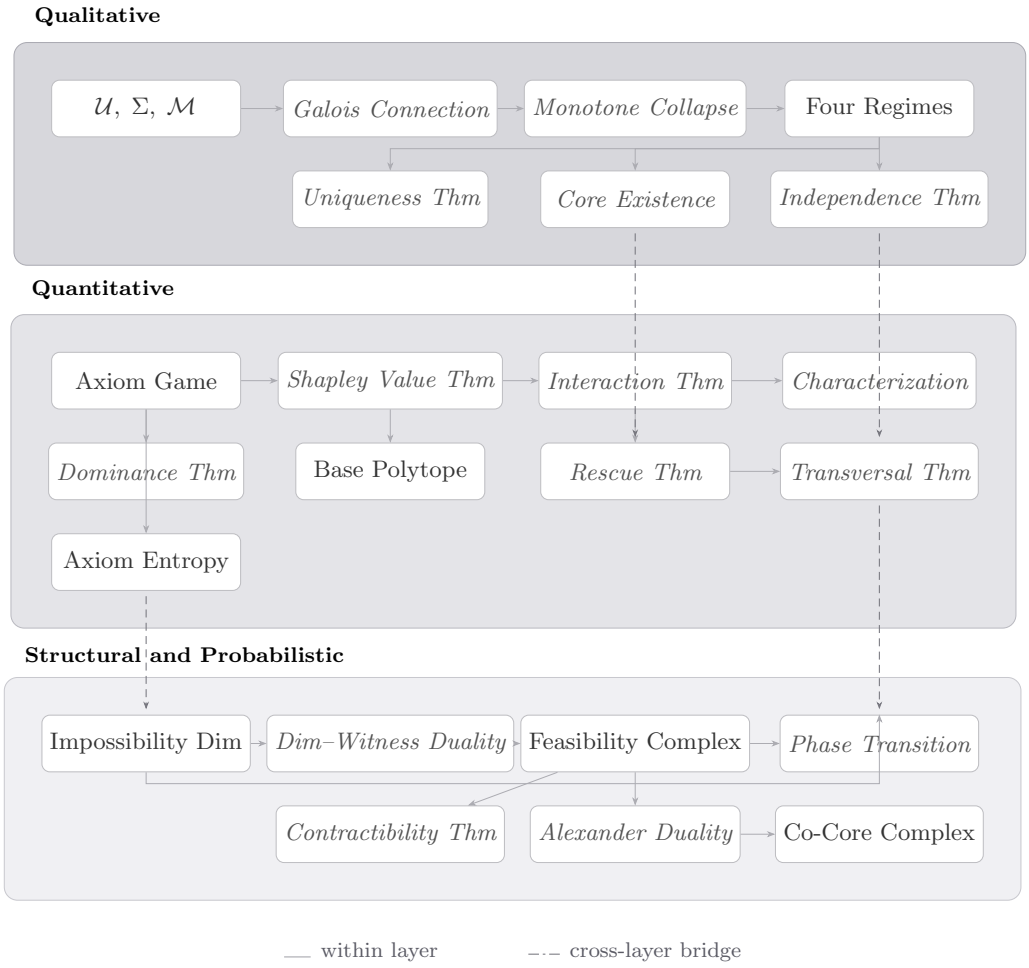


Figure 1: Synoptic map of the paper's architecture. The three horizontal bands correspond to the qualitative, quantitative, and structural/probabilistic layers. Solid arrows indicate logical dependence within a layer; dashed arrows indicate cross-layer bridge results.

identify an axiom with the set of objects satisfying it. This is a semantic simplification suited to the finite-domain setting and to the combinatorial questions we shall study. No claim is made about syntactic formulation or proof-theoretic content.

Remark 2.4 (Union of Axiom Sets)**.** For any two axiom systems $S, T \subseteq \mathcal{P}(\mathcal{U}) \setminus \{\emptyset\}$, the model class of their union satisfies $\mathcal{M}(S \cup T) = \mathcal{M}(S) \cap \mathcal{M}(T)$. Indeed,

$$\mathcal{M}(S \cup T) = \bigcap_{A \in S \cup T} A = \Big(\bigcap_{A \in S} A\Big) \cap \Big(\bigcap_{A \in T} A\Big) = \mathcal{M}(S) \cap \mathcal{M}(T).$$

We shall use this identity repeatedly.

Definition 2.5 (Independence and Redundancy)**.** An axiom $A \in \Sigma$ is *independent* in Σ if $\mathcal{M}(\Sigma \setminus \{A\}) \supsetneq \mathcal{M}(\Sigma)$. The system Σ is *independent* if every $A \in \Sigma$ is independent in Σ. An axiom A is *redundant* in Σ if $\mathcal{M}(\Sigma \setminus \{A\}) = \mathcal{M}(\Sigma)$.

Definition 2.6 (Theory and Definability)**.** The *theory* of a class $C \subseteq \mathcal{U}$ is defined as $\mathrm{Th}(C) = \{A \in \mathcal{P}(\mathcal{U}) \setminus \{\emptyset\} \mid C \subseteq A\}$. A class C is *definable* if $C = \mathcal{M}(\Sigma)$ for some axiom system Σ. An axiom system Σ is *logically closed* if $\Sigma = \mathrm{Th}(\mathcal{M}(\Sigma))$.

We equip $\mathcal{U}$ with a probability measure μ, i.e., a function $\mu\colon \mathcal{P}(\mathcal{U}) \to [0,1]$ with $\mu(\mathcal{U}) = 1$ and μ finitely additive on disjoint sets. Since $\mathcal{U}$ is finite, every subset is automatically measurable. This is the setting of Section 8 and later; Sections 4–7 require no measure.

Table 2 collects the principal notation used throughout the paper for the reader's convenience.

Symbol	Meaning
$\mathcal{U}$, $N = \lvert\mathcal{U}\rvert$	Finite domain (universe) and its cardinality
Σ	Axiom system (finite family of nonempty subsets of $\mathcal{U}$)
$\mathcal{M}(\Sigma)$	Model class $\bigcap_{A\in\Sigma} A$
$\mathrm{Th}(C)$	Theory of a class C (all axioms containing C)
$v_\mu(S)$	Collapse function $1 - \mu(\mathcal{M}(S))$
$\phi_A(\Sigma, \mu)$	Axiom Shapley value of axiom A
$\iota_\mu(A, B \mid S)$	Pairwise interaction of A and B conditional on S
Δ_A	Full-coalition marginal $v_\mu(\Sigma) - v_\mu(\Sigma \setminus \{A\})$
$\dim(\Sigma)$	Impossibility dimension (minimum core size)
$d(\Sigma)$	Impossibility depth (minimum axiom drops to restore possibility)
$\mathcal{C}(\Sigma)$	Family of impossibility cores
$\tau(\mathcal{C}(\Sigma))$	Transversal number of the core family
$K(\Sigma)$	Feasibility complex (simplicial complex of possible sub-systems)
$K^\vee(\Sigma)$	Co-core complex (Alexander dual of $K(\Sigma)$)
$H(\Sigma, \mu)$	Axiom entropy

Table 2: Principal notation.

3 A Running Example

To build intuition before the general theory, we introduce a small concrete example that we shall revisit in each subsequent section.

Example 3.1 (The Four-Axiom System on Six Objects)**.** Let $\mathcal{U} = \{1,2,3,4,5,6\}$ and consider the axiom system $\Sigma = \{A_1, A_2, A_3, A_4\}$ defined by

$$A_1 = \{1,2,3\}, \quad A_2 = \{4,5,6\}, \quad A_3 = \{1,2,4,5\}, \quad A_4 = \{1,3,4,6\}.$$

The model class is $\mathcal{M}(\Sigma) = A_1 \cap A_2 \cap A_3 \cap A_4 = \emptyset$, so the system is in the impossibility regime. Indeed, note that $A_1 \cap A_2 = \emptyset$, so the sub-system $\{A_1, A_2\}$ is already impossible, making it an impossibility core of size 2. No other pair of axioms has empty intersection, so this is the unique minimal core. The axioms A_1 and A_2 are independent (dropping either restores possibility), while A_3 and A_4 are redundant in Σ (in the technical sense of Definition 2.5), since $\mathcal{M}(\Sigma \setminus \{A_3\}) = \mathcal{M}(\Sigma \setminus \{A_4\}) = \emptyset = \mathcal{M}(\Sigma)$. In the impossibility regime, an axiom A is redundant in Σ precisely when A lies outside every impossibility core of Σ, so that removing A leaves the core (and hence the impossibility) intact.

We shall equip $\mathcal{U}$ with the uniform measure $\mu(\{i\}) = 1/6$ and revisit this system in subsequent sections to illustrate the Galois connection (Section 5), the cooperative game and Shapley value (Section 8), the Rescue and Transversal Theorems (Section 9), the impossibility dimension (Section 10), and the feasibility complex (Section 11).

4 The Monotone Collapse Theorem

We begin with the fundamental structural property of the model-class map.

Theorem 4.1 (Monotone Collapse)**.** *Let $\Sigma, \Sigma' \subseteq \mathcal{P}(\mathcal{U}) \setminus \{\emptyset\}$ be axiom systems.*

(i) Antitonicity. *If $\Sigma \subseteq \Sigma'$, then $\mathcal{M}(\Sigma') \subseteq \mathcal{M}(\Sigma)$.*

(ii) Strict Collapse. *If $\Sigma \subseteq \Sigma'$ and some $A \in \Sigma' \setminus \Sigma$ is independent in Σ', then $\mathcal{M}(\Sigma') \subsetneq \mathcal{M}(\Sigma)$.*

(iii) Chain Structure. *Whenever each A_k is independent in $\{A_1, A_2, \ldots, A_k\}$, the induced model classes form a strictly decreasing chain:*

$$\mathcal{U} = \mathcal{M}(\emptyset) \supsetneq \mathcal{M}(\{A_1\}) \supsetneq \mathcal{M}(\{A_1, A_2\}) \supsetneq \cdots .$$

(iv) Termination. *Since $\mathcal{U}$ is finite, every strictly decreasing chain of model classes has length at most $|\mathcal{U}|$.*

Proof. For (i), if $f \in \mathcal{M}(\Sigma')$, then $f \in A$ for every $A \in \Sigma' \supseteq \Sigma$, so $f \in \mathcal{M}(\Sigma)$.

For (ii), by independence of A in Σ', there exists $g \in \mathcal{M}(\Sigma' \setminus \{A\}) \setminus A$. Since $\Sigma \subseteq \Sigma' \setminus \{A\}$, part (i) gives $g \in \mathcal{M}(\Sigma)$; since $g \notin A$, we have $g \notin \mathcal{M}(\Sigma')$.

For (iii), we apply (ii) at each step to $\Sigma = \{A_1, A_2, \ldots, A_{k-1}\}$ and $\Sigma' = \{A_1, A_2, \ldots, A_k\}$.

For (iv), since $\mathcal{U}$ is finite with $|\mathcal{U}| = N$, each strict inclusion step removes at least one element from the model class, so the chain from $\mathcal{U}$ to $\emptyset$ has length at most N. □

4.1 The Four Regimes

The model class of an axiom system falls into one of four mutually exclusive regimes.

Definition 4.2 (The Four Regimes)**.** Let Σ be an axiom system. We say Σ is in the

(i) *vacuous regime* if $\mathcal{M}(\Sigma) = \mathcal{U}$;

(ii) *possibility regime* if $2 \leq |\mathcal{M}(\Sigma)| < |\mathcal{U}|$;

(iii) *uniqueness regime* if $|\mathcal{M}(\Sigma)| = 1$;

(iv) *impossibility regime* if $\mathcal{M}(\Sigma) = \emptyset$.

Under the standing assumption $N \geq 2$, the four regimes are mutually exclusive and exhaustive, with the possibility regime occupying the strict middle ground $2 \leq |\mathcal{M}(\Sigma)| < |\mathcal{U}|$. In the colloquial sense, the uniqueness regime stands at the boundary between possibility and impossibility; it is the last station before the model class reaches $\emptyset$.

Remark 4.3 (The Degenerate One-Element Domain)**.** If $N = 1$, then the only nonempty axiom is $\mathcal{U}$ itself, and every axiom system has $\mathcal{M}(\Sigma) = \mathcal{U}$; the vacuous and uniqueness regimes coincide. If $N = 2$, then the possibility regime is empty, since $|\mathcal{M}(\Sigma)|$ can only be 0, 1, or 2. The framework is nontrivial precisely for $N \geq 3$, and we shall assume this whenever a regime distinction is needed.

As independent axioms are added, the model class travels strictly downward in the lattice $(\mathcal{P}(\mathcal{U}), \subseteq)$, starting in the vacuous regime and terminating either in the uniqueness regime (a characterization theorem) or in the impossibility regime. Figure 2 illustrates this process.

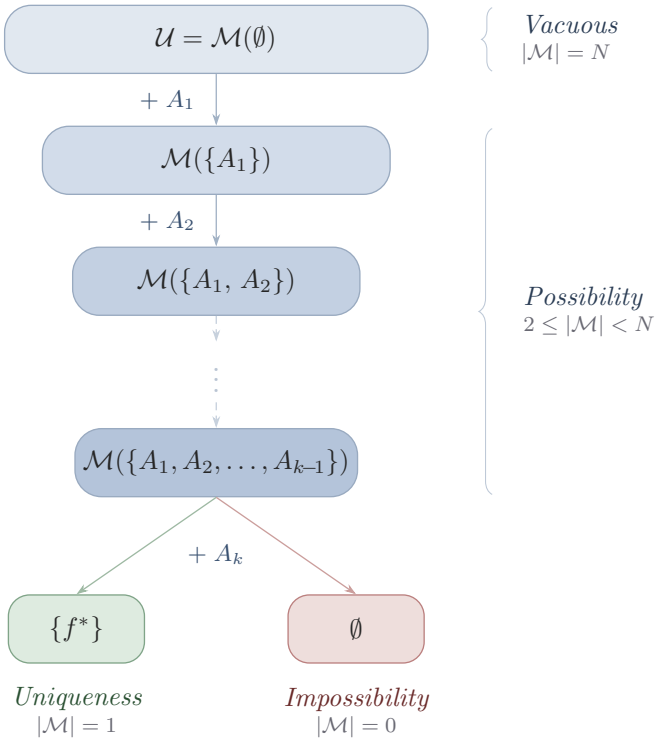


Figure 2: The monotone collapse of the model class. As independent axioms $A_1, A_2, \ldots$ are added, the model class $\mathcal{M}$ shrinks strictly (Theorem 4.1). The process terminates in one of four regimes (Definition 4.2), depending on whether the final model class is all of $\mathcal{U}$ (vacuous), has at least two elements (possibility), is a singleton (uniqueness), or is empty (impossibility). The decreasing width of the nodes represents the strict containment $\mathcal{M}(\Sigma') \subsetneq \mathcal{M}(\Sigma)$ at each step.

5 The Galois Connection

The maps $\mathcal{M}$ and Th are not merely monotone; they form a Galois connection whose fixed points characterize the logically closed systems and the definable classes.

Theorem 5.1 (Galois Connection)**.** *The pair* $(\mathcal{M}, \mathrm{Th})$ *forms a Galois connection between axiom systems ordered by inclusion and classes ordered by reverse inclusion; that is, between*

$(\mathcal{P}(\mathcal{P}(\mathcal{U}) \setminus \{\emptyset\}), \subseteq)$ *and* $(\mathcal{P}(\mathcal{U}), \supseteq)$*. In particular, for every axiom system* Σ *and every class* $C \subseteq \mathcal{U}$*,*

$$C \subseteq \mathcal{M}(\Sigma) \quad \Longleftrightarrow \quad \Sigma \subseteq \mathrm{Th}(C).$$

The following additional properties hold.

(i) $\Sigma \subseteq \mathrm{Th}(\mathcal{M}(\Sigma))$ *and* $C \subseteq \mathcal{M}(\mathrm{Th}(C))$ *for all* Σ *and* C*.*

(ii) The composites $\Sigma \mapsto \mathrm{Th}(\mathcal{M}(\Sigma))$ *and* $C \mapsto \mathcal{M}(\mathrm{Th}(C))$ *are closure operators.*

(iii) An axiom system Σ *is logically closed if and only if* $\Sigma = \mathrm{Th}(\mathcal{M}(\Sigma))$*. A class* C *is definable if and only if* $C = \mathcal{M}(\mathrm{Th}(C))$*.*

(iv) The maps $\mathcal{M}$ *and* Th *restrict to mutually inverse order-reversing bijections between logically closed axiom systems and definable classes.*

Proof. For the main equivalence, $C \subseteq \mathcal{M}(\Sigma) = \bigcap_{A \in \Sigma} A$ holds if and only if $C \subseteq A$ for every $A \in \Sigma$, which holds if and only if $\Sigma \subseteq \mathrm{Th}(C)$.

For (i), setting $C = \mathcal{M}(\Sigma)$ in the main equivalence gives $\mathcal{M}(\Sigma) \subseteq \mathcal{M}(\Sigma)$, hence $\Sigma \subseteq \mathrm{Th}(\mathcal{M}(\Sigma))$. Similarly, setting $\Sigma = \mathrm{Th}(C)$ gives $C \subseteq \mathcal{M}(\mathrm{Th}(C))$.

For (ii), we verify the three closure-operator axioms (i.e., extensivity, monotonicity, idempotence) for the composite $\mathrm{cl}\colon \Sigma \mapsto \mathrm{Th}(\mathcal{M}(\Sigma))$. Extensivity ($\Sigma \subseteq \mathrm{cl}(\Sigma)$) is part (i). Monotonicity follows from the fact that $\mathcal{M}$ is antitone and Th is antitone (in the reverse-inclusion order on classes), so their composite is order-preserving. Concretely, if $\Sigma \subseteq \Sigma'$, then $\mathcal{M}(\Sigma') \subseteq \mathcal{M}(\Sigma)$ by Theorem 4.1(i), hence $\mathrm{Th}(\mathcal{M}(\Sigma)) \subseteq \mathrm{Th}(\mathcal{M}(\Sigma'))$, i.e., $\mathrm{cl}(\Sigma) \subseteq \mathrm{cl}(\Sigma')$. For idempotence, note that $\mathrm{cl}(\mathrm{cl}(\Sigma)) = \mathrm{Th}(\mathcal{M}(\mathrm{Th}(\mathcal{M}(\Sigma))))$. By part (i) applied to the class $C = \mathcal{M}(\Sigma)$, we have $\mathcal{M}(\Sigma) \subseteq \mathcal{M}(\mathrm{Th}(\mathcal{M}(\Sigma)))$. By antitonicity of $\mathcal{M}$, we have $\mathcal{M}(\mathrm{Th}(\mathcal{M}(\Sigma))) \subseteq \mathcal{M}(\Sigma)$ (since $\Sigma \subseteq \mathrm{Th}(\mathcal{M}(\Sigma))$ by part (i), antitonicity gives $\mathcal{M}(\mathrm{Th}(\mathcal{M}(\Sigma))) \subseteq \mathcal{M}(\Sigma)$). Hence, $\mathcal{M}(\mathrm{Th}(\mathcal{M}(\Sigma))) = \mathcal{M}(\Sigma)$, and applying Th to both sides gives $\mathrm{cl}(\mathrm{cl}(\Sigma)) = \mathrm{cl}(\Sigma)$. The proof for $C \mapsto \mathcal{M}(\mathrm{Th}(C))$ is symmetric.

For (iii), Σ is logically closed if and only if $\Sigma = \mathrm{cl}(\Sigma) = \mathrm{Th}(\mathcal{M}(\Sigma))$, which is the definition combined with the fact that the fixed points of a closure operator are precisely the closed elements. For the model side, C is definable if and only if $C = \mathcal{M}(\Sigma)$ for some Σ, and by idempotence $\mathcal{M}(\mathrm{Th}(C)) = \mathcal{M}(\mathrm{Th}(\mathcal{M}(\Sigma))) = \mathcal{M}(\Sigma) = C$, so C is a fixed point. Conversely, if $C = \mathcal{M}(\mathrm{Th}(C))$, then C is definable (with $\Sigma = \mathrm{Th}(C)$).

For (iv), by part (iii), the maps $\mathcal{M}$ and Th restrict to functions between logically closed systems and definable classes. We verify that these restrictions are mutually inverse. If Σ is logically closed, then $\mathrm{Th}(\mathcal{M}(\Sigma)) = \Sigma$, so $\mathrm{Th} \circ \mathcal{M}$ is the identity on logically closed systems. If C is definable, then $\mathcal{M}(\mathrm{Th}(C)) = C$, so $\mathcal{M} \circ \mathrm{Th}$ is the identity on definable classes. These restricted maps are order-reversing because $\mathcal{M}$ and Th are individually antitone. □

Remark 5.2 (Definability on Finite Domains)**.** In the finite-domain setting, every subset of $\mathcal{U}$ is definable. Indeed, for any nonempty $C \subseteq \mathcal{U}$, the singleton system $\{C\}$ satisfies $\mathcal{M}(\{C\}) = C$, and $\emptyset$ is definable whenever $|\mathcal{U}| \geq 2$ (take any two disjoint nonempty axioms, which exist since $\mathcal{U}$ has at least two elements). The fixed-point structure on the model side of the Galois connection is therefore degenerate, and the interesting structure lies in the axiom side and in the quantitative and topological layers developed below.

Example 5.3 (Running Example: Galois Connection)**.** Returning to Example 3.1, consider the sub-system $S = \{A_3, A_4\}$. The model class is $\mathcal{M}(S) = A_3 \cap A_4 = \{1, 4\}$. The theory of this class is $\mathrm{Th}(\{1, 4\}) = \{A \in \mathcal{P}(\mathcal{U}) \mid \{1, 4\} \subseteq A\}$, which contains A_3 and A_4 but also many other subsets of $\mathcal{U}$. The closure $\mathrm{Th}(\mathcal{M}(S))$ is therefore strictly larger than S, confirming that S is not logically closed. Observe that $\mathcal{M}(\mathrm{Th}(\{1, 4\})) = \{1, 4\}$, so the class $\{1, 4\}$ is definable and hence is a fixed point of the closure operator $C \mapsto \mathcal{M}(\mathrm{Th}(C))$.

6 The Uniqueness Theorem

We now characterize when an axiom system singles out exactly one object, the setting of every characterization theorem in the axiomatic tradition.

Theorem 6.1 (Uniqueness)**.** *Let $f^* \in \mathcal{U}$. The following are equivalent.*

(i) There exists an axiom system Σ such that $\mathcal{M}(\Sigma) = \{f^*\}$.

(ii) The singleton $\{f^*\}$ is a definable class.

(iii) $\mathcal{M}(\mathrm{Th}(\{f^*\})) = \{f^*\}$, and $\mathrm{Th}(\{f^*\})$ is the unique logically closed system whose model class is $\{f^*\}$.

Moreover, in any axiom system Σ with $\mathcal{M}(\Sigma) = \{f^*\}$, the following hold.

(a) Existence. *Every axiom in Σ is satisfied by f^*.*

(b) Uniqueness. *The axioms jointly exclude every object other than f^*.*

(c) Canonical closure. *The system $\mathrm{Th}(\{f^*\})$ is the unique logically closed axiom system with model class $\{f^*\}$, and every axiom system Σ with $\mathcal{M}(\Sigma) = \{f^*\}$ satisfies $\Sigma \subseteq \mathrm{Th}(\{f^*\})$.*

Proof. The equivalences (i)–(iii) follow from Theorem 5.1(iii)–(iv). Indeed, (i)⇒(ii) is the definition of definability. For (ii)⇒(iii), if $\{f^*\}$ is definable, then $\{f^*\} = \mathcal{M}(\mathrm{Th}(\{f^*\}))$ by Theorem 5.1(iii), and the system $\mathrm{Th}(\{f^*\})$ is logically closed by construction (it is a fixed point of the closure operator $\Sigma \mapsto \mathrm{Th}(\mathcal{M}(\Sigma))$). Uniqueness of the logically closed system follows from the bijection of Theorem 5.1(iv). For (iii)⇒(i), the system $\mathrm{Th}(\{f^*\})$ witnesses definability.

For (a), if $\mathcal{M}(\Sigma) = \{f^*\}$, then $f^* \in A$ for every $A \in \Sigma$, by definition of $\mathcal{M}(\Sigma) = \bigcap_{A \in \Sigma} A$.

For (b), let $g \in \mathcal{U}$ with $g \neq f^*$. Since $\mathcal{M}(\Sigma) = \{f^*\}$, we have $g \notin \mathcal{M}(\Sigma)$, so there exists some $A \in \Sigma$ with $g \notin A$. Hence, the axioms jointly exclude every object other than f^*.

For (c), by Theorem 5.1(i), $\Sigma \subseteq \mathrm{Th}(\mathcal{M}(\Sigma)) = \mathrm{Th}(\{f^*\})$. The system $\mathrm{Th}(\{f^*\})$ is logically closed because $\mathrm{Th}(\mathcal{M}(\mathrm{Th}(\{f^*\}))) = \mathrm{Th}(\{f^*\})$ by Theorem 5.1(ii) (idempotence of the closure operator). If Σ' is any other logically closed system with $\mathcal{M}(\Sigma') = \{f^*\}$, then $\Sigma' = \mathrm{Th}(\mathcal{M}(\Sigma')) = \mathrm{Th}(\{f^*\})$ by Theorem 5.1(iii), so $\mathrm{Th}(\{f^*\})$ is the unique logically closed system with model class $\{f^*\}$. □

Remark 6.2 (Structure of the Uniqueness Proof)**.** Theorem 6.1 makes the two-part structure of every characterization theorem explicit. Existence (a) is proved by verifying, axiom by

axiom, that f^* satisfies each one. Uniqueness (b) is proved by showing that the axioms force $f = f^*$, typically by a constructive derivation. Part (c) establishes that $\mathrm{Th}(\{f^*\})$ is the canonical closed system, and that every axiom system characterizing f^* is contained in it. In practice, one seeks a small, interpretable subset of $\mathrm{Th}(\{f^*\})$ that still forces uniqueness.

7 Impossibility Cores and Independence

We prove two theorems that are, in practice, the workhorse results of the axiomatic tradition.

Theorem 7.1 (Impossibility Core Existence)**.** *Let Σ be an axiom system on a finite domain $\mathcal{U}$ with $\mathcal{M}(\Sigma) = \emptyset$. Then Σ contains a* minimal *impossible sub-system $\Sigma_0 \subseteq \Sigma$, i.e., a sub-system satisfying $\mathcal{M}(\Sigma_0) = \emptyset$ such that $\mathcal{M}(\Sigma_0 \setminus \{A\}) \neq \emptyset$ for every $A \in \Sigma_0$. In particular, every impossible system contains an impossibility core.*

Proof. Since Σ is finite (by definition) and impossible, the family $\mathcal{I} = \{\Sigma' \subseteq \Sigma \mid \mathcal{M}(\Sigma') = \emptyset\}$ is nonempty (it contains Σ) and finite, so it contains an element Σ_0 of minimum cardinality. By minimality of $|\Sigma_0|$, no proper sub-system of Σ_0 is impossible, i.e., $\mathcal{M}(\Sigma_0 \setminus \{A\}) \neq \emptyset$ for every $A \in \Sigma_0$. □

Remark 7.2. In the present finite-domain setting, every axiom system Σ is finite by definition, so the existence of an impossibility core is immediate from the well-ordering of finite sets. The content of the theorem is that it establishes the terminology and the structural guarantee invoked throughout Sections 9–14. Once we relax the finiteness assumptions, the analogous statement can fail without additional compactness hypotheses. For example, if $\mathcal{U} = \mathbb{N}$ and $A_n = \mathbb{N} \setminus \{n\}$ for each $n \in \mathbb{N}$, then the family $\{A_n\}_{n\in\mathbb{N}}$ (which is infinite, hence outside our standing assumption that Σ is finite) satisfies $\mathcal{M}(\{A_n\}_{n\in\mathbb{N}}) = \emptyset$, yet every finite sub-family has nonempty intersection. The system therefore admits no finite impossibility core.

Theorem 7.3 (Independence Witnessing)**.** *Let Σ be an axiom system and let $A \in \Sigma$. The following are equivalent:*

(i) The axiom A is independent in Σ.

(ii) There exists a witness object $w_A \in \mathcal{M}(\Sigma \setminus \{A\}) \setminus A$.

(iii) $\mathcal{M}(\Sigma \setminus \{A\}) \not\subseteq A$.

In particular, Σ is independent if and only if, for every $A \in \Sigma$, such a witness w_A exists.

Proof. We show (i)⇔(ii)⇔(iii). For (i)⇒(iii), if A is independent, then $\mathcal{M}(\Sigma\setminus\{A\}) \supsetneq \mathcal{M}(\Sigma)$. Since $\mathcal{M}(\Sigma) = \mathcal{M}(\Sigma \setminus \{A\}) \cap A$ (by Remark 2.4, because $\Sigma = (\Sigma \setminus \{A\}) \cup \{A\}$), the strict inclusion $\mathcal{M}(\Sigma \setminus \{A\}) \supsetneq \mathcal{M}(\Sigma \setminus \{A\}) \cap A$ implies $\mathcal{M}(\Sigma \setminus \{A\}) \not\subseteq A$. For (iii)⇒(ii), if $\mathcal{M}(\Sigma\setminus\{A\}) \not\subseteq A$, then there exists $w_A \in \mathcal{M}(\Sigma\setminus\{A\})$ with $w_A \notin A$, so $w_A \in \mathcal{M}(\Sigma\setminus\{A\})\setminus A$. For (ii)⇒(i), the witness w_A satisfies $w_A \in \mathcal{M}(\Sigma\setminus\{A\})$ (it satisfies all axioms except possibly A) and $w_A \notin A$ (so $w_A \notin \mathcal{M}(\Sigma)$), giving $\mathcal{M}(\Sigma \setminus \{A\}) \supsetneq \mathcal{M}(\Sigma)$. □

8 The Cooperative Game on the Axiom Set

We now develop one of the central contributions of this paper. An axiom system Σ naturally defines a cooperative game on the player set Σ, where the "players" are the axioms and their joint "worth" is the degree to which they collapse the model class. This game-theoretic structure provides a canonical attribution of credit (or blame) to each axiom, resolves the ambiguity in the order of axiom application, and yields a quantitative theory of axiom strength with provably good properties.

Throughout this section, we fix a finite axiom system $\Sigma = \{A_1, A_2, \ldots, A_n\}$ and a probability measure μ on $\mathcal{U}$ with $\mu(\mathcal{U}) = 1$ and $\mu(A) > 0$ for every $A \in \Sigma$. Note that $\mu(A) > 0$ does not exclude null axioms (those with zero marginal collapse in every coalition); an axiom $A = \mathcal{U}$ satisfies $\mu(A) = 1$ yet contributes nothing.

8.1 The Axiom Game

We now define the central object of this section.

Definition 8.1 (The Axiom Game)**.** The *axiom game* of (Σ, μ) is the cooperative game (Σ, v_μ) whose characteristic function is

$$v_\mu(S) = \mu(\mathcal{U}) - \mu(\mathcal{M}(S)) = 1 - \mu\left(\bigcap_{A \in S} A\right), \qquad S \subseteq \Sigma,$$

with the convention $v_\mu(\emptyset) = 0$. We call $v_\mu(S)$ the *collapse* induced by the sub-system S.

Intuitively, $v_\mu(S)$ measures the fraction of the universe that the axioms in S jointly exclude. When $\mathcal{M}(\Sigma) = \{f^*\}$, the full axiom system achieves $v_\mu(\Sigma) = 1 - \mu(\{f^*\})$; when $\mathcal{M}(\Sigma) = \emptyset$, it achieves $v_\mu(\Sigma) = 1$.

Proposition 8.2 (Properties of the Axiom Game)**.** *The axiom game (Σ, v_μ) satisfies the following.*

(i) Normalization. $v_\mu(\emptyset) = 0$.

(ii) Monotonicity. *If $S \subseteq T \subseteq \Sigma$, then $v_\mu(S) \leq v_\mu(T)$.*

(iii) Subadditivity. $v_\mu(S \cup T) \leq v_\mu(S) + v_\mu(T)$ *for all $S, T \subseteq \Sigma$.*

(iv) Submodularity (Concavity). $v_\mu(S \cup T) + v_\mu(S \cap T) \leq v_\mu(S) + v_\mu(T)$ *for all $S, T \subseteq \Sigma$.*

Proof. Normalization (i) follows from $\mathcal{M}(\emptyset) = \mathcal{U}$ and $\mu(\mathcal{U}) = 1$.

Monotonicity (ii) follows from Theorem 4.1(i). Indeed, $S \subseteq T$ implies $\mathcal{M}(T) \subseteq \mathcal{M}(S)$, so $\mu(\mathcal{M}(T)) \leq \mu(\mathcal{M}(S))$, and hence $v_\mu(T) \geq v_\mu(S)$.

Subadditivity (iii) follows from (iv) together with (i) and (ii). Indeed, submodularity gives $v_\mu(S \cup T) + v_\mu(S \cap T) \leq v_\mu(S) + v_\mu(T)$, and $v_\mu(S \cap T) \geq 0$ by normalization and monotonicity, whence $v_\mu(S \cup T) \leq v_\mu(S) + v_\mu(T)$. We give a direct proof. Since we have

$$\mathcal{M}(S \cup T) = \mathcal{M}(S) \cap \mathcal{M}(T) \quad \text{and} \quad \mu(\mathcal{M}(S) \cup \mathcal{M}(T)) \leq \mu(\mathcal{U}) = 1,$$

inclusion-exclusion gives $\mu(\mathcal{M}(S)\cap\mathcal{M}(T)) \geq \mu(\mathcal{M}(S))+\mu(\mathcal{M}(T))-1$. Negating and adding 1, we obtain

$$v_\mu(S\cup T) = 1-\mu(\mathcal{M}(S)\cap\mathcal{M}(T)) \leq 1-[\mu(\mathcal{M}(S))+\mu(\mathcal{M}(T))-1] = v_\mu(S)+v_\mu(T).$$

For submodularity (iv), note that $\mathcal{M}(S\cup T) = \mathcal{M}(S)\cap\mathcal{M}(T)$ by Remark 2.4. Furthermore, by antitonicity of $\mathcal{M}$ applied to the axiom-set inclusion $S\cap T\subseteq S$ (where $S\cap T$ denotes the intersection of S and T as subsets of Σ, not as subsets of $\mathcal{U}$), we get $\mathcal{M}(S)\subseteq\mathcal{M}(S\cap T)$, and similarly $\mathcal{M}(T)\subseteq\mathcal{M}(S\cap T)$, so

$$\mathcal{M}(S\cap T)\supseteq\mathcal{M}(S)\cup\mathcal{M}(T) \quad \text{and} \quad \mu(\mathcal{M}(S\cap T))\geq\mu(\mathcal{M}(S)\cup\mathcal{M}(T)).$$

We can now apply the principle of inclusion-exclusion, to get

$$\mu(\mathcal{M}(S)\cap\mathcal{M}(T))+\mu(\mathcal{M}(S)\cup\mathcal{M}(T)) = \mu(\mathcal{M}(S))+\mu(\mathcal{M}(T)).$$

It follows that

$$\begin{aligned} v_\mu(S\cup T)+v_\mu(S\cap T) &= 2-\mu(\mathcal{M}(S)\cap\mathcal{M}(T))-\mu(\mathcal{M}(S\cap T)) \\ &\leq 2-\mu(\mathcal{M}(S)\cap\mathcal{M}(T))-\mu(\mathcal{M}(S)\cup\mathcal{M}(T)) \\ &= 2-\mu(\mathcal{M}(S))-\mu(\mathcal{M}(T)) \\ &= v_\mu(S)+v_\mu(T). \end{aligned}$$

□

The submodularity of v_μ is a fundamental structural property of the axiom game, derived purely from the intersection structure of the model-class map. In fact, submodularity is equivalent to the condition that all pairwise marginal interactions are nonpositive, which is precisely the content of the Axiom Interaction Theorem proved below. In this sense, Proposition 8.2(iv) is the global statement of diminishing returns. The collapse from merging two coalitions never exceeds the sum of their individual collapses. The Axiom Interaction Theorem makes this law pointwise and explicit.

8.2 The Axiom Shapley Value

Since (Σ, v_μ) is a well-defined cooperative game, we may apply the Shapley value to it.

Definition 8.3 (Axiom Shapley Value)**.** The *axiom Shapley value* of the axiom system (Σ,μ) is the vector $\phi(\Sigma,\mu)\in\mathbb{R}^\Sigma$ defined by

$$\phi_A(\Sigma,\mu) = \sum_{S\subseteq\Sigma\setminus\{A\}} \frac{|S|!\,(n-|S|-1)!}{n!}\big[v_\mu(S\cup\{A\})-v_\mu(S)\big], \qquad A\in\Sigma,$$

where $n=|\Sigma|$.

Thus, $\phi_A(\Sigma,\mu)$ is the average marginal collapse contributed by adding axiom A to a uniformly random sub-system of $\Sigma\setminus\{A\}$. In a uniqueness context with $\mathcal{M}(\Sigma)=\{f^*\}$, the vector $\phi(\Sigma,\mu)$ assigns to each axiom its average contribution to forcing the characterization. In an impossibility context with $\mathcal{M}(\Sigma)=\emptyset$, it assigns to each axiom its average contribution to the contradiction.

We now state the main theorem of this section.

Theorem 8.4 (Axiom Shapley Value)**.** *Let (Σ, v_μ) be the axiom game. The axiom Shapley value $\phi(\Sigma, \mu)$ is the unique function on axiom systems satisfying the following four axioms.*

(i) Efficiency. *$\sum_{A\in\Sigma} \phi_A(\Sigma, \mu) = v_\mu(\Sigma)$; that is, the total attribution equals the total collapse.*

(ii) Symmetry. *If two axioms $A, B \in \Sigma$ satisfy $v_\mu(S \cup \{A\}) = v_\mu(S \cup \{B\})$ for every $S \subseteq \Sigma \setminus \{A, B\}$, then $\phi_A(\Sigma, \mu) = \phi_B(\Sigma, \mu)$.*

(iii) Linearity. *The Shapley value is linear in the characteristic function v_μ.*

(iv) Null axiom. *If $A \in \Sigma$ satisfies $v_\mu(S \cup \{A\}) = v_\mu(S)$ for every $S \subseteq \Sigma \setminus \{A\}$ (i.e., A contributes zero marginal collapse in every coalition), then $\phi_A(\Sigma, \mu) = 0$.*

The following further properties are consequences of the axiom game structure.

(a) Nonnegativity. *$\phi_A(\Sigma, \mu) \geq 0$ for every $A \in \Sigma$, with $\phi_A(\Sigma, \mu) = 0$ if and only if A contributes zero marginal collapse in every sub-system (i.e., $v_\mu(S \cup \{A\}) = v_\mu(S)$ for all $S \subseteq \Sigma \setminus \{A\}$).*

(b) Independence detection. *If an axiom A is independent in Σ (in the sense of Definition 2.5), then for every measure μ satisfying $\mu(\mathcal{M}(\Sigma \setminus \{A\}) \setminus \mathcal{M}(\Sigma)) > 0$, it holds that $\phi_A(\Sigma, \mu) > 0$. In fact, independence implies a strictly positive full-coalition marginal under any such μ, and hence a positive Shapley value. The converse does not hold in general. An axiom may be redundant in Σ (i.e., $\mathcal{M}(\Sigma \setminus \{A\}) = \mathcal{M}(\Sigma)$) yet still receive a positive Shapley value, because it may contribute positive marginal collapse in sub-coalitions that do not contain all of $\Sigma \setminus \{A\}$.*

Proof. The uniqueness in (i)–(iv) is Shapley's theorem (Shapley, 1953) applied to the game (Σ, v_μ); the axiom Shapley value of Definition 8.3 is the unique function satisfying efficiency, symmetry, linearity in the game, and the null-player property. Note that linearity in v_μ implies linearity in μ as well, since $v_{\alpha\mu+(1-\alpha)\mu'}(S) = \alpha\, v_\mu(S) + (1-\alpha)\, v_{\mu'}(S)$ for any $S \neq \emptyset$. In particular,

$$\phi(\Sigma, \alpha\mu + (1-\alpha)\mu') = \alpha\, \phi(\Sigma, \mu) + (1-\alpha)\, \phi(\Sigma, \mu').$$

For (a), monotonicity of v_μ (Proposition 8.2(ii)) gives every marginal contribution $v_\mu(S \cup \{A\}) - v_\mu(S) \geq 0$, so the weighted average $\phi_A \geq 0$. The null condition holds if and only if every marginal contribution is zero.

For (b), if A is independent in Σ, then $\mathcal{M}(\Sigma \setminus \{A\}) \supsetneq \mathcal{M}(\Sigma)$, so for any μ with $\mu(\mathcal{M}(\Sigma \setminus \{A\}) \setminus \mathcal{M}(\Sigma)) > 0$ we have $v_\mu(\Sigma) > v_\mu(\Sigma \setminus \{A\})$, and the full-coalition marginal of A is strictly positive. Since ϕ_A is a weighted average of nonnegative marginals including this strictly positive one, we obtain $\phi_A > 0$. For the failure of the converse, note that a redundant axiom A satisfies $\mathcal{M}(\Sigma \setminus \{A\}) = \mathcal{M}(\Sigma)$ and hence has zero full-coalition marginal, but may contribute strictly positive marginals in sub-coalitions $S \subsetneq \Sigma \setminus \{A\}$; see Example 8.10 for an instance of this phenomenon (the axioms A_3 and A_4 are redundant in Σ yet have $\phi_{A_3} = \phi_{A_4} = 1/6 > 0$). □

Remark 8.5 (Monotone Attribution and Diminishing Returns)**.** The Shapley value is the natural measure of individual axiom contribution in a setting of diminishing returns.

Nonnegativity of every marginal gives $\phi_A \geq 0$, and submodularity of v_μ means all pairwise interactions are nonpositive, which is the content of Theorem 8.8 below.

Remark 8.6 (Interpretation in the Four Regimes)**.** In the uniqueness regime, the efficiency property reads $\sum_A \phi_A = 1 - \mu(\{f^*\})$. The axioms together account for the full collapse from $\mathcal{U}$ to $\{f^*\}$, and each axiom receives its average marginal share. In the impossibility regime, $\sum_A \phi_A = 1$. The axioms together account for the full collapse to $\emptyset$, and the axiom Shapley value distributes blame for the impossibility. The axiom A with the largest ϕ_A is the one most responsible, on average, for the collapse. It is important to note that this does not automatically identify the best axiom to drop in a rescue operation; as shown in Section 9 (see Remark 9.5 and Example 9.10), the correct rescue criterion is the full-coalition marginal Δ_A, which may rank axioms differently.

8.3 The Axiom Interaction Theorem

The submodularity of the axiom game (Proposition 8.2(iv)) is the global statement of diminishing returns. We now make this law pointwise and explicit at the level of individual axiom pairs. The fundamental finding is a universal law of *diminishing marginal returns*. Each axiom contributes less collapse when others are already in force.

Definition 8.7 (Axiom Interaction)**.** The *pairwise interaction* of two axioms $A, B \in \Sigma$ conditional on a sub-system $S \subseteq \Sigma \setminus \{A, B\}$ is

$$\iota_\mu(A, B \mid S) = v_\mu(S \cup \{A, B\}) - v_\mu(S \cup \{A\}) - v_\mu(S \cup \{B\}) + v_\mu(S).$$

The *total pairwise interaction* of A and B is

$$\iota_\mu(A, B) = \sum_{S \subseteq \Sigma \setminus \{A,B\}} \frac{|S|!\,(n - |S| - 2)!}{(n-1)!}\, \iota_\mu(A, B \mid S),$$

the average of $\iota_\mu(A, B \mid S)$ over all sub-systems S, weighted by the Shapley-type combinatorial weights for the reduced $(n-1)$-player game on $\Sigma \setminus \{B\}$ (cf. the interaction index of Grabisch and Roubens 1999).

Theorem 8.8 (Axiom Interaction)**.** *Let $A, B \in \Sigma$ and $S \subseteq \Sigma \setminus \{A, B\}$. Write $C = \mathcal{M}(S)$, $C_A = \mathcal{M}(S \cup \{A\})$, $C_B = \mathcal{M}(S \cup \{B\})$, $C_{AB} = \mathcal{M}(S \cup \{A, B\})$. Note that $C_{AB} = C_A \cap C_B$ and $C_{AB} \subseteq C_A, C_B \subseteq C$.*

(i) Explicit formula.

$$\iota_\mu(A, B \mid S) = \mu(C_A \cup C_B) - \mu(C) = -\mu(C \setminus (C_A \cup C_B)).$$

The second equality holds because $C_A \cup C_B \subseteq C$, so the interaction is the negative μ-measure of the set of objects in the current model class that are excluded by both *axioms simultaneously.*

(ii) Diminishing returns. *$\iota_\mu(A, B \mid S) \leq 0$, which is immediate from the formula above. Equivalently,*

$$v_\mu(S \cup \{A, B\}) - v_\mu(S \cup \{A\}) \leq v_\mu(S \cup \{B\}) - v_\mu(S).$$

(iii) Characterization of equality. *$\iota_\mu(A, B \mid S) = 0$ if and only if $C_A \cup C_B = C$ μ-almost surely, i.e., the two axioms together already cover the entire model class C without any synergy.*

(iv) Strict diminishing returns. *$\iota_\mu(A, B \mid S) < 0$ whenever $\mu(C \setminus (C_A \cup C_B)) > 0$, i.e., whenever both axioms individually fail to cover some positive-measure part of the current model class.*

(v) Total interaction. *Assume $|\Sigma| \geq 2$. Then $\iota_\mu(A, B) \leq 0$, with equality if and only if $C_A \cup C_B = C$ μ-a.s. for every $S \subseteq \Sigma \setminus \{A, B\}$.*

Proof. For (i), we substitute $v_\mu(\cdot) = 1 - \mu(\mathcal{M}(\cdot))$ and $C_{AB} = C_A \cap C_B$ to obtain

$$\begin{aligned}\iota_\mu(A, B \mid S) &= [1 - \mu(C_{AB})] - [1 - \mu(C_A)] - [1 - \mu(C_B)] + [1 - \mu(C)] \\ &= \mu(C_A) + \mu(C_B) - \mu(C_A \cap C_B) - \mu(C).\end{aligned}$$

By inclusion-exclusion, $\mu(C_A) + \mu(C_B) - \mu(C_A \cap C_B) = \mu(C_A \cup C_B)$. It therefore holds that $\iota_\mu(A, B \mid S) = \mu(C_A \cup C_B) - \mu(C)$.

For (ii), since $C_A, C_B \subseteq C$, we have $C \setminus (C_A \cup C_B) \subseteq C$, hence $\mu(C \setminus (C_A \cup C_B)) \geq 0$, and the result follows from part (i). The equivalent formulation follows by rearranging the definition of ι_μ.

For (iii), $\iota_\mu(A, B \mid S) = \mu(C_A \cup C_B) - \mu(C) = 0$ if and only if $\mu(C_A \cup C_B) = \mu(C)$, if and only if $C_A \cup C_B = C$ μ-a.s.

For (iv), $\iota_\mu(A, B \mid S) = -\mu(C \setminus (C_A \cup C_B)) < 0$ when $\mu(C \setminus (C_A \cup C_B)) > 0$.

For (v), $\iota_\mu(A, B) \leq 0$ because it is a convex combination (with strictly positive weights) of nonpositive terms; the equality characterization follows from (iii) applied to each S, since every term must vanish individually. □

Remark 8.9 (The Universal Law of Diminishing Axiom Returns)**.** Theorem 8.8 establishes a universal law. In any axiom system, every pair of axioms exhibits diminishing marginal returns, derived purely from the intersection structure of the model-class map and the monotonicity of μ. The equality case $\iota_\mu(A, B \mid S) = 0$ occurs in the *coverage-exhaustive* configuration $C_A \cup C_B = C$, where the survivor sets together cover $\mathcal{M}(S)$. The strict case $\iota_\mu(A, B \mid S) < 0$ is the generic one, arising whenever a positive-measure set of objects in $\mathcal{M}(S)$ is excluded by both A and B simultaneously. In fact, Proposition 8.2(iv) and Theorem 8.8 are two formulations of the same structural property. A set function v is submodular if and only if all its pairwise interaction terms $\iota(A, B \mid S)$ are nonpositive; this is a standard equivalence in the theory of submodular functions. The former gives the inequality on arbitrary coalition pairs; the latter unpacks it as an explicit formula and identifies the exact gap $\mu(C \setminus (C_A \cup C_B))$.

Example 8.10 (Running Example: Axiom Game and Shapley Value)**.** Returning to Example 3.1, we equip $\mathcal{U} = \{1, 2, 3, 4, 5, 6\}$ with the uniform measure $\mu(\{i\}) = 1/6$. The singleton collapse values are

$$v_\mu(\{A_1\}) = v_\mu(\{A_2\}) = 1/2 \quad \text{and} \quad v_\mu(\{A_3\}) = v_\mu(\{A_4\}) = 1/3.$$

Since $A_1 \cap A_2 = \emptyset$, we have $v_\mu(\{A_1, A_2\}) = 1$. Averaging marginal contributions over all $4! = 24$ orderings yields the axiom Shapley values

$$\phi_{A_1} = \tfrac{1}{3}, \quad \phi_{A_2} = \tfrac{1}{3}, \quad \phi_{A_3} = \tfrac{1}{6}, \quad \phi_{A_4} = \tfrac{1}{6},$$

which sum to $1 = v_\mu(\Sigma)$, confirming efficiency. The axioms A_1 and A_2 bear the largest blame, reflecting that they form the impossibility core, while A_3 and A_4 are redundant (in the sense that their removal does not restore possibility) yet still receive positive Shapley values because they contribute marginal collapse in sub-coalitions that do not yet contain the core. By Theorem 8.8, the pairwise interaction $\iota_\mu(A_1, A_2 \mid \emptyset) = \mu(A_1 \cup A_2) - \mu(\mathcal{U}) = 1 - 1 = 0$, since $A_1 \cup A_2 = \mathcal{U}$; the two core axioms are coverage-exhaustive relative to the empty coalition.

As a contrast, consider the sub-system $\{A_3, A_4\}$, which lies in the *possibility regime*: its model class is $\mathcal{M}(\{A_3, A_4\}) = \{1, 4\}$ and its collapse is $v_\mu(\{A_3, A_4\}) = 2/3$. Here, it holds that $v_\mu(\{A_3\}) = v_\mu(\{A_4\}) = 1/3$, so each axiom has the same individual collapse, and by symmetry it also holds that $\phi_{A_3}(\{A_3, A_4\}, \mu) = \phi_{A_4}(\{A_3, A_4\}, \mu) = 1/3$. The total attribution $2/3$ is equal to $v_\mu(\{A_3, A_4\}) < 1$, reflecting that the axioms jointly narrow the model class from six objects to two but do not force uniqueness or impossibility. This illustrates the axiom Shapley value in a non-degenerate possibility regime, where the efficiency condition distributes only a partial collapse.

8.4 Axiom Dominance and Robust Rankings

The Shapley value assigns each axiom a numerical credit, but a natural question is whether the resulting *ranking* of axioms depends on the choice of measure μ or is a robust feature of the axiom system's structure. We now give an exact formula for the Shapley-value difference of two axioms and derive a clean sufficient condition under which the ranking is invariant over a class of measures.

Definition 8.11 (Uniform Restrictiveness)**.** Let (Σ, v_μ) be an axiom game. An axiom $A \in \Sigma$ is said to be *uniformly more restrictive* than $B \in \Sigma$, written $A \succeq_\mu B$, if it holds that $\mu(\mathcal{M}(S \cup \{A\})) \leq \mu(\mathcal{M}(S \cup \{B\}))$ for every $S \subseteq \Sigma \setminus \{A, B\}$.

In other words, A admits fewer models than B in every context, regardless of which other axioms are already in force.

Theorem 8.12 (Axiom Dominance)**.** *Let (Σ, v_μ) be an axiom game with $|\Sigma| = n$, and let $A, B \in \Sigma$ be distinct.*

(i) Exact difference formula.

$$\phi_A(\Sigma, \mu) - \phi_B(\Sigma, \mu) = \sum_{S \subseteq \Sigma \setminus \{A,B\}} \frac{|S|!\,(n - |S| - 2)!}{(n-1)!} \big[\mu(\mathcal{M}(S \cup \{B\})) - \mu(\mathcal{M}(S \cup \{A\}))\big].$$

The right-hand side is a convex combination (positive weights summing to 1) of the model-class differences $\mu(\mathcal{M}(S \cup \{B\})) - \mu(\mathcal{M}(S \cup \{A\}))$ over all contexts S.

(ii) Dominance implies ranking. *If $A \succeq_\mu B$, then $\phi_A(\Sigma,\mu) \geq \phi_B(\Sigma,\mu)$, with equality if and only if $\mu(\mathcal{M}(S\cup\{A\})) = \mu(\mathcal{M}(S\cup\{B\}))$ for every $S \subseteq \Sigma\setminus\{A,B\}$. Consequently, if the dominance chain $A_1 \succeq_\mu A_2 \succeq_\mu \cdots \succeq_\mu A_n$ holds for every μ in a class $\mathcal{Q}$ of measures, then the Shapley ranking $\phi_{A_1}(\Sigma,\mu) \geq \phi_{A_2}(\Sigma,\mu) \geq \cdots \geq \phi_{A_n}(\Sigma,\mu)$ holds for every $\mu \in \mathcal{Q}$.*

Proof. For (i), we split the Shapley formula $\phi_A = \sum_{S\subseteq\Sigma\setminus\{A\}} w_{|S|}[v_\mu(S\cup\{A\}) - v_\mu(S)]$ according to whether $B \in S$ or $B \notin S$. The sub-systems $S \subseteq \Sigma\setminus\{A\}$ not containing B are indexed by $S \subseteq \Sigma\setminus\{A,B\}$, and those containing B are $T\cup\{B\}$ for $T \subseteq \Sigma\setminus\{A,B\}$. Subtracting the analogous decomposition of ϕ_B and combining terms gives

$$\begin{aligned}\phi_A - \phi_B &= \sum_{S\subseteq\Sigma\setminus\{A,B\}} w_{|S|}\big[v_\mu(S\cup\{A\}) - v_\mu(S\cup\{B\})\big] \\ &\quad + \sum_{S\subseteq\Sigma\setminus\{A,B\}} w_{|S|+1}\big[v_\mu(S\cup\{A\}) - v_\mu(S\cup\{B\})\big] \\ &= \sum_{S\subseteq\Sigma\setminus\{A,B\}} (w_{|S|} + w_{|S|+1})\big[v_\mu(S\cup\{A\}) - v_\mu(S\cup\{B\})\big],\end{aligned}$$

where $w_s = s!(n-s-1)!/n!$. A direct calculation gives $w_s + w_{s+1} = s!(n-s-2)!/(n-1)!$. Substituting $v_\mu(S\cup\{A\}) - v_\mu(S\cup\{B\}) = \mu(\mathcal{M}(S\cup\{B\})) - \mu(\mathcal{M}(S\cup\{A\}))$ yields the formula. The weights are positive and sum to 1 (they are the Shapley weights for an $(n-1)$-player game on $\Sigma\setminus\{B\}$).

Part (ii) is immediate from (i). If $A \succeq_\mu B$, then $\mu(\mathcal{M}(S\cup\{B\})) - \mu(\mathcal{M}(S\cup\{A\})) \geq 0$ for every term in the convex combination, so their weighted sum is nonnegative; equality forces every term to vanish. The final claim about $\mathcal{Q}$ then follows by applying (ii) pointwise for each $\mu \in \mathcal{Q}$. □

Remark 8.13 (Interpretation)**.** Theorem 8.12(ii) is the "desirability" property of the Shapley value (Shapley, 1953); the contribution here is the formulation in terms of model-class sizes and the robust-ranking consequence (iii), which transforms measure dependence into a transparent structural condition. In practice, the dominance chain is far easier to verify than a full Shapley computation and certifies the ranking for *all* measures in the class simultaneously.

8.5 The Geometry of Attribution: Polymatroids and the Base Polytope

By Proposition 8.2, the axiom game (Σ, v_μ) is in particular a *polymatroid*, i.e., a normalized, monotone, submodular set function. This places the axiom Shapley value in a rich geometric framework.

For each ordering $\sigma\colon \{1,2,\ldots,n\} \to \Sigma$ of the axioms, define the *marginal vector* $m^\sigma \in \mathbb{R}^\Sigma$ by

$$m^\sigma_A = v_\mu(\{\sigma(1),\sigma(2),\ldots,\sigma(k)\}) - v_\mu(\{\sigma(1),\sigma(2),\ldots,\sigma(k-1)\}),$$

where k is the index with $\sigma(k) = A$. In other words, m^σ_A is the marginal collapse contributed by A when the axioms are added in the order σ. The *base polytope* of v_μ is

$$B(v_\mu) = \Big\{x \in \mathbb{R}^\Sigma \;\Big|\; \textstyle\sum_{A\in\Sigma} x_A = v_\mu(\Sigma),\ \sum_{A\in S} x_A \leq v_\mu(S) \text{ for all } S \subseteq \Sigma\Big\}.$$

Proposition 8.14 (Axiom Shapley Value as Centroid)**.** *The vertices of the base polytope $B(v_\mu)$ are precisely the distinct marginal vectors m^σ, one for each ordering σ of the axioms (with possible repetitions when different orderings yield the same vector). The axiom Shapley value is the centroid of the $n!$ marginal vectors, counting multiplicities:*

$$\phi_A(\Sigma, \mu) = \frac{1}{n!} \sum_{\sigma \in S_n} m_A^\sigma \quad \textit{for every } A \in \Sigma.$$

Proof. We verify that each marginal vector m^σ lies in $B(v_\mu)$ and is a vertex, and then derive the centroid formula.

First, $m^\sigma \in B(v_\mu)$. The efficiency constraint $\sum_{A \in \Sigma} m_A^\sigma = v_\mu(\Sigma)$ holds because the marginal contributions telescope,

$$\begin{aligned} \sum_{A \in \Sigma} m_A^\sigma &= \sum_{j=1}^{n} \big[v_\mu(\{\sigma(1), \sigma(2), \ldots, \sigma(j)\}) - v_\mu(\{\sigma(1), \sigma(2), \ldots, \sigma(j-1)\}) \big] \\ &= v_\mu(\Sigma) - v_\mu(\emptyset) = v_\mu(\Sigma). \end{aligned}$$

For the coalition constraints, let $S \subseteq \Sigma$ and write $S = \{\sigma(j_1), \sigma(j_2), \ldots, \sigma(j_r)\}$ with $j_1 < j_2 < \cdots < j_r$. By submodularity of v_μ, each marginal $m_{\sigma(j_\ell)}^\sigma$ is at most the marginal of $\sigma(j_\ell)$ relative to the smaller coalition $\{\sigma(j_1), \sigma(j_2), \ldots, \sigma(j_{\ell-1})\}$, so $\sum_{A \in S} m_A^\sigma \le v_\mu(S)$.

Second, m^σ is a vertex of $B(v_\mu)$. For the ordering σ, define $S_j = \{\sigma(1), \sigma(2), \ldots, \sigma(j)\}$. The vector m^σ satisfies the equality $\sum_{A \in S_j} m_A^\sigma = v_\mu(S_j)$ for every $j = 1, 2, \ldots, n$, giving n tight constraints that are linearly independent by the chain structure of the S_j (together with efficiency). Since $B(v_\mu)$ lies in $\mathbb{R}^n$, these n tight constraints determine a unique point, which is therefore a vertex. Conversely, every vertex of $B(v_\mu)$ arises from some ordering by the greedy algorithm (Edmonds, 1970).

The centroid formula follows from Definition 8.3. The Shapley value of A is by definition equal to $\phi_A = \sum_{S \subseteq \Sigma \setminus \{A\}} \frac{|S|!(n-|S|-1)!}{n!} [v_\mu(S \cup \{A\}) - v_\mu(S)]$. Each ordering σ with $\sigma(k) = A$ contributes exactly one marginal

$$m_A^\sigma = v_\mu(\{\sigma(1), \sigma(2), \ldots, \sigma(k)\}) - v_\mu(\{\sigma(1), \sigma(2), \ldots, \sigma(k-1)\}),$$

and the number of orderings placing A in position k with a given predecessor set S of size $k-1$ is $|S|!(n-|S|-1)!$. Summing over all $n!$ orderings with equal weight $1/n!$ gives $\phi_A = \frac{1}{n!} \sum_\sigma m_A^\sigma$. □

Remark 8.15 (Geometric Meaning)**.** The base polytope $B(v_\mu)$ is the set of all attribution vectors that are collectively efficient ($\sum_A x_A = v_\mu(\Sigma)$) and individually bounded above by marginal contributions. Among all points in $B(v_\mu)$, the axiom Shapley value $\phi(\Sigma, \mu)$ is the unique one that is symmetric under permutations of axioms with equal marginal contributions. In the impossibility regime, $B(v_\mu)$ is a convex polytope inside the hyperplane $\sum_A x_A = 1$; its vertices are the $n!$ greedy blame vectors, and ϕ is their average. This gives a canonical geometric meaning to the notion that the Shapley value is the "fairest" attribution.

8.6 Characterization of Axiom Games

Proposition 8.2 shows that every axiom game is normalized, monotone, and submodular. A natural question is whether these three conditions are also sufficient, i.e., whether every normalized, monotone, submodular set function with values in $[0,1]$ arises from some axiom system and measure. The answer turns out to be negative for $n \geq 3$, and the correct characterization involves a condition strictly stronger than submodularity.

Definition 8.16 (Möbius Coefficients)**.** Let $b\colon 2^\Sigma \to \mathbb{R}$ be a set function with $|\Sigma| = n$. The *Möbius coefficients* of b are the values

$$c(T) = \sum_{S \supseteq T} (-1)^{|S|-|T|}\, b(S), \qquad T \subseteq \Sigma.$$

We say b has *nonnegative Möbius coefficients* if $c(T) \geq 0$ for every $T \subseteq \Sigma$.

Theorem 8.17 (Characterization of Axiom Games)**.** *A function $v\colon 2^\Sigma \to [0,1]$ is an axiom game (that is, $v = v_\mu$ for some domain $\mathcal{U}$, probability measure μ on $\mathcal{U}$, and axiom system $\Sigma \subseteq \mathcal{P}(\mathcal{U}) \setminus \{\emptyset\}$) if and only if the following two conditions hold.*

(i) $v(\emptyset) = 0$.

(ii) The function $b = 1 - v$ has nonnegative Möbius coefficients in the sense of Definition 8.16.

In particular, every axiom game is monotone and submodular (Proposition 8.2), and these properties follow from condition (ii). However, for $|\Sigma| \geq 3$, condition (ii) is strictly stronger than monotonicity and submodularity combined.

Proof. For the forward direction (necessity), let $v = v_\mu$ be an axiom game. Note that $b(S) = 1 - v(S) = \mu(\mathcal{M}(S)) = \mu(\bigcap_{i \in S} A_i)$ for every $S \subseteq \Sigma$. By Möbius inversion on the Boolean lattice (with respect to the superset ordering), the Möbius coefficient $c(T)$ of b satisfies

$$c(T) = \mu\Big(\bigcap_{i \in T} A_i \;\cap\; \bigcap_{j \notin T} (\mathcal{U} \setminus A_j) \Big),$$

which is the μ-measure of the set of objects satisfying exactly the axioms indexed by T. Indeed, the sets $\{\bigcap_{i\in T} A_i \cap \bigcap_{j\notin T}(\mathcal{U}\setminus A_j)\}_{T\subseteq\Sigma}$ partition $\mathcal{U}$, and $b(S) = \sum_{T \supseteq S} c(T)$ holds because every object satisfying all axioms in S belongs to exactly one cell of the partition, namely the cell indexed by the full set of axioms it satisfies. Since μ is a probability measure, $c(T) \geq 0$.

For the converse (sufficiency), let $v\colon 2^{[n]} \to [0,1]$ satisfy (i)–(ii). We construct an explicit axiom system realizing v. Let $\mathcal{U}$ be a set of 2^n elements, one element u_T for each subset $T \subseteq [n]$. Define the axioms by $A_i = \{u_T \mid i \in T\}$ for each $i \in [n]$, and define the probability measure $\mu(\{u_T\}) = c(T)$, where $c(T)$ is the Möbius coefficient of $b = 1 - v$. By condition (ii), $c(T) \geq 0$ for all T, and $\sum_{T \subseteq [n]} c(T) = b(\emptyset) = 1 - v(\emptyset) = 1$, so μ is indeed a probability measure. For any $S \subseteq [n]$, we have $\mathcal{M}(S) = \bigcap_{i\in S} A_i = \{u_T \mid S \subseteq T\}$, so $\mu(\mathcal{M}(S)) = \sum_{T \supseteq S} c(T) = b(S)$ by Möbius inversion, and $v_\mu(S) = 1 - b(S) = v(S)$. Each axiom A_i is nonempty (it contains at least $u_{\{i\}}$), and $v = v_\mu$ as required.

It remains to verify that condition (ii) implies monotonicity and submodularity of v, and that the converse fails for $n \geq 3$. For monotonicity, note that $b(S) = \sum_{T \supseteq S} c(T)$, so $S \subseteq S'$ implies $b(S) \geq b(S')$, i.e., $v(S) \leq v(S')$. For submodularity, let $S_1, S_2 \subseteq [n]$ and write

$$f(T) = \mathbf{1}_{T \supseteq S_1} + \mathbf{1}_{T \supseteq S_2} - \mathbf{1}_{T \supseteq S_1 \cup S_2} - \mathbf{1}_{T \supseteq S_1 \cap S_2}.$$

We verify that $f(T) \in \{0, -1\}$ for every T by exhausting the four cases determined by the containment relations between T and the pair S_1, S_2. *Case 1*: $T \supseteq S_1 \cup S_2$. Then $T \supseteq S_1$, $T \supseteq S_2$, $T \supseteq S_1 \cup S_2$, and $T \supseteq S_1 \cap S_2$, so all four indicators equal 1 and $f(T) = 1+1-1-1 = 0$. *Case 2*: $T \supseteq S_1$ but $T \not\supseteq S_2$ (and symmetrically $T \supseteq S_2$ but $T \not\supseteq S_1$). Then $T \supseteq S_1 \cap S_2$ (since $S_1 \cap S_2 \subseteq S_1 \subseteq T$) but $T \not\supseteq S_1 \cup S_2$, giving $f(T) = 1+0-0-1 = 0$. *Case 3*: $T \not\supseteq S_1 \cap S_2$. Then $T \not\supseteq S_1$ and $T \not\supseteq S_2$ and $T \not\supseteq S_1 \cup S_2$, so all four indicators are 0 and $f(T) = 0$. *Case 4*: $T \supseteq S_1 \cap S_2$ but $T \not\supseteq S_1$ and $T \not\supseteq S_2$. Then $T \not\supseteq S_1 \cup S_2$, so $f(T) = 0 + 0 - 0 - 1 = -1$. In particular, $f(T) = -1$ occurs precisely in Case 4, where T covers the common part of S_1 and S_2 but fails to extend to either one individually. Since $c(T) \geq 0$ and $f(T) \leq 0$, the sum $b(S_1) + b(S_2) - b(S_1 \cup S_2) - b(S_1 \cap S_2) = \sum_T c(T)\, f(T) \leq 0$, which is supermodularity of b and hence submodularity of v.

For the strictness claim, consider $n = 3$ with $v(\{i\}) = 0.4$ for all i, $v(\{i, j\}) = 0.7$ for all pairs, and $v(\{1, 2, 3\}) = 0.8$. One verifies directly that v is normalized, monotone, and submodular. However,

$$c(\emptyset) = b(\emptyset) - \sum_i b(\{i\}) + \sum_{i<j} b(\{i, j\}) - b(\{1, 2, 3\}) = 1 - 3(0.6) + 3(0.3) - 0.2 = -0.1 < 0,$$

so condition (ii) fails, and v is not an axiom game. □

Remark 8.18 (Connection to Belief Functions)**.** Since axiom games are contained in the class of polymatroids (Edmonds, 1970) (i.e., normalized, monotone, submodular set functions) but do not exhaust it, the characterization shows that the intersection structure of the model-class map imposes constraints beyond submodularity. In fact, the nonnegative Möbius condition (ii) is equivalent to requiring that $b = 1 - v$ be *completely monotone* in the sense that all iterated discrete differences are nonnegative, or equivalently that b be a *belief function* in the sense of Dempster–Shafer theory (Shafer, 1976). This places axiom games in a well-studied class of set functions that arises independently in decision theory and artificial intelligence, where the Möbius coefficients are interpreted as *basic probability assignments* over focal elements. In the special case $n \leq 2$, condition (ii) reduces to monotonicity and submodularity, so the distinction matters only for three or more axioms.

Remark 8.19 (Realization Is Not Unique; Domain Size)**.** The sufficiency proof constructs an *explicit* witness realization on a domain of size 2^n (one atom per subset of the axiom index set). This realization is far from unique. Indeed, two triples $(\mathcal{U}, \Sigma, \mu)$ and $(\mathcal{U}', \Sigma', \mu')$ can induce the same axiom game whenever the induced Möbius coefficients agree, and many different domains and measures yield the same coefficients. The construction above is exponential in the number of axioms, and whether every axiom game admits a realization on a domain of polynomial size in n is an open question.

Remark 8.20 (Standing Assumption on Positive Measure)**.** The standing assumption of this section requires $\mu(A) > 0$ for every $A \in \Sigma$. Theorem 8.17 is stated without this restriction;

functions with $v(\{i\}) = 1$ are included provided condition (ii) holds. The construction in the sufficiency proof still produces a valid axiom system in this case, since each atom $u_{\{i\}}$ belongs to A_i, making A_i nonempty as a set, but the resulting measure may assign $\mu(A_i) = 0$, which lies outside the standing assumption. In practice, this boundary case is benign. An axiom with $\mu(A_i) = 0$ excludes a full-measure subset of the universe on its own, so $v_\mu(\{A_i\}) = 1$ and the Shapley formula continues to apply; the axiom simply carries a large individual weight reflecting its severe restrictiveness.

Remark 8.21 (What the Characterization Excludes)**.** The forward direction of Theorem 8.17 has a noteworthy consequence, independent of the Möbius condition. A cooperative game in which two axioms *amplify* each other's collapse (i.e., the marginal contribution of one axiom *increases* when the other is already present) would require superadditive or supermodular interactions. No axiom system on any domain, under any measure, can exhibit such synergistic interactions. In fact, the intersection structure of the model-class map forces all pairwise interactions to be nonpositive (Theorem 8.8), so diminishing returns is not an empirical regularity but a mathematical necessity. This exclusion is relevant in mechanism design and social choice, where one might informally speak of axioms "reinforcing" one another; Proposition 8.2(iv) shows that such reinforcement cannot occur at the level of marginal contributions to model-class collapse.

8.7 Axiom Entropy

The axiom Shapley value assigns to each axiom a nonnegative weight summing to $v_\mu(\Sigma)$. When $v_\mu(\Sigma) > 0$, we may normalize these weights to a probability distribution and measure how evenly the explanatory burden is spread.

Definition 8.22 (Axiom Entropy)**.** Let (Σ, v_μ) be an axiom game with $v_\mu(\Sigma) > 0$. The *axiom Shapley measure* is the probability distribution $\pi\colon \Sigma \to [0, 1]$ defined by

$$\pi_A = \frac{\phi_A(\Sigma, \mu)}{v_\mu(\Sigma)}, \qquad A \in \Sigma.$$

The *axiom entropy* of (Σ, μ) is

$$H(\Sigma, \mu) = -\sum_{A \in \Sigma} \pi_A \ln \pi_A,$$

where ln denotes the natural logarithm, with the convention $0 \ln 0 = 0$.

The axiom entropy measures how evenly the axioms share responsibility for collapsing the model class. When H is small, a few axioms dominate; when H is large, the collapse is distributed democratically.

Theorem 8.23 (Axiom Entropy Bounds)**.** *Let (Σ, v_μ) be an axiom game with $v_\mu(\Sigma) > 0$ and $|\Sigma| = n$.*

(i) Minimum. *$H(\Sigma, \mu) = 0$ if and only if all axioms except one are null, i.e., there exists $A \in \Sigma$ such that $\phi_B(\Sigma, \mu) = 0$ for every $B \neq A$. Equivalently, some single axiom carries the entire collapse burden.*

(ii) MAXIMUM. *$H(\Sigma, \mu) \leq \ln n$, with equality if and only if $\phi_A(\Sigma, \mu) = v_\mu(\Sigma)/n$ for every $A \in \Sigma$, i.e., every axiom bears equal responsibility.*

(iii) NULL-AXIOM INVARIANCE. *$H(\Sigma, \mu) = H(\Sigma', \mu)$, where $\Sigma' \subseteq \Sigma$ is the sub-system obtained by removing all null axioms (those with $\phi_A = 0$). In particular, axiom entropy is an invariant of the irredundant core of Σ.*

(iv) IMPOSSIBILITY INTERPRETATION. *The axiom Shapley measure π is a probability distribution on Σ for every game with $v_\mu(\Sigma) > 0$, by construction. In the impossibility regime ($v_\mu(\Sigma) = 1$), the efficiency condition reads $\sum_A \phi_A = 1$, so $\pi_A = \phi_A$ without any normalization; each axiom's Shapley value is directly its share of the blame. The entropy $H(\Sigma, \mu)$ then measures the dispersal of impossibility blame across the axiom set.*

Proof. For (i), note that $H(\Sigma, \mu) = -\sum_A \pi_A \ln \pi_A \geq 0$ always, since $-\pi_A \ln \pi_A \geq 0$ for $\pi_A \in [0, 1]$. Moreover, $H = 0$ if and only if every term vanishes, which occurs if and only if each $\pi_A \in \{0, 1\}$. Since the π_A sum to 1, exactly one equals 1 and the rest equal 0, i.e., π is a point mass at some A. This means $\phi_A = v_\mu(\Sigma)$ and $\phi_B = 0$ for $B \neq A$. By Theorem 8.4(a), $\phi_B = 0$ if and only if B is null, i.e., $v_\mu(S \cup \{B\}) = v_\mu(S)$ for every $S \subseteq \Sigma \setminus \{B\}$.

For (ii), the function $t \mapsto -t \ln t$ is strictly concave on $[0, 1]$. By the concavity of the entropy function (or equivalently by Jensen's inequality applied to the convex function $t \mapsto t \ln t$), $H(\Sigma, \mu) \leq \ln n$, with equality if and only if $\pi_A = 1/n$ for all $A \in \Sigma$.

For (iii), null axioms have $\phi_A = 0$, hence $\pi_A = 0$, and contribute $0 \ln 0 = 0$ to the entropy sum. Removing them does not change H. Since the remaining axioms have the same normalized weights (the normalization constant $v_\mu(\Sigma)$ is unchanged), $H(\Sigma, \mu) = H(\Sigma', \mu)$.

For (iv), when $v_\mu(\Sigma) = 1$, the efficiency condition is $\sum_A \phi_A = 1$, so $\pi_A = \phi_A / v_\mu(\Sigma) = \phi_A$ without any normalization; the interpretation as a blame distribution is exact. □

Example 8.24 (Entropy in the Symmetric Case)**.** Consider the canonical impossible system of dimension k. Let $\mathcal{U} = \{1, 2, \ldots, k\}$ and $A_i = \mathcal{U} \setminus \{i\}$ for each i. Under the uniform measure $\mu(\{i\}) = 1/k$, every axiom has $v_\mu(\{A_i\}) = 1 - (k-1)/k = 1/k$, and by symmetry $\phi_{A_i} = 1/k$ for all i. Thus, $\pi_{A_i} = 1/k$ and $H = \ln k = \ln(\dim(\Sigma))$. The entropy equals the natural logarithm of the impossibility dimension. The more axioms are required to produce the contradiction, the more evenly the blame is spread, and the higher the entropy.

Remark 8.25 (Entropy and Impossibility Depth)**.** The example above suggests a general principle. Systems of higher impossibility dimension tend to have higher maximum entropy, since the blame cannot concentrate on fewer axioms than are needed for the core contradiction. The precise statement concerns the support of π, not a lower bound on H itself. For any impossible Σ with $\dim(\Sigma) = k$ and any measure μ, every axiom in an impossibility core has positive Shapley value (it is independent in the core, hence makes a strictly positive marginal contribution to the full coalition), so the effective support of π has size at least k. However, $H(\Sigma, \mu) \geq \ln k$ does not hold in general. The miniature Arrow-type impossibility of Section 12.2 has $\dim(\Sigma) = 3$ (each of its two cores has size 3) and axiom entropy $H \approx 1.31$ nats, close to the theoretical maximum $\ln 4 \approx 1.39$ nats. What is true universally is that at least k atoms carry positive mass, so the maximum possible entropy given the support constraint is $\ln(\text{support size}) \geq \ln k$; the actual entropy may be anywhere in $[0, \ln n]$.

9 The Rescue Theorem

When a system is impossible ($\mathcal{M}(\Sigma) = \emptyset$), the analyst faces a fundamental question: how can possibility be restored? We now classify rescue operations, identify the optimal single-axiom drop, and characterize the weakening condition for restoring possibility.

Definition 9.1 (Rescue Operations)**.** Let Σ be an impossible axiom system (that is, $\mathcal{M}(\Sigma) = \emptyset$). A *rescue* of Σ is a modified axiom system Σ' on a domain $\mathcal{U}'$ with $\mathcal{M}_{\mathcal{U}'}(\Sigma') \neq \emptyset$. We distinguish three canonical types.

(i) *Axiom drop*: $\mathcal{U}' = \mathcal{U}$ and $\Sigma' = \Sigma \setminus \{A\}$ for some $A \in \Sigma$.

(ii) *Axiom weakening*: $\mathcal{U}' = \mathcal{U}$ and $\Sigma' = (\Sigma \setminus \{A\}) \cup \{A'\}$ for some $A \in \Sigma$ and $A' \supsetneq A$ (a strictly weaker version of A).

(iii) *Domain restriction*: $\mathcal{U}' \subsetneq \mathcal{U}$ and $\Sigma' = \{A \cap \mathcal{U}' \mid A \in \Sigma\}$ (the axioms restricted to a subdomain).

A rescue is *minimal* if no strictly smaller modification (in terms of μ-measure surrendered) achieves $\mathcal{M}_{\mathcal{U}'}(\Sigma') \neq \emptyset$.

Definition 9.2 (Impossibility Depth)**.** Let Σ be an impossible axiom system ($\mathcal{M}(\Sigma) = \emptyset$). The *impossibility depth* of Σ is

$$d(\Sigma) = \min\{|T| \mid T \subseteq \Sigma,\ \mathcal{M}(\Sigma \setminus T) \neq \emptyset\},$$

i.e., the minimum size of a set $T \subseteq \Sigma$ of axioms whose removal restores possibility. Note that $d(\Sigma) \geq 1$ always, and $d(\Sigma) = 1$ if and only if $\mathcal{M}(\Sigma \setminus \{A\}) \neq \emptyset$ for some $A \in \Sigma$.

Remark 9.3 (Rescue Decomposition)**.** Any rescue of an impossible system modifies either the domain, the axioms, or both. Restricting the domain gives type (iii) of Definition 9.1. Modifying axioms on the same domain amounts to replacing each axiom A by a weakly larger set $A' \supseteq A$ (weakening, type (ii)) or removing it entirely (type (i), the limit of weakening to $A' = \mathcal{U}$). In this sense, every rescue decomposes, up to logical equivalence, into a combination of these three types.

Theorem 9.4 (Rescue)**.** *Let Σ be an impossible axiom system with $\mathcal{M}(\Sigma) = \emptyset$.*

(i) Optimal single-axiom drop. *Among all single-axiom drops, the one that maximally reduces the residual collapse $v_\mu(\Sigma \setminus \{A\})$ is the axiom A^* maximizing the full-coalition marginal $\Delta_A = v_\mu(\Sigma) - v_\mu(\Sigma \setminus \{A\})$. In general, A^* need not coincide with the Shapley-maximal axiom* $\arg\max_A \phi_A(\Sigma, \mu)$.

(ii) Rescue by weakening. *Suppose $d(\Sigma) = 1$, and assume additionally that every rescuing axiom A (i.e., every A with $\mathcal{M}(\Sigma \setminus \{A\}) \neq \emptyset$) satisfies $\mu(\mathcal{M}(\Sigma \setminus \{A\})) > 0$. (This condition is automatic whenever μ has full support, i.e., $\mu(\{u\}) > 0$ for every $u \in \mathcal{U}$.) Then the maximizer A^* from (i) has $\Delta_{A^*} > 0$, whereas any non-rescuing A has $\mathcal{M}(\Sigma \setminus \{A\}) = \emptyset$ and hence $\Delta_A = 0$; the maximizer A^* is therefore among the rescuing axioms and satisfies $\mathcal{M}(\Sigma \setminus \{A^*\}) \neq \emptyset$. An axiom weakening $A' \supseteq A^*$*

restores possibility if and only if $A' \cap \mathcal{M}(\Sigma \setminus \{A^*\}) \neq \emptyset$. *The* canonical weakening *is* $A^\dagger = A^* \cup \mathcal{M}(\Sigma \setminus \{A^*\})$, *the weakest replacement that preserves the full model class of the remaining axioms, i.e.,* $\mathcal{M}((\Sigma \setminus \{A^*\}) \cup \{A^\dagger\}) = \mathcal{M}(\Sigma \setminus \{A^*\})$. *The* minimum-cost weakening *adds a single element* $f \in \mathcal{M}(\Sigma \setminus \{A^*\}) \setminus A^*$ *of minimum* μ*-weight to* A^*, *yielding* $A' = A^* \cup \{f\}$ *at cost* $\mu(\{f\})$. *When* $d(\Sigma) > 1$, *no single-axiom weakening suffices, and one must weaken or drop multiple axioms.*

Proof. For (i), among axiom-drop rescues, the residual collapse after dropping A is $v_\mu(\Sigma\setminus\{A\})$. We seek to minimize $v_\mu(\Sigma \setminus \{A\})$, equivalently to maximize $\Delta_A = v_\mu(\Sigma) - v_\mu(\Sigma \setminus \{A\})$, the full-coalition marginal contribution of A. The claim that A^* need not coincide with $\arg\max_A \phi_A$ is established by Example 9.10.

For (ii), suppose $d(\Sigma) = 1$, so that some axiom A satisfies $\mathcal{M}(\Sigma \setminus \{A\}) \neq \emptyset$. For any such rescuing axiom, $v_\mu(\Sigma \setminus \{A\}) = 1 - \mu(\mathcal{M}(\Sigma \setminus \{A\})) < 1 = v_\mu(\Sigma)$ (using the hypothesis $\mu(\mathcal{M}(\Sigma \setminus \{A\})) > 0$), so $\Delta_A > 0$. For any non-rescuing axiom A (one with $\mathcal{M}(\Sigma \setminus \{A\}) = \emptyset$), we have $v_\mu(\Sigma \setminus \{A\}) = 1 = v_\mu(\Sigma)$, so $\Delta_A = 0$. Hence, the maximizer A^* of Δ_A is necessarily among the rescuing axioms. An axiom weakening $A' \supseteq A^*$ restores possibility if and only if $\mathcal{M}((\Sigma \setminus \{A^*\}) \cup \{A'\}) = \mathcal{M}(\Sigma \setminus \{A^*\}) \cap A' \neq \emptyset$, which holds precisely when A' contains at least one element of $\mathcal{M}(\Sigma \setminus \{A^*\})$. The canonical weakening $A^\dagger = A^* \cup \mathcal{M}(\Sigma \setminus \{A^*\})$ achieves $\mathcal{M}((\Sigma \setminus \{A^*\}) \cup \{A^\dagger\}) = \mathcal{M}(\Sigma \setminus \{A^*\}) \neq \emptyset$, and it is the weakest (i.e., largest) such replacement because $A^\dagger \supseteq A'$ for every weakening A' that preserves the full surviving model class. The minimum-cost weakening adds a single element $f \in \mathcal{M}(\Sigma \setminus \{A^*\}) \setminus A^*$ of minimum μ-weight, yielding $A' = A^* \cup \{f\}$ at cost $\mu(\{f\})$. When $d(\Sigma) > 1$, every axiom drop leaves at least one impossibility core intact (by definition of d), so no single weakening or drop suffices. □

Remark 9.5 (Shapley Value Versus Full-Coalition Marginal)**.** The Shapley value $\phi_A(\Sigma, \mu)$ is the *average* marginal contribution of A, averaging over all orderings of Σ. The full-coalition marginal $\Delta_A = v_\mu(\Sigma) - v_\mu(\Sigma \setminus \{A\})$ is the marginal contribution of A when it is the *last* axiom added. By submodularity of v_μ (Proposition 8.2(iv)), the last-added marginal is the *smallest* among all marginals of A, so $\Delta_A \leq \phi_A$. The two quantities can rank axioms differently. The axiom with the largest average contribution (ϕ) need not be the one whose removal most reduces the full-coalition collapse (Δ). Example 9.10 below demonstrates this divergence explicitly.

Corollary 9.6 (Rescue Cost)**.** *Let* Σ *be an impossible axiom system. The minimum residual collapse achievable by a single axiom drop equals*

$$\min_{A \in \Sigma} v_\mu(\Sigma \setminus \{A\}) = v_\mu(\Sigma) - \max_{A \in \Sigma} \Delta_A.$$

Proof. This follows from Theorem 9.4(i). Among all single-axiom drops, the one minimizing the residual collapse $v_\mu(\Sigma \setminus \{A\})$ is the one maximizing $\Delta_A = v_\mu(\Sigma) - v_\mu(\Sigma \setminus \{A\})$. □

In the running example (Example 3.1), the minimum residual collapse equals $1-1/6 = 5/6$, achieved by dropping either A_1 or A_2.

Remark 9.7 (Domain Restriction Cannot Rescue Impossibility)**.** One might expect domain restriction (type (iii) of Definition 9.1) to provide an alternative rescue strategy. However, in

the extensional setup of this paper, restricting to a subdomain $\mathcal{U}' \subseteq \mathcal{U}$ cannot restore possibility. Indeed, the restricted model class is

$$\mathcal{M}_{\mathcal{U}'}(\Sigma') = \bigcap_{A \in \Sigma} (A \cap \mathcal{U}') = \mathcal{U}' \cap \bigcap_{A \in \Sigma} A = \mathcal{U}' \cap \mathcal{M}(\Sigma),$$

which is empty whenever $\mathcal{M}(\Sigma) = \emptyset$. The difficulty is that our axioms are extensional objects (subsets), not intensional properties. In the standard formulations of Arrow's or Nash's theorem, restricting the domain changes the *meaning* of each axiom (e.g., the set of social welfare functions satisfying IIA changes when the set of alternatives changes), so the restricted axioms are genuinely different subsets of the restricted domain. This intensional phenomenon lies outside the scope of the present framework.

The impossibility depth $d(\Sigma)$ has a clean combinatorial characterization in terms of the family of impossibility cores, which we now develop. Let $\mathcal{C}(\Sigma)$ denote the family of all impossibility cores of Σ. A *transversal* (or hitting set) of $\mathcal{C}(\Sigma)$ is a set $T \subseteq \Sigma$ that intersects every core: $T \cap C \neq \emptyset$ for all $C \in \mathcal{C}(\Sigma)$. The *transversal number* $\tau(\mathcal{C}(\Sigma))$ is the minimum size of a transversal.

Theorem 9.8 (Impossibility Depth as Transversal Number)**.** *For any impossible axiom system Σ,*

$$d(\Sigma) = \tau(\mathcal{C}(\Sigma)).$$

That is, the minimum number of axioms whose removal restores possibility equals the minimum size of a hitting set of the family of impossibility cores.

Proof. We show that a set $T \subseteq \Sigma$ restores possibility (i.e., $\mathcal{M}(\Sigma \setminus T) \neq \emptyset$) if and only if T is a transversal of $\mathcal{C}(\Sigma)$.

If T is a transversal, then every core $C \in \mathcal{C}(\Sigma)$ satisfies $C \not\subseteq \Sigma \setminus T$, so $\Sigma \setminus T$ contains no impossibility core. By Theorem 7.1, a system is impossible if and only if it contains a finite impossibility core; since $\Sigma \setminus T$ contains none, $\mathcal{M}(\Sigma \setminus T) \neq \emptyset$.

Conversely, if T is not a transversal, there exists a core $C \in \mathcal{C}(\Sigma)$ with $C \cap T = \emptyset$, so $C \subseteq \Sigma \setminus T$. Since C is impossible, $\mathcal{M}(\Sigma \setminus T) \subseteq \mathcal{M}(C) = \emptyset$.

Taking the minimum over all T, the minimum number of drops restoring possibility is exactly $\tau(\mathcal{C}(\Sigma))$. □

Remark 9.9 (Depth, Dimension, and Duality)**.** Theorem 9.8 closes a triangle of dualities. The impossibility dimension $\dim(\Sigma) = k$ is the minimum core size (the smallest hyperedge of the compatibility hypergraph). The transversal number $\tau(\mathcal{C}(\Sigma))$ is the minimum number of vertices needed to hit every core (a covering problem). The Shapley value ϕ_A measures A's average marginal contribution to the collapse, which is a different quantity from either. These three invariants are genuinely distinct; the following example demonstrates their divergence.

Example 9.10 (Greedy-Shapley Rescue is Multi-Step Suboptimal)**.** Let $\mathcal{U} = \{1, 2, 3, 4, 5, 6\}$ with uniform measure $\mu(\{i\}) = 1/6$, and let

$$A_1 = \{1\}, \quad A_2 = \{1, 2, 3\}, \quad A_3 = \{4, 5, 6\}, \quad A_4 = \{1, 4\}.$$

The system $\Sigma = \{A_1, A_2, A_3, A_4\}$ is impossible since $\mathcal{M}(\Sigma) = A_1 \cap A_2 \cap A_3 \cap A_4 = \emptyset$. Its impossibility cores are $\{A_1, A_3\}$ (since $A_1 \cap A_3 = \emptyset$) and $\{A_2, A_3\}$ (since $A_2 \cap A_3 = \emptyset$); no other pair is impossible. By Theorem 9.8, $d(\Sigma) = \tau(\mathcal{C}(\Sigma)) = 1$. Indeed, the single axiom A_3 hits both cores (since $A_3 \in \{A_1, A_3\}$ and $A_3 \in \{A_2, A_3\}$), so dropping A_3 alone restores possibility. Indeed, $\mathcal{M}(\{A_1, A_2, A_4\}) = A_1 \cap A_2 \cap A_4 = \{1\} \neq \emptyset$.

We compute the axiom Shapley values from the sub-system collapse values. The singleton collapses are $v(\{A_1\}) = 5/6$, $v(\{A_2\}) = 1/2$, $v(\{A_3\}) = 1/2$, $v(\{A_4\}) = 2/3$. The pair collapses are $v(\{A_1, A_2\}) = 5/6$, $v(\{A_1, A_3\}) = 1$, $v(\{A_1, A_4\}) = 5/6$, $v(\{A_2, A_3\}) = 1$, $v(\{A_2, A_4\}) = 5/6$, $v(\{A_3, A_4\}) = 5/6$. The triple collapses are $v(\{A_1, A_2, A_3\}) = 1$, $v(\{A_1, A_2, A_4\}) = 1$, $v(\{A_1, A_3, A_4\}) = 1$, $v(\{A_2, A_3, A_4\}) = 1$. Applying the Shapley weights $w_0 = w_3 = 1/4$ and $w_1 = w_2 = 1/12$ yields

$$\phi_{A_1} = \tfrac{11}{36}, \quad \phi_{A_3} = \tfrac{10}{36}, \quad \phi_{A_4} = \tfrac{8}{36}, \quad \phi_{A_2} = \tfrac{7}{36},$$

with sum $1 = v_\mu(\Sigma)$. The full-coalition marginals are

$$\Delta_{A_3} = v_\mu(\Sigma) - v_\mu(\{A_1, A_2, A_4\}) = 1 - 5/6 = 1/6 \quad \text{and} \quad \Delta_{A_1} = 1 - v_\mu(\{A_2, A_3, A_4\}) = 0,$$

since $\{A_2, A_3\}$ is a core contained in $\{A_2, A_3, A_4\}$. Thus, A_1 has the highest Shapley value but the *lowest* full-coalition marginal; dropping it does not even restore possibility.

The Shapley-greedy rescue strategy drops A_1 first (highest ϕ). The residual system $\{A_2, A_3, A_4\}$ remains impossible, since it contains the core $\{A_2, A_3\}$. A second greedy drop removes one of A_2 or A_3, and possibly a third is needed. The optimal strategy drops A_3 in one step.

This example illustrates the distinction of Remark 9.5. The Shapley value ranks axioms by average contribution across all sub-coalitions, not by their structural role in the core family. The axiom A_1 has a high Shapley value because $v_\mu(\{A_1\}) = 5/6$ is large (it alone eliminates most of the universe), but it appears in only one core. The axiom A_3, with the second-highest ϕ, is the unique minimum transversal and the correct first drop.

Remark 9.11 (Finding the Minimum Transversal)**.** Computing $\tau(\mathcal{C}(\Sigma))$ is equivalent to the minimum hitting set problem, which is NP-hard in general (it contains vertex cover as a special case). However, for axiom systems arising in practice, the family $\mathcal{C}(\Sigma)$ is small and the transversal can be found by inspection. A useful algorithmic observation is that an axiom A appearing in every impossibility core must belong to every transversal and can therefore be included greedily, leaving a system with strictly fewer cores. Iterating this rule when it applies yields a partial reduction of the transversal problem, though not in general an optimal solution.

Example 9.12 (Running Example: Rescue and Transversal)**.** In Example 3.1, the only impossibility core is $\{A_1, A_2\}$. By Theorem 9.8, $d(\Sigma) = \tau(\mathcal{C}(\Sigma)) = 1$, since dropping either A_1 or A_2 hits the sole core and restores possibility. Indeed, $\mathcal{M}(\Sigma \setminus \{A_1\}) = A_2 \cap A_3 \cap A_4 = \{4\}$ and $\mathcal{M}(\Sigma \setminus \{A_2\}) = A_1 \cap A_3 \cap A_4 = \{1\}$. By Theorem 9.4(i), the optimal single-axiom drop maximizes the full-coalition marginal $\Delta_A = v_\mu(\Sigma) - v_\mu(\Sigma \setminus \{A\})$. We have $\Delta_{A_1} = \Delta_{A_2} = 1/6$, so both drops are equally effective. Note that $\Delta_{A_3} = \Delta_{A_4} = 0$, since dropping A_3 or A_4 does not restore possibility (the core $\{A_1, A_2\}$ remains). This illustrates the distinction between the Shapley ranking ($\phi_{A_3} = \phi_{A_4} = 1/6 > 0$) and the full-coalition marginal ($\Delta_{A_3} = \Delta_{A_4} = 0$).

10 The Impossibility Dimension

So far, impossibility has been treated as a binary phenomenon. A system either has a model or it does not. We now introduce a finer invariant that measures how many axioms are needed, at minimum, to produce a contradiction, and we prove a fundamental theorem connecting this invariant to the minimum size of the domain needed to certify independence.

Definition 10.1 (Impossibility Dimension)**.** Let Σ be an impossible axiom system (that is, $\mathcal{M}(\Sigma) = \emptyset$). The *impossibility dimension* of Σ is

$$\dim(\Sigma) = \min\bigl\{|\Sigma_0| \;\bigm|\; \Sigma_0 \subseteq \Sigma,\ \mathcal{M}(\Sigma_0) = \emptyset\bigr\},$$

the minimum cardinality of an impossibility core of Σ. Note that any impossible sub-system of minimum cardinality is automatically a core in the sense of Theorem 7.1, since if some axiom could be removed while preserving impossibility, the resulting sub-system would be impossible with strictly smaller cardinality.

The impossibility dimension is a well-defined positive integer by Theorem 7.1, since some finite sub-system witnesses the impossibility and the minimum is attained.

Remark 10.2 (Notation)**.** Throughout this paper, $\dim(\Sigma)$ always denotes the *impossibility dimension* of Definition 10.1, i.e., the minimum core size. This is distinct from the topological dimension $\dim(K(\Sigma))$ of the feasibility complex introduced in Section 11, which equals the maximum face size minus one. The two invariants are related but not equal. For instance, when Σ is a core of size n, we have $\dim(\Sigma) = n$ but $\dim(K(\Sigma)) = n - 2$. Context will always make clear which quantity is meant.

Example 10.3 (Running Example: Impossibility Dimension)**.** In Example 3.1, the smallest impossible sub-system is $\{A_1, A_2\}$ (since $A_1 \cap A_2 = \emptyset$), so $\dim(\Sigma) = 2$. By Proposition 10.5(i), this is the minimum possible value. By Theorem 10.6(i), the two independence witnesses for $\{A_1, A_2\}$ must be distinct. Indeed, $w_1 \in \mathcal{M}(\{A_2\}) \setminus A_1 = \{4,5,6\} \setminus \{1,2,3\} = \{4,5,6\}$ and $w_2 \in \mathcal{M}(\{A_1\}) \setminus A_2 = \{1,2,3\} \setminus \{4,5,6\} = \{1,2,3\}$, so any such pair is indeed distinct, confirming the theorem.

Example 10.4 (Dimension of the Miniature and Classical Arrow Systems)**.** The miniature Arrow-type system of Section 12.2 has $\dim(\Sigma) = 3$, with two overlapping impossibility cores of size 3. Classical Arrow's system $\Sigma = \{A_{\mathrm{UD}}, A_{\mathrm{P}}, A_{\mathrm{IIA}}, A_{\mathrm{ND}}\}$ is generally understood to satisfy $\dim(\Sigma) = 4$, so the full system is an impossibility core, since dropping any one axiom restores possibility (as verified by the four independence witnesses standard in the social-choice literature). This is the maximal possible dimension for a four-axiom system and reflects the fact that the classical Arrow contradiction requires all four axioms simultaneously. The status of the classical Arrow claim within the framework of this paper is heuristic; see Section 13.1.

Proposition 10.5 (Basic Properties of the Impossibility Dimension)**.** *Let Σ be impossible.*

(i) Lower bound. $\dim(\Sigma) \geq 2$*, since $\mathcal{M}(\{A\}) = A \neq \emptyset$ for any axiom A (recall that axioms are nonempty by Definition 2.1).*

(ii) Direct contradiction. $\dim(\Sigma) = 2$ *if and only if* Σ *contains two contradictory axioms* A *and* B *with* $A \cap B = \emptyset$.

(iii) Core independence. *Every impossibility core is an independent system: each axiom in the core is essential, and removing any one of them restores possibility.*

(iv) Upper bound. $\dim(\Sigma) \leq |\Sigma|$, *with equality if and only if* Σ *itself is an impossibility core.*

Proof. For (i), $\mathcal{M}(\{A\}) = A$; since $A \neq \emptyset$ by definition, no single axiom is impossible, so $\dim(\Sigma) \geq 2$.

For (ii), $\dim(\Sigma) = 2$ if and only if some two-axiom sub-system $\{A, B\}$ is impossible, i.e., $A \cap B = \mathcal{M}(\{A, B\}) = \emptyset$.

For (iii), let Σ_0 be an impossibility core with $|\Sigma_0| = \dim(\Sigma)$. If some $A \in \Sigma_0$ were redundant in Σ_0, then $\mathcal{M}(\Sigma_0 \setminus \{A\}) = \mathcal{M}(\Sigma_0) = \emptyset$, contradicting minimality of $|\Sigma_0|$. Hence, every $A \in \Sigma_0$ is independent, i.e., $\mathcal{M}(\Sigma_0 \setminus \{A\}) \neq \emptyset$.

For (iv), this is immediate from the definition and part (iii). □

The central theorem of this section connects the impossibility dimension to the minimum size of a domain needed to certify independence of an impossibility core.

Theorem 10.6 (Impossibility Dimension and Witness Domain)**.** *Let* $\Sigma_0 = \{A_1, A_2, \ldots, A_k\}$ *be an independent impossible system (*$\mathcal{M}(\Sigma_0) = \emptyset$ *and* $\mathcal{M}(\Sigma_0 \setminus \{A_i\}) \neq \emptyset$ *for each* i*).*

(i) Witness domain lower bound. *The domain* $\mathcal{U}$ *must contain at least* k *objects. Specifically, the* k *independence witnesses* $w_1, w_2, \ldots, w_k$ *(where* $w_i \in \mathcal{M}(\Sigma_0 \setminus \{A_i\}) \setminus A_i$*) are pairwise distinct.*

(ii) Witness domain upper bound. *Conversely, any domain of size at least* k *can support an independent impossible system of dimension* k*.*

(iii) The Dimension-Witness Duality. *For any impossible axiom system* Σ *on domain* $\mathcal{U}$*, we have* $\dim(\Sigma) \leq |\mathcal{U}|$*. Conversely, for every integer* k *with* $2 \leq k \leq |\mathcal{U}|$*, there exists an impossible axiom system on some subdomain of* $\mathcal{U}$ *of size exactly* k *whose impossibility dimension equals* k*.*

Proof. For (i), let $w_i \in \mathcal{M}(\Sigma_0 \setminus \{A_i\}) \setminus A_i$ be the independence witness for A_i. We claim $w_i \neq w_j$ for $i \neq j$. Indeed, since $w_i \in \mathcal{M}(\Sigma_0 \setminus \{A_i\})$ and $j \neq i$, we have $w_i \in A_j$ (because $A_j \in \Sigma_0 \setminus \{A_i\}$ and w_i satisfies every axiom in $\Sigma_0 \setminus \{A_i\}$). On the other hand, $w_j \notin A_j$ by definition of w_j. Hence, $w_i \in A_j$ but $w_j \notin A_j$, so $w_i \neq w_j$.

For (ii), let $\mathcal{U} = \{u_1, u_2, \ldots, u_k\}$ and define $A_i = \mathcal{U} \setminus \{u_i\}$ for each $i = 1, 2, \ldots, k$. Then $\mathcal{M}(\{A_1, A_2, \ldots, A_k\}) = \bigcap_i (\mathcal{U} \setminus \{u_i\}) = \emptyset$, and for each i we have

$$\mathcal{M}(\{A_1, A_2, \ldots, A_k\} \setminus \{A_i\}) = \mathcal{U} \setminus \bigcup_{j \neq i} \{u_j\} = \{u_i\} \neq \emptyset.$$

So the family $\{A_1, A_2, \ldots, A_k\}$ is an independent impossible system of dimension k on a domain of size k.

For (iii), the upper bound $\dim(\Sigma) \leq |\mathcal{U}|$ follows by applying (i) to any impossibility core of Σ. The converse construction follows from (ii) applied to any k-element subdomain of $\mathcal{U}$. □

Remark 10.7 (The Witness Distinctness Argument)**.** The key step in Theorem 10.6(i) is that independence witnesses are pairwise distinct. The intuition is that each witness w_i is "approved" by all axioms except A_i; in particular, w_i is approved by A_j for $j \neq i$. Meanwhile, w_j is not approved by A_j. So w_i and w_j differ on their attitude toward A_j. The former is approved, the latter is not. They cannot be the same object.

This argument has a striking consequence. The minimum domain size needed for an independent impossible system of dimension k is exactly k. Neither more nor fewer objects are required. In this sense, the impossibility dimension is not merely a property of the axiom system but an intrinsic measure of the minimum "social complexity" needed to exhibit the contradiction.

Corollary 10.8 (Arrow's Theorem Revisited)**.** *Arrow's impossibility system has dimension* 4*. By Theorem 10.6(i), it requires at least* 4 *distinct objects in the domain (i.e.,* 4 *distinct social welfare functions serving as independence witnesses) to exhibit the four independence witnesses. Indeed, in the standard formulation with at least* 3 *alternatives and* 2 *individuals, the four witnesses exist and are distinct.*

10.1 The Impossible System as a Hypergraph

The impossibility dimension has a natural combinatorial interpretation. Associate to Σ the *compatibility hypergraph* $H(\Sigma)$. Its vertex set is Σ, and a subset $S \subseteq \Sigma$ is a hyperedge if and only if $\mathcal{M}(S) = \emptyset$ (the axioms in S are jointly incompatible). The impossibility dimension $\dim(\Sigma)$ is then the minimum size of a hyperedge in $H(\Sigma)$.

Proposition 10.9 (Hypergraph Monotonicity)**.** *The compatibility hypergraph $H(\Sigma)$ is an* upset *(upward-closed family): if S is a hyperedge and $S \subseteq T \subseteq \Sigma$, then T is also a hyperedge.*

Proof. If $\mathcal{M}(S) = \emptyset$ and $S \subseteq T$, then $\mathcal{M}(T) \subseteq \mathcal{M}(S) = \emptyset$ by Theorem 4.1(i), so T is a hyperedge. □

The minimal hyperedges of $H(\Sigma)$ are the impossibility cores. Proposition 10.9 says that once a sub-system is impossible, every extension is too. The impossibility dimension is the minimum cardinality of a minimal hyperedge.

11 The Feasibility Complex

The compatibility hypergraph of Σ records which sub-systems are impossible. Its complement, the collection of possible sub-systems, is a simplicial complex whose topology encodes the impossibility structure. We develop this now.

Definition 11.1 (Feasibility Complex)**.** The *feasibility complex* of Σ is

$$K(\Sigma) = \big\{ S \subseteq \Sigma \;\big|\; \mathcal{M}(S) \neq \emptyset \big\},$$

the collection of all possible sub-systems of Σ, ordered by inclusion.

Proposition 11.2 (Simplicial Structure)**.** *The set $K(\Sigma)$ is a simplicial complex, i.e., it is closed under taking subsets. The facets of $K(\Sigma)$ (its maximal faces) are precisely the maximal possible sub-systems of Σ.*

Proof. If $\mathcal{M}(S) \neq \emptyset$ and $T \subseteq S$, then $\mathcal{M}(T) \supseteq \mathcal{M}(S) \neq \emptyset$, so $T \in K(\Sigma)$. The characterization of facets is immediate. A face $S \in K(\Sigma)$ is maximal if and only if no proper extension $S \cup \{A\}$ (with $A \in \Sigma \setminus S$) is in $K(\Sigma)$, i.e., if and only if $\mathcal{M}(S \cup \{A\}) = \emptyset$ for every $A \in \Sigma \setminus S$. □

The simplicial complex $K(\Sigma)$ is a sub-complex of the full $(n-1)$-simplex Δ^{n-1} on vertex set Σ. It is a proper sub-complex exactly when Σ is impossible. The full simplex Δ^{n-1} corresponds to the case where every sub-system is possible, i.e., Σ is in the possibility regime.

Theorem 11.3 (Topology of the Feasibility Complex)**.** *Let Σ be an impossible axiom system.*

(i) Dimension bound. $\dim(K(\Sigma)) \leq n-2$ *(using topological dimension, i.e., maximum face size minus one), with equality if and only if Σ has some possible sub-system of size $n-1$.*

(ii) Core case. *If Σ is an impossibility core ($\dim(\Sigma) = n$, i.e., every proper sub-system is possible), then $K(\Sigma)$ equals the boundary of Δ^{n-1} and is homeomorphic to the $(n-2)$-dimensional sphere S^{n-2}.*

(iii) Euler characteristic. *When Σ is an impossibility core of size n, we compute $\chi(K(\Sigma)) = \sum_{j=0}^{n-2}(-1)^j f_j$, where f_j denotes the number of j-dimensional faces (that is, faces of cardinality $j+1$). The result is*

$$\chi(K(\Sigma)) = 1 + (-1)^n = \chi(S^{n-2}) = \begin{cases} 2 & \text{if } n \text{ is even,} \\ 0 & \text{if } n \text{ is odd.} \end{cases}$$

(iv) General impossible case. *For an arbitrary impossible Σ, the feasibility complex $K(\Sigma)$ is a proper sub-complex of $\partial\Delta^{n-1}$. The precise relationship between the topology of $K(\Sigma)$ and the impossibility cores is given by the Alexander Duality Theorem (Theorem 11.11 below), which identifies $\widetilde{H}_i(K(\Sigma); \daleth) \cong \widetilde{H}_{n-i-3}(K^\vee(\Sigma); \daleth)$, where $K^\vee(\Sigma)$ is the co-core complex whose maximal faces are the complements $\Sigma \setminus C$ of the impossibility cores $C \in \mathcal{C}(\Sigma)$.*

Proof. For (i), since Σ is impossible, $\Sigma \notin K(\Sigma)$, so every face has size at most $n-1$, giving topological dimension at most $n-2$. Equality holds if and only if some $(n-1)$-element sub-system is possible.

For (ii), if Σ is an impossibility core, every proper subset of Σ is possible (by definition of core and Proposition 10.5(iii)), so $K(\Sigma) = \{S \subsetneq \Sigma\} = \partial\Delta^{n-1}$, the boundary of the $(n-1)$-simplex. The boundary of Δ^{n-1} is homeomorphic to S^{n-2}, a simplicial sphere of dimension $n-2$.

For (iii), the complex $K(\Sigma) = \partial\Delta^{n-1}$ has exactly $\binom{n}{k+1}$ faces of dimension k, for each $k = 0, 1, \ldots, n-2$. Equivalently, there are $\binom{n}{s}$ faces of cardinality s for $s = 1, 2, \ldots, n-1$. Writing the Euler characteristic as $\chi = \sum_{s=1}^{n-1}(-1)^{s-1}\binom{n}{s}$ and applying the binomial identity $\sum_{s=0}^{n}(-1)^s\binom{n}{s} = 0$, we obtain $\sum_{s=1}^{n-1}(-1)^s\binom{n}{s} = -1-(-1)^n$ and hence $\chi(K(\Sigma)) = 1+(-1)^n$, which equals 2 for even n and 0 for odd n, matching $\chi(S^{n-2})$.

For (iv), when Σ is not a core, it has proper impossible sub-systems, which cause additional "holes" in $K(\Sigma)$ relative to the boundary of the full simplex. The topology of $K(\Sigma)$ in this case depends on the precise structure of the impossibility cores and their overlaps. □

Example 11.4 (Running Example: Feasibility Complex)**.** In Example 3.1, the only impossibility core is $\{A_1, A_2\}$. The feasibility complex $K(\Sigma)$ consists of all subsets of Σ that do not contain both A_1 and A_2. Its maximal faces are $\{A_1, A_3, A_4\}$ and $\{A_2, A_3, A_4\}$, so $K(\Sigma)$ is the union of two 2-simplices sharing the common edge $\{A_3, A_4\}$. This complex is contractible (it deformation retracts onto the shared edge), consistent with Theorem 11.7, since $d(\Sigma) = 1$. The f-vector is $(4, 5, 2)$ and the Euler characteristic is $\chi = 4 - 5 + 2 = 1$, consistent with contractibility.

Example 11.5 (Arrow's System)**.** Arrow's system has $n = 4$ axioms and $\dim(\Sigma) = 4$, so it is an impossibility core. By Theorem 11.3(ii), $K(\Sigma)$ is homeomorphic to S^2 (the ordinary 2-sphere). The Euler characteristic is $\chi = 2$, matching $\chi(S^2) = 2$. The faces of $K(\Sigma)$ consist of four vertices (the single axioms), six edges (the pairs), and four triangles (the triples). Geometrically, the feasibility complex of Arrow's system is the boundary of a tetrahedron. Figure 3 illustrates this structure.

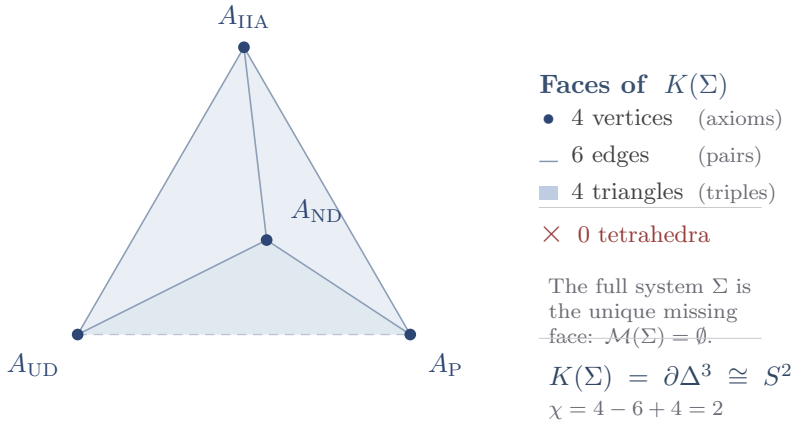


Figure 3: The feasibility complex of Arrow's four-axiom system of social choice theory. Since $\Sigma = \{A_{\text{UD}}, A_{\text{P}}, A_{\text{IIA}}, A_{\text{ND}}\}$ is an impossibility core of dimension 4, every proper subsystem is possible, and $K(\Sigma)$ consists of all proper subsets of Σ. This is the boundary of the 3-simplex (a tetrahedron), homeomorphic to the 2-sphere S^2. The unique missing face is the full system Σ itself, representing the impossibility $\mathcal{M}(\Sigma) = \emptyset$. The dashed edge indicates a hidden edge in the 3-dimensional perspective.

Remark 11.6 (Combinatorial Duality)**.** The feasibility complex $K(\Sigma)$ and the compatibility hypergraph $H(\Sigma)$ are complementary. Indeed, we have that $S \in K(\Sigma)$ if and only if $S \notin H(\Sigma)$. Proposition 10.9 says $H(\Sigma)$ is an upset; Proposition 11.2 says $K(\Sigma)$ is a downset (simplicial complex). Together they partition 2^Σ into the possible and the impossible subsystems. The minimal elements of $H(\Sigma)$ are the impossibility cores, which are exactly the "missing faces" of $K(\Sigma)$ relative to Δ^{n-1}. They are the smallest sub-systems not in $K(\Sigma)$. In this sense, the impossibility cores are the "holes" that prevent $K(\Sigma)$ from being the full simplex, and the impossibility dimension $\dim(\Sigma)$ is the size of the smallest hole.

We now establish a sufficient topological condition for rescuability.

Theorem 11.7 (Contractibility of the Feasibility Complex)**.** *Let Σ be an impossible axiom system. If $d(\Sigma) = 1$, then $K(\Sigma)$ is contractible.*

Proof. Since $d(\Sigma) = 1$, Theorem 9.8 gives $\tau(\mathcal{C}(\Sigma)) = 1$, so some single axiom $A^* \in \Sigma$ belongs to every impossibility core. Consider the face-poset endomap $r\colon K(\Sigma) \to K(\Sigma)$ defined by $r(S) = S \setminus \{A^*\}$. Since $r(S) \subseteq S$ and $K(\Sigma)$ is closed under taking subsets, $r(S) \in K(\Sigma)$ for every S, and the image of r lies in the sub-complex of faces avoiding A^*. Because A^* belongs to every impossibility core, the reduced system $\Sigma \setminus \{A^*\}$ contains no impossibility core, so every sub-system of $\Sigma \setminus \{A^*\}$ is possible. It follows that $K(\Sigma \setminus \{A^*\})$ is the full simplex $\Delta^{|\Sigma|-2}$ on vertex set $\Sigma \setminus \{A^*\}$, and is therefore contractible.

The map r is order-preserving on the face poset of $K(\Sigma)$ and satisfies $r(S) \subseteq S$ for every S, i.e., $r \leq \mathrm{id}_{K(\Sigma)}$ pointwise. We now invoke the following standard fact from combinatorial topology (the *order homotopy lemma*; see, e.g., Kozlov 2008, Chapter 13): if P is a finite poset and $f, g\colon P \to P$ are two order-preserving maps satisfying $f(x) \leq g(x)$ for every $x \in P$, then the induced maps on the geometric realization of P are homotopic. In particular, $f \leq \mathrm{id}_P$ implies that $|f|$ is homotopic to the identity. Applied to r and $\mathrm{id}_{K(\Sigma)}$, this yields $|r| \simeq |\mathrm{id}_{K(\Sigma)}|$. The map r factors as $r = \iota \circ r'$, where $\iota\colon K(\Sigma \setminus \{A^*\}) \hookrightarrow K(\Sigma)$ is the inclusion and $r'\colon K(\Sigma) \to K(\Sigma \setminus \{A^*\})$ is the corestriction of r. In particular, $\iota \circ r' \simeq \mathrm{id}_{K(\Sigma)}$, and since $r' \circ \iota = \mathrm{id}$, the inclusion ι is a homotopy equivalence. Hence, $K(\Sigma)$ is homotopy equivalent to the contractible complex $K(\Sigma \setminus \{A^*\})$, so $K(\Sigma)$ is itself contractible. □

Remark 11.8 (The Converse Fails)**.** One might expect the converse to hold, namely that $K(\Sigma)$ is contractible only if $d(\Sigma) = 1$. However, this is false. Let $\Sigma = \{A_1, A_2, A_3, A_4, A_5\}$ with impossibility cores $C_1 = \{A_1, A_2\}$ and $C_2 = \{A_3, A_4\}$. Since $C_1 \cap C_2 = \emptyset$, we have $d(\Sigma) = 2$, yet $K(\Sigma)$ is contractible because A_5 belongs to no core and therefore acts as a cone point. The Alexander Duality Theorem below gives a complete explanation of this phenomenon.

The Contractibility Theorem identifies a sufficient condition for acyclicity, but gives no information about the homology of $K(\Sigma)$ when the condition fails. We now prove a structural theorem that computes the homology of the feasibility complex entirely in terms of the overlap pattern of the impossibility cores. The key is Alexander duality, one of the fundamental results in combinatorial topology (see, e.g., Kozlov, 2008, Chapter 5).

Definition 11.9 (The Co-Core Complex)**.** Let Σ be an impossible axiom system with core family $\mathcal{C}(\Sigma)$. The *co-core complex* of Σ is

$$K^\vee(\Sigma) = \big\{ S \subseteq \Sigma \;\big|\; \mathcal{M}(\Sigma \setminus S) = \emptyset \big\},$$

the collection of all subsets whose complement in Σ is impossible.

Proposition 11.10 (Structure of the Co-Core Complex)**.** *The co-core complex $K^\vee(\Sigma)$ is a simplicial complex. Its maximal faces are the sets $\Sigma \setminus C$ for $C \in \mathcal{C}(\Sigma)$, one per impossibility core.*

Proof. If $S \in K^\vee(\Sigma)$ and $S' \subseteq S$, then $\Sigma \setminus S' \supseteq \Sigma \setminus S$, so $\mathcal{M}(\Sigma \setminus S') \subseteq \mathcal{M}(\Sigma \setminus S) = \emptyset$ by Theorem 4.1(i), giving $S' \in K^\vee(\Sigma)$. For the facet characterization, observe that S is maximal in $K^\vee(\Sigma)$ if and only if $\Sigma \setminus S$ is impossible but $\Sigma \setminus (S \cup \{A\})$ is possible for every

$A \in \Sigma \setminus S$. Since impossibility is upward-closed (Proposition 10.9), single-element removal suffices to verify minimality, so this condition is equivalent to $\Sigma \setminus S$ being an impossibility core. □

Theorem 11.11 (Alexander Duality of the Feasibility Complex)**.** *Let Σ be an impossible axiom system with $|\Sigma| = n$ and let $\mathbb{F}$ be any field. The co-core complex $K^\vee(\Sigma)$ is the Alexander dual of $K(\Sigma)$, and their reduced homology groups are related by*

$$\widetilde{H}_i\big(K(\Sigma);\, \mathbb{F}\big) \;\cong\; \widetilde{H}_{n-i-3}\big(K^\vee(\Sigma);\, \mathbb{F}\big) \qquad \text{for all } i \geq -1.$$

In particular, the homology of the feasibility complex is entirely determined by the overlap pattern of the impossibility cores.

Proof. Note that the Alexander dual of a simplicial complex K on vertex set V is equal to $K^\vee = \{S \subseteq V \mid V \setminus S \notin K\}$. In our setting, $V \setminus S \notin K(\Sigma)$ means $\mathcal{M}(V \setminus S) = \emptyset$, which is precisely the defining condition of $K^\vee(\Sigma)$. The isomorphism then follows from the Alexander duality theorem for simplicial complexes (see, e.g., Kozlov, 2008, Chapter 5), which asserts $\widetilde{H}_i(K;\, \mathbb{F}) \cong \widetilde{H}_{n-i-3}(K^\vee;\, \mathbb{F})$ over any field $\mathbb{F}$. □

Theorem 11.11 has several immediate consequences.

Corollary 11.12 (Homological Consequences)**.** *Let Σ be an impossible axiom system with $|\Sigma| = n$.*

(i) Top Betti number. *$\widetilde{\beta}_{n-2}(K(\Sigma)) = 1$ if Σ is itself an impossibility core, and $\widetilde{\beta}_{n-2}(K(\Sigma)) = 0$ otherwise.*

(ii) Acyclicity from a non-core axiom. *If some axiom $A \in \Sigma$ belongs to no impossibility core, then $K(\Sigma)$ is acyclic (all reduced homology vanishes over every field).*

(iii) Acyclicity from $d(\Sigma) = 1$. *If $d(\Sigma) = 1$, then $K(\Sigma)$ is acyclic. This follows from the Contractibility Theorem 11.7, which gives the stronger conclusion that $K(\Sigma)$ is contractible.*

Proof. For (i), by Theorem 11.11, $\widetilde{\beta}_{n-2}(K(\Sigma)) = \widetilde{\beta}_{-1}(K^\vee(\Sigma))$. The reduced Betti number $\widetilde{\beta}_{-1}$ equals 1 if and only if the complex is the void complex (containing no vertices), and 0 otherwise. Now $K^\vee(\Sigma)$ has a vertex $\{A\}$ if and only if $\mathcal{M}(\Sigma \setminus \{A\}) = \emptyset$, i.e., removing A preserves impossibility. If Σ is a core, removing any axiom restores possibility, so $K^\vee(\Sigma)$ has no vertices and $\widetilde{\beta}_{-1}(K^\vee) = 1$. If Σ is not a core, some axiom can be removed while preserving impossibility, so $K^\vee(\Sigma)$ has at least one vertex and $\widetilde{\beta}_{-1}(K^\vee) = 0$.

For (ii), if A belongs to no core, then $A \in \Sigma \setminus C$ for every $C \in \mathcal{C}(\Sigma)$, so A lies in every maximal face of $K^\vee(\Sigma)$. Thus, A is a cone point of $K^\vee(\Sigma)$, so $K^\vee(\Sigma)$ is contractible and all its reduced homology vanishes. By Theorem 11.11, all reduced homology of $K(\Sigma)$ vanishes.

For (iii), this is a consequence of the Contractibility Theorem 11.7 rather than of Alexander duality per se. That theorem establishes that $K(\Sigma)$ is contractible whenever $d(\Sigma) = 1$, and a contractible complex is acyclic. We note that (iii) does not follow from (ii) in general, since $d(\Sigma) = 1$ means some axiom belongs to *every* core, which does not imply that some axiom belongs to *no* core. □

Example 11.13 (Alexander Duality in the Paper's Examples)**.** We verify Theorem 11.11 against the examples developed above.

In the running example (Example 3.1), the sole core is $\{A_1, A_2\}$, so $K^\vee(\Sigma)$ has unique maximal face $\{A_3, A_4\}$, a single edge. This is contractible, so Corollary 11.12(ii) gives $K(\Sigma)$ acyclic, consistent with the Euler characteristic $\chi = 1$ computed in Example 11.4.

For Arrow's system (Example 10.4), Σ itself is the sole core, so $K^\vee(\Sigma)$ is the void complex (no vertices). Corollary 11.12(i) gives $\widetilde{\beta}_2(K(\Sigma)) = 1$, confirming that $K(\Sigma) \cong S^2$.

Regarding Example 11.14 below ($n = 5$, cores $C_1 = \{A_1, A_2, A_3\}$ and $C_2 = \{A_4, A_5\}$), the co-core complex $K^\vee(\Sigma)$ has maximal faces $\{A_4, A_5\}$ and $\{A_1, A_2, A_3\}$. Since $C_1 \cup C_2 = \Sigma$, these two faces are disjoint, so $K^\vee(\Sigma)$ is disconnected with $\widetilde{\beta}_0(K^\vee) = 1$. By Theorem 11.11, $\widetilde{\beta}_2(K(\Sigma)) = \widetilde{\beta}_0(K^\vee) = 1$, recovering $K(\Sigma) \cong S^2$ without any boundary-map computation.

Example 11.14 (Disjoint Cores and Nontrivial Homology)**.** Let $\Sigma = \{A_1, A_2, A_3, A_4, A_5\}$ and let $K(\Sigma)$ be the simplicial complex on vertex set $[5]$ with six maximal faces, namely $\{1, 2, 4\}$, $\{1, 2, 5\}$, $\{1, 3, 4\}$, $\{1, 3, 5\}$, $\{2, 3, 4\}$, and $\{2, 3, 5\}$. This is the boundary of the bipyramid over the triangle $\{1, 2, 3\}$ with poles 4 and 5. The minimal non-faces are $\{1, 2, 3\}$ and $\{4, 5\}$, giving two impossibility cores $C_1 = \{A_1, A_2, A_3\}$ and $C_2 = \{A_4, A_5\}$. Since $C_1 \cap C_2 = \emptyset$, we have $d(\Sigma) = 2$. Every axiom belongs to at least one core, so $K(\Sigma)$ has no cone point. By Example 11.13, the Alexander dual $K^\vee(\Sigma)$ is disconnected, and $\widetilde{\beta}_2(K(\Sigma)) = 1$, confirming that $K(\Sigma) \cong S^2$. The minimum rescue requires two drops (one from each core); dropping A_3 and A_4 leaves $\{A_1, A_2, A_5\}$, which is possible.

Remark 11.15 (Connection to the Rescue Theorem)**.** Theorem 11.7 gives a topological interpretation of the condition $d(\Sigma) = 1$, and Theorem 11.11 reveals the deeper structure. The Rescue Theorem (Theorem 9.4) and the Transversal Theorem (Theorem 9.8) are the combinatorial shadows of this duality. The possibility side $K(\Sigma)$ and the impossibility side $K^\vee(\Sigma)$ are Alexander duals, and their homologies are locked together. When $d \geq 2$ and every axiom belongs to some core, the nontrivial homology of $K(\Sigma)$ encodes a genuine topological obstruction to possibility.

12 A Catalogue of Fully Derived Examples

We now illustrate the framework with a catalogue of worked examples that exercise the full machinery of Sections 8–9. Every example in this section is fully derived. The domain is explicit and finite, and every model-class size, Shapley value, interaction, and homological invariant is obtained by direct enumeration rather than by stipulation. The examples are chosen to cover distinct regimes and distinct structural phenomena of the theory. Section 12.1 treats a Condorcet impossibility of dimension 3 with two overlapping cores. Section 12.2 constructs a miniature Arrow-type impossibility of dimension 3 on a 9-element domain of two-profile choice functions. Section 12.3 constructs a miniature Nash bargaining uniqueness on a 16-element domain of solution selectors. Section 12.4 exhibits a dimension-4 impossibility core and verifies the Dimension-Witness Duality on it. Section 12.5 exhibits a five-axiom system with two disjoint cores whose feasibility complex has nontrivial second Betti number, verifying Alexander duality outside the trivial regime. Section 12.6 exhibits a six-axiom system with three pairwise-disjoint cores, whose co-core complex has a nontrivial 1-cycle (in contrast to the mere disconnectedness of the two-core case), demonstrating a structurally

different mechanism of Alexander duality. The extrapolations to the classical infinite-domain theorems of Arrow, Nash, and the Shapley value axiomatization are discussed separately in Section 13, and are not treated as theorems of this paper.

12.1 The Condorcet Cycle With a Positional Constraint

Consider the six strict linear orders on three alternatives $\{a, b, c\}$, listed in Table 3. We take $\mathcal{U} = \{abc, acb, bac, bca, cab, cba\}$ and equip it with the uniform measure $\mu(\{x\}) = 1/6$. Each element of $\mathcal{U}$ is a ranking of the three alternatives (so bca denotes $b \succ c \succ a$), and every model-class size below is obtained by direct inspection.

Ranking	Interpretation
abc	$a \succ b \succ c$
acb	$a \succ c \succ b$
bac	$b \succ a \succ c$
bca	$b \succ c \succ a$
cab	$c \succ a \succ b$
cba	$c \succ b \succ a$

Table 3: The six strict linear orders on $\{a, b, c\}$.

We define four axioms, each a subset of $\mathcal{U}$. The first three are pairwise majority constraints forming a Condorcet cycle (a beats b beats c beats a), and the fourth is a positional constraint on a.

$$A_1 = \{abc, acb, cab\}, \qquad A_2 = \{abc, bac, bca\},$$
$$A_3 = \{bca, cab, cba\}, \qquad A_4 = \{abc, acb, bac, cab\}.$$

In words, A_1 requires $a \succ b$, A_2 requires $b \succ c$, A_3 requires $c \succ a$, and A_4 requires that a is not ranked last.

Impossibility and cores. The full model class is $\mathcal{M}(\Sigma) = A_1 \cap A_2 \cap A_3 \cap A_4 = \emptyset$, so the system is in the impossibility regime. No pair of axioms has empty intersection (one verifies that $A_1 \cap A_2 = \{abc\}$, $A_1 \cap A_3 = \{cab\}$, $A_2 \cap A_3 = \{bca\}$, $A_1 \cap A_4 = A_1$, $A_2 \cap A_4 = \{abc, bac\}$, and $A_3 \cap A_4 = \{cab\}$), so the impossibility dimension is at least 3. In fact, the system has exactly two impossibility cores, both of size 3. The Condorcet core is $C_1 = \{A_1, A_2, A_3\}$, since $A_1 \cap A_2 \cap A_3 = \emptyset$ and every proper sub-pair has nonempty intersection. The positional core is $C_2 = \{A_2, A_3, A_4\}$, since $A_2 \cap A_3 \cap A_4 = \emptyset$ and every proper sub-pair has nonempty intersection. No other triple is impossible (one checks that $\mathcal{M}(\{A_1, A_2, A_4\}) = \{abc\}$ and $\mathcal{M}(\{A_1, A_3, A_4\}) = \{cab\}$), so the full core family is $\mathcal{C}(\Sigma) = \{C_1, C_2\}$, with $\dim(\Sigma) = 3$.

The axiom game. Table 4 gives the model-class sizes and collapse values for all subsystems. Every entry is obtained by direct enumeration.

Note that $v_\mu(\{A_1, A_4\}) = v_\mu(\{A_1\}) = 1/2$, so A_4 contributes zero marginal collapse beyond A_1; the structural reason for this equality shall emerge below.

Sub-system S	$\mu(\mathcal{M}(S))$	$v_\mu(S)$
$\emptyset$	1	0
$\{A_1\}$	1/2	1/2
$\{A_2\}$	1/2	1/2
$\{A_3\}$	1/2	1/2
$\{A_4\}$	2/3	1/3
$\{A_1, A_2\}$	1/6	5/6
$\{A_1, A_3\}$	1/6	5/6
$\{A_1, A_4\}$	1/2	1/2
$\{A_2, A_3\}$	1/6	5/6
$\{A_2, A_4\}$	1/3	2/3
$\{A_3, A_4\}$	1/6	5/6
$\{A_1, A_2, A_3\}$	0	1
$\{A_1, A_2, A_4\}$	1/6	5/6
$\{A_1, A_3, A_4\}$	1/6	5/6
$\{A_2, A_3, A_4\}$	0	1
Σ	0	1

Table 4: Sub-system model-class sizes and collapse values for the Condorcet system, derived by direct enumeration.

Shapley values. We derive the axiom Shapley values directly from Table 4. For each axiom A and each predecessor set $S \subseteq \Sigma \setminus \{A\}$, the marginal contribution is equal to $\Delta_A(S) = v_\mu(S \cup \{A\}) - v_\mu(S)$, and the Shapley weight is $w_{|S|} = |S|!(3 - |S|)!/4!$, giving $w_0 = w_3 = 1/4$ and $w_1 = w_2 = 1/12$. We compute each axiom's weighted marginal sum in turn.

For A_1, the eight marginal contributions in predecessor sets of size $0, 1, 2, 3$ are

$$\begin{aligned} \Delta_{A_1}(\varnothing) &= \tfrac{1}{2}, & \Delta_{A_1}(\{A_2\}) &= \tfrac{1}{3}, \\ \Delta_{A_1}(\{A_3\}) &= \tfrac{1}{3}, & \Delta_{A_1}(\{A_4\}) &= \tfrac{1}{6}, \\ \Delta_{A_1}(\{A_2, A_3\}) &= \tfrac{1}{6}, & \Delta_{A_1}(\{A_2, A_4\}) &= \tfrac{1}{6}, \\ \Delta_{A_1}(\{A_3, A_4\}) &= 0, & \Delta_{A_1}(\{A_2, A_3, A_4\}) &= 0, \end{aligned}$$

giving $\phi_{A_1} = \frac{1}{4} \cdot \frac{1}{2} + \frac{1}{12}(\frac{1}{3} + \frac{1}{3} + \frac{1}{6}) + \frac{1}{12}(\frac{1}{6} + \frac{1}{6} + 0) + \frac{1}{4} \cdot 0 = \frac{8}{36}$.

For A_2, the marginals are $\frac{1}{2}$, $\frac{1}{3}$, $\frac{1}{3}$, $\frac{1}{3}$, $\frac{1}{6}$, $\frac{1}{3}$, $\frac{1}{6}$, $\frac{1}{6}$ (in the same predecessor-size order over subsets of $\{A_1, A_3, A_4\}$), giving $\phi_{A_2} = \frac{11}{36}$. For A_3, the marginals are $\frac{1}{2}$, $\frac{1}{3}$, $\frac{1}{3}$, $\frac{1}{2}$, $\frac{1}{6}$, $\frac{1}{3}$, $\frac{1}{3}$, $\frac{1}{6}$, giving $\phi_{A_3} = \frac{12}{36}$. For A_4, the marginals are $\frac{1}{3}$, 0, $\frac{1}{6}$, $\frac{1}{3}$, 0, 0, $\frac{1}{6}$, 0, giving $\phi_{A_4} = \frac{5}{36}$. These values sum to $(8 + 11 + 12 + 5)/36 = 1 = v_\mu(\Sigma)$, confirming efficiency. The ranking is $\phi_{A_3} > \phi_{A_2} > \phi_{A_1} > \phi_{A_4}$.

The blame ranking is a genuine discovery. At first glance, A_1, A_2, and A_3 are symmetric. Each is a pairwise majority constraint, and each is satisfied by exactly three of the six rankings. One might therefore expect them to share equal blame. They do not. The axiom A_3 receives the most blame (33.3%), A_2 is close behind (30.6%), and A_1 receives the least (22.2%) among the three cycle axioms. This asymmetry is not an artifact of

Axiom	ϕ_A	Decimal	Blame share
$c \succ a$ (A_3)	12/36	0.333	33.3%
$b \succ c$ (A_2)	11/36	0.306	30.6%
$a \succ b$ (A_1)	8/36	0.222	22.2%
a not last (A_4)	5/36	0.139	13.9%
Total	1	1.000	100%

Table 5: Axiom Shapley values for the Condorcet system (derived, not stipulated).

the measure or an assumption built into the inputs; it is a consequence of the structural position of A_4. Indeed, a direct check shows that $A_1 \subseteq A_4$; any ranking placing a above b certainly does not place a last. No such containment holds among the other cycle axioms, and this breaks the apparent symmetry. Concretely, whenever A_4 is already in force, the positional constraint A_4 partially absorbs A_1's restrictive power, so the marginal collapse from adding A_1 is smaller. In contrast, A_3 conflicts most sharply with A_4 (their intersection is the single ranking cab), so adding A_3 contributes substantial marginal collapse in every context containing A_4. The axiom A_2 falls in between, interacting moderately with A_4 (intersection $\{abc, bac\}$, two elements). This graduated blame ranking is invisible without the Shapley computation, and would not be predicted by counting models alone.

Pairwise interactions. The pairwise interactions at the empty coalition (Theorem 8.8) are:

$$\iota_\mu(A_1, A_4 \mid \emptyset) = -\tfrac{1}{3}, \qquad \iota_\mu(A_3, A_4 \mid \emptyset) = 0, \qquad \iota_\mu(A_i, A_j \mid \emptyset) = -\tfrac{1}{6} \text{ otherwise.}$$

The strongest diminishing returns occur between A_1 and A_4 ($\iota = -1/3$), reflecting their containment relation. Since $A_1 \subseteq A_4$, axiom A_1 already enforces everything A_4 does, so their combined collapse is no more than A_1's alone. The interaction $\iota_\mu(A_3, A_4 \mid \emptyset) = 0$ means that A_3 and A_4 are coverage-exhaustive. Indeed, $A_3 \cup A_4 = \mathcal{U}$, so together they cover the entire universe. In fact, every ranking either has $c \succ a$ (and belongs to A_3) or has a not last (and belongs to A_4), or both.

Rescue analysis. The full-coalition marginals are $\Delta_{A_2} = \Delta_{A_3} = 1/6$ and $\Delta_{A_1} = \Delta_{A_4} = 0$. By Theorem 9.8, the impossibility depth is $d(\Sigma) = \tau(\mathcal{C}(\Sigma)) = 1$, since A_2 and A_3 each belong to both cores C_1 and C_2. Dropping A_2 restores possibility, giving $\mathcal{M}(\{A_1, A_3, A_4\}) = \{cab\}$, a uniqueness result (the ranking $c \succ a \succ b$). Dropping A_3 also restores possibility, giving $\mathcal{M}(\{A_1, A_2, A_4\}) = \{abc\}$ (the ranking $a \succ b \succ c$). Dropping A_1 or A_4 alone does not help, because the surviving system still contains a core. Note the divergence between the Shapley ranking and the rescue criterion. The axiom A_1 has a higher Shapley value (8/36) than A_4 (5/36), yet both have the same full-coalition marginal ($\Delta = 0$), and neither is a viable single-axiom drop. This confirms the general observation of Remark 9.5.

The feasibility complex and Alexander duality. We now analyze the topology of the feasibility complex $K(\Sigma)$ in detail, illustrating the Alexander Duality Theorem (Theorem 11.11) on this example. The possible sub-systems are all subsets of Σ that

contain neither $C_1 = \{A_1, A_2, A_3\}$ nor $C_2 = \{A_2, A_3, A_4\}$ as a subset. The maximal possible sub-systems (facets of $K(\Sigma)$) are

$$\{A_1, A_2, A_4\}, \quad \{A_1, A_3, A_4\}, \quad \text{and} \quad \{A_2, A_3\}.$$

Indeed, $\mathcal{M}(\{A_1, A_2, A_4\}) = \{abc\} \neq \emptyset$ and $\mathcal{M}(\{A_1, A_3, A_4\}) = \{cab\} \neq \emptyset$ (both derived in Table 4), and each becomes impossible upon adding the remaining axiom; the pair $\{A_2, A_3\}$ with $\mathcal{M} = \{bca\}$ is maximal because both $\{A_1, A_2, A_3\}$ and $\{A_2, A_3, A_4\}$ are impossible. The f-vector is $(4, 6, 2)$ (four vertices, six edges, two triangles), and the Euler characteristic is $\chi = 4 - 6 + 2 = 0$. Note that $\chi(S^1) = 0$, which already suggests the homotopy type.

The co-core complex $K^\vee(\Sigma)$ has, by Proposition 11.10, maximal faces $\Sigma \setminus C_1 = \{A_4\}$ and $\Sigma \setminus C_2 = \{A_1\}$. In other words, $K^\vee(\Sigma)$ consists of two isolated vertices $\{A_1\}$ and $\{A_4\}$, with no edge between them. This complex is disconnected, with $\widetilde{\beta}_0(K^\vee(\Sigma)) = 1$. Note that A_2 and A_3 appear in both cores, so neither belongs to any maximal face of $K^\vee(\Sigma)$; they are absent from the co-core complex entirely. The isolated-vertex structure reflects the fact that the two cores overlap in $\{A_2, A_3\}$, preventing their complements from sharing any common axiom.

Applying Theorem 11.11 with $n = 4$ and $i = 1$, we obtain

$$\widetilde{H}_1\big(K(\Sigma);\, \beth\big) \;\cong\; \widetilde{H}_{4-1-3}\big(K^\vee(\Sigma);\, \beth\big) \;=\; \widetilde{H}_0\big(K^\vee(\Sigma);\, \beth\big) \;\cong\; \beth.$$

All other reduced Betti numbers of $K(\Sigma)$ vanish (one checks that $\widetilde{\beta}_0(K) = 0$ since K is connected, and $\widetilde{\beta}_2(K) = \widetilde{\beta}_{-1}(K^\vee) = 0$ since $K^\vee$ is nonempty). Hence, $K(\Sigma) \simeq S^1$. The feasibility complex is homotopy equivalent to a circle, with a single nontrivial 1-cycle. Geometrically, $K(\Sigma)$ consists of two triangles $\{A_1, A_2, A_4\}$ and $\{A_1, A_3, A_4\}$ sharing the edge $\{A_1, A_4\}$, together with the maximal edge $\{A_2, A_3\}$ bridging the two triangles along their unshared vertices. The resulting complex has the topology of a band that closes into a loop. The two missing triangles (the cores C_1 and C_2) are the "holes" in the boundary of the full 3-simplex that create this loop. Figure 4 illustrates this structure.

This example illustrates the power of Alexander duality. The nontrivial 1-cycle in $K(\Sigma)$ is not obvious from direct inspection of the feasibility complex, but it follows immediately from the observation that $K^\vee(\Sigma)$ is disconnected, a property that can be read off from the core family $\mathcal{C}(\Sigma) = \{C_1, C_2\}$ by checking that $(\Sigma \setminus C_1) \cap (\Sigma \setminus C_2) = \Sigma \setminus (C_1 \cup C_2) = \emptyset$.

The axiom entropy is $H \approx 1.34$ nats, close to the maximum $\ln 4 \approx 1.39$, reflecting the relatively even distribution of blame.

12.2 A Miniature Impossibility of the Arrow Type

We now construct a fully derived Arrow-type impossibility on a small explicit finite domain of choice functions. The construction captures the qualitative essence of Arrow's theorem, namely the incompatibility of Pareto-efficiency, structural symmetry, and non-dictatorship, without relying on stipulation. Every model-class size below is obtained by direct enumeration of a 9-element domain.

Let the alternative set be $\{a, b, c\}$, and consider two voters. Fix two profiles P_1 and P_2 over these alternatives. The profile P_1 has voter 1 ranking $a \succ b \succ c$ and voter 2 ranking $b \succ a \succ c$. The profile P_2 is obtained from P_1 by swapping the two voters; that is, voter

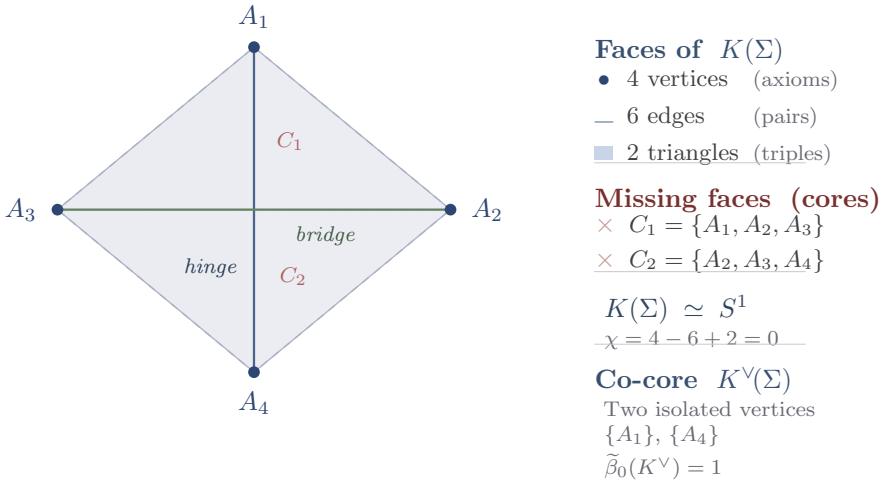


Figure 4: The feasibility complex of the Condorcet system. The two shaded triangles $\{A_1, A_2, A_4\}$ (right) and $\{A_1, A_3, A_4\}$ (left) are the maximal possible sub-systems, sharing the vertical hinge edge $\{A_1, A_4\}$. The horizontal bridge edge $\{A_2, A_3\}$ is a maximal face of dimension 1, connecting the two triangles along their unshared vertices and closing the complex into a loop. The two dashed triangles are the impossibility cores $C_1 = \{A_1, A_2, A_3\}$ and $C_2 = \{A_2, A_3, A_4\}$, whose absence creates a nontrivial 1-cycle. The resulting complex has the homotopy type $K(\Sigma) \simeq S^1$ with $\chi = 0$. By Alexander duality, it holds that $\widetilde{H}_1(K(\Sigma)) \cong \widetilde{H}_0(K^\vee(\Sigma)) \cong \neg$, since the co-core complex consists of two disconnected vertices.

1 ranks $b \succ a \succ c$ and voter 2 ranks $a \succ b \succ c$. Note that P_2 is also the profile obtained from P_1 by swapping the alternatives a and b (since swapping voters and swapping the alternatives a and b produce the same profile in this symmetric setup, a coincidence we shall exploit below).

Let $\mathcal{U}$ denote the set of all choice functions $c\colon \{P_1, P_2\} \to \{a, b, c\}$. Every such function is specified by the pair $(c(P_1), c(P_2))$, so $|\mathcal{U}| = 9$. We write elements of $\mathcal{U}$ as ordered pairs; for example, the choice function assigning a to P_1 and b to P_2 is denoted (a, b). We equip $\mathcal{U}$ with the uniform measure $\mu(\{c\}) = 1/9$.

We now define four axioms, each expressing a standard requirement on choice functions.

$$
\begin{aligned}
A_{\mathrm{P}} &= \big\{\, c \in \mathcal{U} \;\big|\; c(P_1) \in \{a, b\} \text{ and } c(P_2) \in \{a, b\} \,\big\}, \\
A_{\mathrm{A}} &= \big\{\, c \in \mathcal{U} \;\big|\; c(P_1) = c(P_2) \,\big\}, \\
A_{\mathrm{N}} &= \big\{\, c \in \mathcal{U} \;\big|\; c(P_2) = \sigma(c(P_1)) \,\big\}, \\
A_{\mathrm{D}} &= \mathcal{U} \setminus \big\{\, (a, b),\ (b, a) \,\big\},
\end{aligned}
$$

where σ denotes the transposition on $\{a, b, c\}$ that swaps a and b and fixes c. The axiom A_{P} is the Pareto axiom. Alternative c is Pareto-dominated by both a and b in each of the two profiles, so a Pareto-efficient choice function never selects c. The axiom A_{A} is anonymity. Since P_2 is obtained from P_1 by permuting the voters, an anonymous choice function must return the same outcome on both profiles. The axiom A_{N} is neutrality with respect to the

swap $a \leftrightarrow b$. Since P_2 is also obtained from P_1 by swapping the alternatives a and b, a neutral choice function must return outcomes on P_1 and P_2 that are related by the same swap. The axiom A_{D} is non-dictatorship. The choice function (a, b) always picks voter 1's top-ranked alternative, and (b, a) always picks voter 2's top-ranked alternative; these two functions are the dictatorships and are excluded by A_{D}.

By direct enumeration, $|A_{\mathrm{P}}| = 4$, $|A_{\mathrm{A}}| = 3$, $|A_{\mathrm{N}}| = 3$, and $|A_{\mathrm{D}}| = 7$. The elements of the four axioms are

$$\begin{aligned}
A_{\mathrm{P}} &= \{(a,a),(a,b),(b,a),(b,b)\},\\
A_{\mathrm{A}} &= \{(a,a),(b,b),(c,c)\},\\
A_{\mathrm{N}} &= \{(a,b),(b,a),(c,c)\},\\
A_{\mathrm{D}} &= \{(a,a),(a,c),(b,b),(b,c),(c,a),(c,b),(c,c)\}.
\end{aligned}$$

Impossibility and cores. The full intersection is $\mathcal{M}(\Sigma) = A_{\mathrm{P}} \cap A_{\mathrm{A}} \cap A_{\mathrm{N}} \cap A_{\mathrm{D}} = \emptyset$, so the system is in the impossibility regime. Indeed, $A_{\mathrm{A}} \cap A_{\mathrm{N}} = \{(c,c)\}$, since anonymity and neutrality together force $c(P_1) = c(P_2) = \sigma(c(P_1))$, and the only fixed point of σ is c; but $(c,c) \notin A_{\mathrm{P}}$, since $c(P_1) = c$ violates Pareto. No pair of axioms has empty intersection, so the impossibility dimension is at least 3. We claim the system has exactly two impossibility cores, both of dimension 3. The first core is $C_1 = \{A_{\mathrm{P}}, A_{\mathrm{A}}, A_{\mathrm{N}}\}$, since $A_{\mathrm{P}} \cap A_{\mathrm{A}} \cap A_{\mathrm{N}} = \emptyset$ (as just shown) and every proper sub-pair has nonempty intersection. The second core is $C_2 = \{A_{\mathrm{P}}, A_{\mathrm{N}}, A_{\mathrm{D}}\}$, since $A_{\mathrm{P}} \cap A_{\mathrm{N}} = \{(a,b),(b,a)\}$ consists of the two dictatorships, and both are excluded by A_{D}; every proper sub-pair has nonempty intersection. The remaining two triples are possible: $A_{\mathrm{P}} \cap A_{\mathrm{A}} \cap A_{\mathrm{D}} = \{(a,a),(b,b)\}$ and $A_{\mathrm{A}} \cap A_{\mathrm{N}} \cap A_{\mathrm{D}} = \{(c,c)\}$, so neither is a core. Hence, $\mathcal{C}(\Sigma) = \{C_1, C_2\}$ and $\dim(\Sigma) = 3$.

The axiom game. Table 6 records the model-class sizes and collapse values for all 16 sub-systems, each entry obtained by direct enumeration.

Shapley values. We derive the axiom Shapley values directly from Table 6, using the Shapley weights $w_0 = w_3 = 1/4$ and $w_1 = w_2 = 1/12$ as in the Condorcet derivation above.

For A_{P}, the eight marginal contributions $\Delta_{A_{\mathrm{P}}}(S) = v_\mu(S \cup \{A_{\mathrm{P}}\}) - v_\mu(S)$ over subsets $S \subseteq \{A_{\mathrm{A}}, A_{\mathrm{N}}, A_{\mathrm{D}}\}$ are, in the order $\emptyset$, singletons, pairs, full complement,

$$\tfrac{5}{9},\quad \tfrac{1}{9},\ \tfrac{1}{9},\ \tfrac{5}{9},\quad \tfrac{1}{9},\ \tfrac{1}{9},\ \tfrac{1}{9},\quad \tfrac{1}{9},$$

giving $\phi_{A_{\mathrm{P}}} = \frac{1}{4}\cdot\frac{5}{9} + \frac{1}{12}(\frac{1}{9}+\frac{1}{9}+\frac{5}{9}) + \frac{1}{12}(\frac{1}{9}+\frac{1}{9}+\frac{1}{9}) + \frac{1}{4}\cdot\frac{1}{9} = \frac{7}{27}$. For A_{A}, the marginals over subsets of $\{A_{\mathrm{P}}, A_{\mathrm{N}}, A_{\mathrm{D}}\}$ are $\frac{6}{9},\ \frac{2}{9},\ \frac{2}{9},\ \frac{4}{9},\ \frac{2}{9},\ 0,\ 0,\ 0$, giving $\phi_{A_{\mathrm{A}}} = \frac{7}{27}$. For A_{N}, the marginals over subsets of $\{A_{\mathrm{P}}, A_{\mathrm{A}}, A_{\mathrm{D}}\}$ are $\frac{6}{9},\ \frac{2}{9},\ \frac{2}{9},\ \frac{6}{9},\ \frac{2}{9},\ \frac{2}{9},\ \frac{2}{9},\ \frac{2}{9}$, giving $\phi_{A_{\mathrm{N}}} = \frac{10}{27}$. For A_{D}, the marginals over subsets of $\{A_{\mathrm{P}}, A_{\mathrm{A}}, A_{\mathrm{N}}\}$ are $\frac{2}{9},\ \frac{2}{9},\ 0,\ \frac{2}{9},\ 0,\ \frac{2}{9},\ 0,\ 0$, giving $\phi_{A_{\mathrm{D}}} = \frac{3}{27}$. These values sum to $(7+7+10+3)/27 = 1 = v_\mu(\Sigma)$, confirming efficiency. The ranking is $\phi_{A_{\mathrm{N}}} > \phi_{A_{\mathrm{P}}} = \phi_{A_{\mathrm{A}}} > \phi_{A_{\mathrm{D}}}$.

Interpretation. Neutrality bears the largest share of the collapse blame, at roughly 37%. Both A_{N} and A_{P} appear in both impossibility cores, so mere core-membership alone does not distinguish them; what does is the Axiom Dominance Theorem (Theorem 8.12).

Sub-system S	$\mu(\mathcal{M}(S))$	$v_\mu(S)$
$\emptyset$	1	0
$\{A_\mathrm{P}\}$	4/9	5/9
$\{A_\mathrm{A}\}$	3/9	6/9
$\{A_\mathrm{N}\}$	3/9	6/9
$\{A_\mathrm{D}\}$	7/9	2/9
$\{A_\mathrm{P}, A_\mathrm{A}\}$	2/9	7/9
$\{A_\mathrm{P}, A_\mathrm{N}\}$	2/9	7/9
$\{A_\mathrm{P}, A_\mathrm{D}\}$	2/9	7/9
$\{A_\mathrm{A}, A_\mathrm{N}\}$	1/9	8/9
$\{A_\mathrm{A}, A_\mathrm{D}\}$	3/9	6/9
$\{A_\mathrm{N}, A_\mathrm{D}\}$	1/9	8/9
$\{A_\mathrm{P}, A_\mathrm{A}, A_\mathrm{N}\}$	0	1
$\{A_\mathrm{P}, A_\mathrm{A}, A_\mathrm{D}\}$	2/9	7/9
$\{A_\mathrm{P}, A_\mathrm{N}, A_\mathrm{D}\}$	0	1
$\{A_\mathrm{A}, A_\mathrm{N}, A_\mathrm{D}\}$	1/9	8/9
Σ	0	1

Table 6: Sub-system model-class sizes and collapse values for the miniature Arrow-type system, derived by direct enumeration on the 9-element domain.

Direct inspection of Table 6 shows that in every context $S \subseteq \{A_\mathrm{A}, A_\mathrm{D}\}$, it holds that $\mu(\mathcal{M}(S \cup \{A_\mathrm{N}\})) \leq \mu(\mathcal{M}(S \cup \{A_\mathrm{P}\}))$; neutrality is uniformly more restrictive than Pareto, so Theorem 8.12(ii) forces $\phi_{A_\mathrm{N}} \geq \phi_{A_\mathrm{P}}$. Intuitively, neutrality alone leaves only three choice functions (namely (a, b), (b, a), and (c, c)), while Pareto alone leaves four; this gap in restrictiveness persists (and is often amplified) when either or both of the remaining axioms are added. Non-dictatorship receives the smallest blame, at roughly 11%, because it is a very mild constraint on its own (excluding only two of the nine choice functions) and contributes zero marginal collapse in the coalition $\{A_\mathrm{A}\}$ (adding A_D to $\{A_\mathrm{A}\}$ does not change the model class, since neither of the two dictatorships is anonymous).

Pairwise interactions. The pairwise interactions at the empty coalition, computed via Theorem 8.8(i) using the identity $\iota_\mu(A, B \mid \emptyset) = \mu(A \cup B) - \mu(\mathcal{U}) = \mu(A \cup B) - 1$, are

$$\begin{aligned} &\iota_\mu(A_\mathrm{P}, A_\mathrm{A} \mid \emptyset) = -\tfrac{4}{9}, \qquad \iota_\mu(A_\mathrm{P}, A_\mathrm{N} \mid \emptyset) = -\tfrac{4}{9}, \qquad \iota_\mu(A_\mathrm{P}, A_\mathrm{D} \mid \emptyset) = 0, \\ &\iota_\mu(A_\mathrm{A}, A_\mathrm{N} \mid \emptyset) = -\tfrac{4}{9}, \qquad \iota_\mu(A_\mathrm{A}, A_\mathrm{D} \mid \emptyset) = -\tfrac{2}{9}, \qquad \iota_\mu(A_\mathrm{N}, A_\mathrm{D} \mid \emptyset) = 0. \end{aligned}$$

Two pairs are coverage-exhaustive at the empty coalition, namely $(A_\mathrm{P}, A_\mathrm{D})$ and $(A_\mathrm{N}, A_\mathrm{D})$. In each case, the union of the two axioms exhausts $\mathcal{U}$. Indeed, $A_\mathrm{P} \cup A_\mathrm{D}$ misses only the pair of choice functions in neither, but A_P contains (a, b) and (b, a) while A_D contains everything else, so the union is $\mathcal{U}$. For $A_\mathrm{N} \cup A_\mathrm{D}$, the two dictatorships (a, b) and (b, a) are excluded from A_D but included in A_N, so the union is again $\mathcal{U}$.

Rescue analysis. Both cores contain A_P and A_N, so the transversal number of $\mathcal{C}(\Sigma)$ is 1, and Theorem 9.8 gives $d(\Sigma) = 1$. Dropping A_P yields $\mathcal{M}(\{A_\mathrm{A}, A_\mathrm{N}, A_\mathrm{D}\}) = \{(c, c)\}$,

a uniqueness result concentrated on the constant choice function selecting the Pareto-dominated alternative. Dropping A_{N} yields $\mathcal{M}(\{A_{\mathrm{P}}, A_{\mathrm{A}}, A_{\mathrm{D}}\}) = \{(a,a),(b,b)\}$, a possibility result with two surviving choice functions (constant a and constant b). Dropping A_{A} leaves the core C_2 intact, and dropping A_{D} leaves C_1 intact; neither of these single drops restores possibility. The axiom entropy is

$$H(\Sigma, \mu) = -\tfrac{7}{27}\ln\tfrac{7}{27} - \tfrac{7}{27}\ln\tfrac{7}{27} - \tfrac{10}{27}\ln\tfrac{10}{27} - \tfrac{3}{27}\ln\tfrac{3}{27} \approx 1.31 \text{ nats},$$

close to the maximum $\ln 4 \approx 1.39$ nats, reflecting the relatively even distribution of blame.

Discussion. This construction is a genuine impossibility theorem on the 9-element domain of two-profile choice functions, with every quantity derived by direct enumeration. It is not Arrow's theorem itself. The full Arrow theorem takes an infinite domain of social welfare functions and uses a fundamentally different axiom (independence of irrelevant alternatives) in place of neutrality. The relationship between the two is the following. Both are instances of the abstract pattern in which Pareto efficiency, a symmetry condition on the response to profiles, and non-dictatorship are jointly inconsistent. The miniature version realizes this pattern in a domain small enough that the entire framework of Sections 8–9 applies without stipulation. Extrapolations to the classical Arrow theorem are discussed in Section 13.

12.3 A Miniature Nash Bargaining Uniqueness

We now construct a fully derived Nash-type uniqueness result on a small explicit finite domain of bargaining solutions. As with the miniature Arrow, the construction captures the qualitative essence of Nash's theorem, namely the joint force of efficiency, symmetry, and scale invariance in singling out a canonical solution, without relying on stipulation.

Consider two two-player bargaining problems with disagreement point $(0,0)$. The first is $B_1 = \{(2,2),(3,0),(0,3),(1,1)\}$, the second is $B_2 = \{(4,4),(6,0),(0,6),(2,2)\}$. Note that $B_2 = T(B_1)$, where $T\colon (x,y) \mapsto (2x,2y)$ is the positive linear map that scales both utilities by a factor of 2. Both problems are symmetric in the sense that $(u,v) \in B_i$ if and only if $(v,u) \in B_i$.

Let $\mathcal{U}$ denote the set of all functions f assigning to each B_i a feasible outcome $f(B_i) \in B_i$. We have $|\mathcal{U}| = 4 \cdot 4 = 16$, and we equip $\mathcal{U}$ with the uniform measure $\mu(\{f\}) = 1/16$. We write elements of $\mathcal{U}$ as ordered pairs $(f(B_1), f(B_2))$.

We now define three axioms, corresponding to Nash's efficiency, symmetry, and scale-invariance conditions.

$$\begin{aligned}
A_{\text{eff}} &= \{\, f \in \mathcal{U} \mid f(B_i) \text{ is Pareto-optimal in } B_i \text{ for } i \in \{1,2\} \,\},\\
A_{\text{sym}} &= \{\, f \in \mathcal{U} \mid f(B_i) \text{ lies on the diagonal for } i \in \{1,2\} \,\},\\
A_{\text{inv}} &= \{\, f \in \mathcal{U} \mid f(B_2) = T(f(B_1)) \,\}.
\end{aligned}$$

The Pareto-optimal points in B_1 are $(2,2)$, $(3,0)$, and $(0,3)$; the point $(1,1)$ is dominated by $(2,2)$. Similarly, the Pareto-optimal points in B_2 are $(4,4)$, $(6,0)$, and $(0,6)$, and $(2,2)$ is dominated by $(4,4)$. The diagonal points in B_1 are $(2,2)$ and $(1,1)$; in B_2 they are $(4,4)$ and $(2,2)$. By direct enumeration, $|A_{\text{eff}}| = 3 \cdot 3 = 9$, $|A_{\text{sym}}| = 2 \cdot 2 = 4$, and $|A_{\text{inv}}| = 4$, since the invariance axiom forces $f(B_2)$ to be the image of $f(B_1)$ under T, and each of the four points of B_1 has its T-image in B_2.

Uniqueness and cores. The full intersection $\mathcal{M}(\Sigma) = A_{\text{eff}} \cap A_{\text{sym}} \cap A_{\text{inv}}$ has size 1, namely the function f^* with $f^*(B_1) = (2,2)$ and $f^*(B_2) = (4,4)$. Indeed, efficiency together with symmetry forces $f(B_1) = (2,2)$, since $(2,2)$ is the unique Pareto-optimal point on the diagonal of B_1, and similarly forces $f(B_2) = (4,4)$. The invariance requirement $f(B_2) = T(f(B_1)) = T(2,2) = (4,4)$ is then automatically satisfied. Hence, the system is in the uniqueness regime with model f^*. Table 7 records the model-class sizes and collapse values for all 8 sub-systems, each obtained by direct enumeration.

Sub-system S	$\mu(\mathcal{M}(S))$	$v_\mu(S)$
$\emptyset$	16/16	0
$\{A_{\text{eff}}\}$	9/16	7/16
$\{A_{\text{sym}}\}$	4/16	12/16
$\{A_{\text{inv}}\}$	4/16	12/16
$\{A_{\text{eff}}, A_{\text{sym}}\}$	1/16	15/16
$\{A_{\text{eff}}, A_{\text{inv}}\}$	3/16	13/16
$\{A_{\text{sym}}, A_{\text{inv}}\}$	2/16	14/16
Σ	1/16	15/16

Table 7: Sub-system model-class sizes and collapse values for the miniature Nash bargaining system, derived by direct enumeration on the 16-element domain.

Shapley values. Averaging the marginal contributions over all $3! = 6$ orderings, we obtain the axiom Shapley values with common denominator 48.

$$\phi_{A_{\text{eff}}} = \tfrac{10}{48}, \qquad \phi_{A_{\text{sym}}} = \tfrac{19}{48}, \qquad \phi_{A_{\text{inv}}} = \tfrac{16}{48}.$$

These values sum to $45/48 = 15/16 = v_\mu(\Sigma)$, confirming efficiency. The ranking is given by $\phi_{A_{\text{sym}}} > \phi_{A_{\text{inv}}} > \phi_{A_{\text{eff}}}$.

Interpretation. Symmetry bears the largest share of the credit for singling out f^*. The structural reason is that symmetry, on its own, reduces the domain to 4 of the 16 solutions, and its interaction with efficiency reduces the model class further to a single point. Efficiency, though it eliminates the largest number of individual solutions ($16 - 9 = 7$), contributes less in the presence of symmetry, since much of what efficiency rules out is already ruled out by symmetry. Invariance is intermediate. Note the contrast with the miniature Arrow above, where the impossibility regime distributes blame more evenly and the largest single share is $10/27 \approx 37\%$. In the miniature Nash uniqueness case, symmetry alone accounts for $19/45 \approx 42\%$ of the collapse.

Dominance ranking. By Theorem 8.12(i), the Shapley-value differences satisfy

$$\begin{aligned}\phi_{A_{\text{sym}}} - \phi_{A_{\text{eff}}} &= \tfrac{1}{2}\big[\mu(\mathcal{M}(\{A_{\text{eff}}\})) - \mu(\mathcal{M}(\{A_{\text{sym}}\}))\big] \\ &\quad + \tfrac{1}{2}\big[\mu(\mathcal{M}(\{A_{\text{eff}}, A_{\text{inv}}\})) - \mu(\mathcal{M}(\{A_{\text{sym}}, A_{\text{inv}}\}))\big] \\ &= \tfrac{1}{2}(\tfrac{9}{16} - \tfrac{4}{16}) + \tfrac{1}{2}(\tfrac{3}{16} - \tfrac{2}{16}) = \tfrac{5}{32} + \tfrac{1}{32} = \tfrac{6}{32} = \tfrac{9}{48}.\end{aligned}$$

In particular, $A_{\text{sym}} \succeq_\mu A_{\text{eff}}$ (symmetry is uniformly more restrictive than efficiency on this domain), so by Theorem 8.12(ii) the ranking $\phi_{A_{\text{sym}}} > \phi_{A_{\text{eff}}}$ is robust across every measure under which this dominance holds.

12.4 A Dimension-4 Impossibility Core

The examples above have impossibility dimension 3. We now exhibit a system in which the impossibility dimension is 4, verifying the Dimension-Witness Duality (Theorem 10.6) at the extreme end of its range on a small domain.

Let $\mathcal{U} = \{u_1, u_2, u_3, u_4\}$ with uniform measure $\mu(\{u_i\}) = 1/4$, and $\Sigma = \{A_1, A_2, A_3, A_4\}$ with $A_i = \mathcal{U} \setminus \{u_i\}$ for each i. The intersection satisfies

$$\mathcal{M}(\Sigma) = \bigcap_{i=1}^{4} (\mathcal{U} \setminus \{u_i\}) = \mathcal{U} \setminus \{u_1, u_2, u_3, u_4\} = \emptyset,$$

so the system is in the impossibility regime. For each proper sub-system $\Sigma \setminus \{A_i\}$, the intersection is $\{u_i\}$, so every proper sub-system is possible. Hence, Σ is itself an impossibility core with $\dim(\Sigma) = 4$.

The independence witnesses are $w_i = u_i$ for each i, and they are pairwise distinct as required by Theorem 10.6(i). Moreover, no impossible system of dimension 4 can exist on a domain of size less than 4, since the four witnesses must be distinct elements of the domain. This confirms the Dimension-Witness Duality (Theorem 10.6(iii)) at $k = 4$.

By symmetry, all four axioms play equivalent roles. Each has $\mu(\mathcal{M}(\{A_i\})) = 3/4$ and $v_\mu(\{A_i\}) = 1/4$. For any pair $\{A_i, A_j\}$, we have $\mu(\mathcal{M}(\{A_i, A_j\})) = 2/4$; for any triple, $\mu(\mathcal{M}) = 1/4$; and $\mu(\mathcal{M}(\Sigma)) = 0$. The axiom Shapley values are $\phi_{A_i} = 1/4$ for each i, distributing the total collapse of 1 equally, and the axiom entropy attains its maximum $H(\Sigma, \mu) = \ln 4 \approx 1.39$ nats. By Theorem 11.3(ii), the feasibility complex $K(\Sigma) = \partial\Delta^3 \cong S^2$ is homeomorphic to the 2-sphere, and its Euler characteristic is $\chi = 4 - 6 + 4 = 2 = \chi(S^2)$.

12.5 Multiple Cores and Nontrivial Higher Homology

We now construct a fully derived example with two disjoint impossibility cores of different sizes, in which the feasibility complex has nontrivial second Betti number. This verifies the Alexander Duality Theorem (Theorem 11.11) in a regime where neither of the easy sufficient conditions of Corollary 11.12 applies, since no axiom is a cone point and the impossibility depth is greater than one.

Let $\mathcal{U} = \{u_1, u_1', u_2, u_2', u_3, u_3'\}$, a domain of six elements, equipped with the uniform measure $\mu(\{u\}) = 1/6$. Consider the axiom system $\Sigma = \{A_1, A_2, A_3, A_4, A_5\}$ defined by

$$A_1 = \{u_2, u_2', u_3, u_3'\}, \qquad A_2 = \{u_1, u_1', u_3, u_3'\}, \qquad A_3 = \{u_1, u_1', u_2, u_2'\},$$
$$A_4 = \{u_1', u_2', u_3'\}, \qquad A_5 = \{u_1, u_2, u_3\}.$$

The first three axioms partition-exclude one primed/unprimed pair each (that is, it holds that $A_i = \mathcal{U} \setminus \{u_i, u_i'\}$ for $i \in \{1, 2, 3\}$). The last two axioms partition $\mathcal{U}$ by prime status, with A_4 containing the primed elements and A_5 containing the unprimed elements.

Cores. We have $A_1 \cap A_2 \cap A_3 = \emptyset$, since every element of $\mathcal{U}$ lies in exactly one of the three pairs $\{u_i, u_i'\}$ and is therefore excluded by A_i. For pairs from $\{A_1, A_2, A_3\}$: $A_1 \cap A_2 = \{u_3, u_3'\}$, and analogously for the other two pairs, each nonempty. Hence, $C_1 = \{A_1, A_2, A_3\}$ is an impossibility core of size 3. Similarly, $A_4 \cap A_5 = \emptyset$, since primed and unprimed elements are disjoint, and each of A_4 and A_5 is nonempty; so $C_2 = \{A_4, A_5\}$ is an impossibility core of size 2. The two cores are disjoint.

We verify that these are the only cores. For $i \in \{1, 2, 3\}$ and $j \in \{4, 5\}$, the intersection $A_i \cap A_j$ contains two elements (one primed and one unprimed for each of the two "surviving" pairs), so no pair with mixed indices is impossible. For triples of the form $\{A_i, A_j, A_k\}$ with $i, j \in \{1, 2, 3\}$ and $k \in \{4, 5\}$, the intersection is $\{u_l, u_l'\} \cap A_k$ (where l is the unique index in $\{1, 2, 3\} \setminus \{i, j\}$), which is $\{u_l'\}$ if $k = 4$ and $\{u_l\}$ if $k = 5$, both nonempty. For triples of the form $\{A_i, A_j, A_k\}$ with $i \in \{1, 2, 3\}$ and $j, k = 4, 5$, the intersection is $A_i \cap A_4 \cap A_5 = A_i \cap \emptyset = \emptyset$, which is impossible but contains the smaller core $\{A_4, A_5\}$ and is therefore not itself minimal. Hence, $\mathcal{C}(\Sigma) = \{C_1, C_2\}$ and $\dim(\Sigma) = 2$.

Impossibility depth. Since $C_1 \cap C_2 = \emptyset$, no single axiom hits both cores, so the transversal number is $\tau(\mathcal{C}(\Sigma)) = 2$, and by Theorem 9.8 the impossibility depth is $d(\Sigma) = 2$. For example, dropping A_1 and A_4 leaves $\{A_2, A_3, A_5\}$, with $\mathcal{M} = A_2 \cap A_3 \cap A_5 = \{u_1\}$, nonempty.

The feasibility complex. The maximal faces of the feasibility complex $K(\Sigma)$ are the six triples

$$\{A_1, A_2, A_4\},\ \{A_1, A_2, A_5\},\ \{A_1, A_3, A_4\},\ \{A_1, A_3, A_5\},\ \{A_2, A_3, A_4\},\ \{A_2, A_3, A_5\},$$

obtained by taking exactly two axioms from $\{A_1, A_2, A_3\}$ (avoiding the core C_1) together with exactly one axiom from $\{A_4, A_5\}$ (avoiding the core C_2). This is the boundary complex of the bipyramid over the triangle $\{A_1, A_2, A_3\}$ with poles A_4 and A_5, first identified in Example 11.14. Its f-vector is $(5, 9, 6)$ and its Euler characteristic is $\chi = 5 - 9 + 6 = 2$, matching $\chi(S^2)$.

Alexander duality. Each of the five axioms belongs to at least one core, so no axiom is a cone point and Corollary 11.12(ii) does not apply. The co-core complex $K^\vee(\Sigma)$ has maximal faces $\Sigma \setminus C_1 = \{A_4, A_5\}$ and $\Sigma \setminus C_2 = \{A_1, A_2, A_3\}$; since these two sets are disjoint, $K^\vee(\Sigma)$ is disconnected with $\widetilde{\beta}_0(K^\vee) = 1$. By Theorem 11.11 with $n = 5$ and $i = 2$, we obtain

$$\widetilde{\beta}_2\big(K(\Sigma)\big) \;=\; \widetilde{\beta}_{n-i-3}\big(K^\vee(\Sigma)\big) \;=\; \widetilde{\beta}_0\big(K^\vee(\Sigma)\big) \;=\; 1,$$

so $K(\Sigma)$ has one nontrivial second homology class over any field, consistent with the homotopy equivalence $K(\Sigma) \simeq S^2$. This example verifies the Alexander Duality Theorem in a regime where neither easy sufficient condition applies, and demonstrates that nontrivial higher Betti numbers do arise in the theory.

Shapley values. The axiom Shapley values, computed by averaging marginal contributions over all $5! = 120$ orderings, are

$$\phi_{A_1} = \phi_{A_2} = \phi_{A_3} = \tfrac{1}{6}, \qquad \phi_{A_4} = \phi_{A_5} = \tfrac{1}{4}.$$

By symmetry, the three cycle axioms A_1, A_2, A_3 receive equal blame, and the two pole axioms A_4, A_5 receive equal blame. Indeed, the two axioms forming the smaller core each carry 25% of the collapse burden, while the three axioms of the larger core each carry 16.7%; the smaller core concentrates blame on its members. The total is $3 \cdot \frac{1}{6} + 2 \cdot \frac{1}{4} = \frac{1}{2} + \frac{1}{2} = 1 = v_\mu(\Sigma)$, confirming efficiency.

12.6 Three Pairwise-Disjoint Cores and the Nerve Theorem

We now construct a six-axiom system with three pairwise-disjoint impossibility cores, in which the co-core complex has a nontrivial 1-cycle and the Alexander duality computation proceeds through the nerve theorem. This example demonstrates a structurally different mechanism from the two-core case of Section 12.5: there, the co-core complex was disconnected ($\widetilde{\beta}_0 = 1$) and Alexander duality mapped this to $\widetilde{\beta}_2(K) = 1$ via the $\widetilde{H}_0$-to-$\widetilde{H}_2$ shift; here, the co-core complex has a loop ($\widetilde{\beta}_1 = 1$), and Alexander duality maps $\widetilde{H}_1$ of the co-core complex to $\widetilde{H}_2$ of the feasibility complex. The final Betti number is the same ($\beta_2 = 1$), but the topological mechanism producing it is entirely different.

Let $\mathcal{U} = \{w_1, w_2, w_3, w_4, w_5, w_6, w_7, w_8\}$, a domain of eight elements, equipped with the uniform measure $\mu(\{w_j\}) = 1/8$. We define six axioms by

$$\begin{aligned} A_1 &= \{w_1, w_2, w_3, w_4\}, & A_2 &= \{w_5, w_6, w_7, w_8\}, \\ A_3 &= \{w_1, w_2, w_5, w_6\}, & A_4 &= \{w_3, w_4, w_7, w_8\}, \\ A_5 &= \{w_1, w_3, w_5, w_7\}, & A_6 &= \{w_2, w_4, w_6, w_8\}. \end{aligned}$$

Observe that $A_1 \cap A_2 = \emptyset$, $A_3 \cap A_4 = \emptyset$, and $A_5 \cap A_6 = \emptyset$, so $C_1 = \{A_1, A_2\}$, $C_2 = \{A_3, A_4\}$, and $C_3 = \{A_5, A_6\}$ are three impossibility cores of size 2. These cores are pairwise disjoint. We verify that no other pair is impossible. In fact, every non-core pair has intersection of size exactly 2. For instance, $A_1 \cap A_3 = \{w_1, w_2\}$, $A_1 \cap A_4 = \{w_3, w_4\}$, $A_1 \cap A_5 = \{w_1, w_3\}$, and so on; the reader may check the remaining nine non-core pairs directly from the definitions.

No hidden cores. A triple $\{A_i, A_j, A_k\}$ is impossible if and only if it contains a core pair. Indeed, each of the eight triples that select one axiom from each pair $\{A_1, A_2\}$, $\{A_3, A_4\}$, $\{A_5, A_6\}$ has intersection of size exactly 1 (a unique "witness" element), hence is possible. Every other triple contains at least one core pair and is therefore impossible. Every subset of Σ of size 4 or more contains at least one core pair (by the pigeonhole principle, since any four axioms drawn from three pairs must include both members of at least one pair), so $v_\mu(S) = 1$ for all $|S| \geq 4$. Hence, $\mathcal{C}(\Sigma) = \{C_1, C_2, C_3\}$ and $\dim(\Sigma) = 2$.

Impossibility depth. Since $C_1 \cap C_2 = C_1 \cap C_3 = C_2 \cap C_3 = \emptyset$, no single axiom hits all three cores, and no pair of axioms does either unless it contains one member of each of two cores. Indeed, the minimum transversal must contain at least one axiom from each core, so $\tau(\mathcal{C}(\Sigma)) = 3$. By Theorem 9.8, $d(\Sigma) = 3$. A minimum rescue requires dropping three axioms (one from each core); for instance, dropping A_2, A_4, and A_6 leaves $\{A_1, A_3, A_5\}$ with $\mathcal{M} = \{w_1\} \neq \emptyset$.

The feasibility complex. The feasibility complex $K(\Sigma)$ has $\Sigma = \{A_1, A_2, A_3, A_4, A_5, A_6\}$ as vertex set, edges consisting of all $\binom{6}{2} - 3 = 12$ non-core pairs, and triangles consisting of the 8 triples with one member from each pair. No face of dimension 3 or higher exists, since every quadruple contains a core. The f-vector is $(6, 12, 8)$ and the Euler characteristic is $\chi = 6 - 12 + 8 = 2$.

Alexander duality via the nerve theorem. The co-core complex $K^\vee(\Sigma)$ has three maximal faces, one per core:

$$\Sigma \setminus C_1 = \{A_3, A_4, A_5, A_6\}, \qquad \Sigma \setminus C_2 = \{A_1, A_2, A_5, A_6\}, \qquad \Sigma \setminus C_3 = \{A_1, A_2, A_3, A_4\}.$$

Each maximal face is a 3-simplex (tetrahedron), hence contractible. The pairwise intersections are

$$(\Sigma \setminus C_1) \cap (\Sigma \setminus C_2) = \{A_5, A_6\}, \qquad (\Sigma \setminus C_1) \cap (\Sigma \setminus C_3) = \{A_3, A_4\},$$
$$(\Sigma \setminus C_2) \cap (\Sigma \setminus C_3) = \{A_1, A_2\},$$

each a single edge (a 1-simplex, hence contractible). The triple intersection is equal to $(\Sigma \setminus C_1) \cap (\Sigma \setminus C_2) \cap (\Sigma \setminus C_3) = \Sigma \setminus (C_1 \cup C_2 \cup C_3) = \emptyset$, since $C_1 \cup C_2 \cup C_3 = \Sigma$. Since the three contractible pieces and their three contractible pairwise intersections form a cover with empty triple intersection, the *nerve theorem* (see, e.g., Kozlov 2008, Chapter 15) identifies the homotopy type of $K^\vee(\Sigma)$ with that of the nerve of the cover. The nerve has three vertices (one per maximal face) and three edges (one per nonempty pairwise intersection), with no triangle (the triple intersection is empty). This nerve is the boundary of a triangle, which is homeomorphic to S^1. Hence, $K^\vee(\Sigma) \simeq S^1$, and $\widetilde{H}_1(K^\vee(\Sigma); \beth) \cong \beth$.

Applying Theorem 11.11 with $n = 6$ gives $\widetilde{H}_i(K(\Sigma)) \cong \widetilde{H}_{3-i}(K^\vee(\Sigma))$. In particular,

$$\widetilde{H}_2\big(K(\Sigma); \beth\big) \;\cong\; \widetilde{H}_1\big(K^\vee(\Sigma); \beth\big) \;\cong\; \beth,$$

and all other reduced homology groups of $K(\Sigma)$ vanish (since $\widetilde{H}_j(K^\vee) = 0$ for $j \neq 1$). We hence have $K(\Sigma) \simeq S^2$, confirming $\chi = 2 = \chi(S^2)$.

Comparison with the two-core case. Both the present example and the two-core example of Section 12.5 yield $\widetilde{\beta}_2(K) = 1$ and $K \simeq S^2$, but the topological mechanisms are different. In the two-core case ($n = 5$, cores of sizes 3 and 2), the co-core complex was disconnected ($\widetilde{\beta}_0(K^\vee) = 1$), and Alexander duality shifted dimensions by $n - 3 = 2$, mapping $\widetilde{H}_0(K^\vee)$ to $\widetilde{H}_2(K)$. In the present three-core case ($n = 6$, three cores of size 2), the co-core complex has a loop ($\widetilde{\beta}_1(K^\vee) = 1$), and Alexander duality shifts dimensions by $n - 3 = 3$, mapping $\widetilde{H}_1(K^\vee)$ to $\widetilde{H}_2(K)$. The same Betti number $\beta_2 = 1$ in the feasibility complex arises from topologically distinct features of the co-core complex (disconnectedness versus a cycle), illustrating that Alexander duality provides a finer invariant than the Betti number alone.

Shapley values. By the symmetry of the construction (the symmetric group S_3 acts on the three pairs, and $\mathbb{Z}/2$ acts within each pair), all six axioms play equivalent roles. Each has $v_\mu(\{A_i\}) = 1/2$, and by symmetry $\phi_{A_i} = 1/6$ for all i, with total $6 \cdot 1/6 = 1$. The axiom entropy attains its maximum $H(\Sigma, \mu) = \ln 6 \approx 1.79$ nats, the highest among all examples in this catalogue.

13 Heuristic Extrapolations to Classical Theorems on Infinite Domains

The examples of Section 12 realize the framework on small finite domains, where every quantity is derived by direct enumeration. The classical theorems of social choice, bargaining, and cooperative game theory typically involve infinite domains (all social welfare functions on a fixed set of alternatives, all bargaining problems on the plane, all cooperative games on a fixed player set), and the framework does not, at present, apply to them as theorems. This section discusses what the framework nevertheless suggests about these classical results, at the level of heuristic parallel and structural hypothesis. The reader is asked to treat this section as commentary, not as a further theoretical development. None of the claims below are theorems of this paper; the framework's actual theorems are those of Sections 4–14, illustrated on the derived examples of Section 12.

13.1 The Arrow Theorem

Arrow's theorem (Arrow, 1951) states that on the domain of all social welfare functions from a set of at least three alternatives and at least two voters to a strict total order, no function satisfies unrestricted domain, Pareto, independence of irrelevant alternatives, and non-dictatorship simultaneously. The domain is infinite (indeed, its cardinality grows factorially with the number of alternatives and exponentially with the number of voters), so the framework of Sections 2–14 does not directly apply. Two heuristic parallels are nevertheless worth stating.

First, the miniature construction of Section 12.2 captures the qualitative pattern of Arrow's theorem, namely the incompatibility of Pareto with a structural symmetry axiom and non-dictatorship. In the miniature, the symmetry axiom is neutrality with respect to a single alternative-swap; in the classical theorem, it is independence of irrelevant alternatives, a much stronger and more restrictive condition. The miniature shows that the pattern is realizable in a small finite domain with all quantitative invariants derivable, which suggests that a corresponding pattern on the infinite Arrow domain is not merely a curiosity of the infinite setting but reflects a structural incompatibility that persists in miniature.

Second, if a canonical probability measure μ^* on the infinite domain of social welfare functions could be identified (for instance, by a limiting procedure on finite restrictions), the framework's Axiom Dominance Theorem (Theorem 8.12) would then apply and provide a canonical attribution. The standard qualitative intuition is that independence of irrelevant alternatives is the most individually restrictive of Arrow's axioms, non-dictatorship the least, and Pareto intermediate; if the dominance chain $A_{\mathrm{IIA}} \succeq_{\mu^*} A_{\mathrm{P}} \succeq_{\mu^*} A_{\mathrm{ND}}$ holds under μ^*, then $\phi_{A_{\mathrm{IIA}}} \geq \phi_{A_{\mathrm{P}}} \geq \phi_{A_{\mathrm{ND}}}$ follows. The identification of μ^* is an open problem; see Section 15 for further discussion.

13.2 Nash's Bargaining Theorem

Nash's bargaining theorem (Nash, 1950) states that on the domain of two-player bargaining problems on the plane, the unique solution satisfying efficiency, symmetry, scale invariance, and independence of irrelevant alternatives is the Nash bargaining solution, which maximizes

the product of utilities above the disagreement point. The miniature construction of Section 12.3 captures the joint force of efficiency, symmetry, and scale invariance in a 16-element domain, where symmetry emerges as the dominant contributor to the collapse. The classical Nash setting adds the fourth axiom of independence of irrelevant alternatives; on the infinite domain, all four are needed for uniqueness.

The miniature does not include an IIA axiom, since with only two bargaining problems related by scaling there is no proper subset relation between the problems and IIA is vacuous. A four-axiom miniature Nash would require a domain with at least one strict inclusion between problems, and the enumeration grows accordingly. The structural hypothesis suggested by the framework is that on the classical Nash domain, the four Shapley values would be positive and their ranking would reflect the relative restrictiveness of the axioms in the presence of the others; empirical calibration would depend on the choice of measure and is left open.

13.3 The Shapley Value Axiomatization

The Shapley value axiomatization (Shapley, 1953) states that on the space of all value operators on n-player cooperative games, the unique operator satisfying efficiency, symmetry, linearity, and the dummy axiom is the Shapley value. Applying our framework to this axiomatization is formally self-referential. The outer Shapley value is the axiom Shapley value of the axiom system for the classical Shapley value; the inner Shapley value is the object being axiomatized. The self-reference is harmless in the sense that the outer and inner Shapley values operate on different games, namely a game on the axiom set of the classical axiomatization for the outer, and a game on the classical player set for the inner.

The classical Shapley domain is infinite (it is the space of all value operators, itself a function space), and a derived treatment on that domain is not available here. On any finite restriction (for instance, restricting to n-player games with integer values in a bounded range), the framework applies without stipulation and yields a canonical attribution of blame to the four Shapley axioms; the resulting numerical values are dependent on the finite restriction chosen. We do not compute a specific instance here.

13.4 Summary

The three parallels above are not theorems of this paper. Rather, they are structural hypotheses about how the framework, if canonical measures on the relevant infinite domains could be identified, would specialize to the classical theorems. The miniature constructions of Section 12 realize the framework without stipulation and demonstrate that the qualitative patterns of the classical theorems (impossibility, uniqueness, and self-reference) are in fact already present in small finite settings and are amenable to complete quantitative analysis. The extension to the classical infinite domains is an open direction, further discussed in Section 15.

14 Phase Transitions in Random Axiom Systems

The previous sections studied a fixed axiom system. We now introduce probability into the axiom set itself and prove that the transition from possibility to impossibility is generically abrupt, exhibiting a *sharp threshold* phenomenon.

Let $\Sigma = \{A_1, A_2, \ldots, A_n\}$ be a collection of n axioms on a fixed domain $\mathcal{U}$ with measure μ. We say that the compatibility hypergraph $H(\Sigma)$ is *vertex-transitive* if, for every pair of axioms $A_i, A_j \in \Sigma$, there exists a permutation π of Σ with $\pi(A_i) = A_j$ that preserves the hyperedge structure, i.e., a subset $S \subseteq \Sigma$ is a hyperedge of $H(\Sigma)$ if and only if $\pi(S)$ is. Intuitively, vertex-transitivity means that every axiom plays a structurally identical role in the impossibility structure, so that the system looks the same from the perspective of any individual axiom. For $p \in [0, 1]$, let Σ_p denote the random sub-system obtained by including each A_i independently with probability p. Define the *impossibility probability*

$$\theta_n(p) = \Pr[\mathcal{M}(\Sigma_p) = \emptyset].$$

Since $\mathcal{M}(\Sigma_p) = \emptyset$ is a monotone event (it holds for Σ_p whenever it holds for a sub-system), $\theta_n(p)$ is a non-decreasing function of p.

Definition 14.1 (Threshold)**.** A function $p^*(n)$ is a *threshold* for impossibility if the following two conditions hold.

$$\theta_n(p) \to 0 \quad \text{as } n \to \infty \text{ whenever } p = o(p^*(n)),$$

$$\theta_n(p) \to 1 \quad \text{as } n \to \infty \text{ whenever } p = \omega(p^*(n)).$$

The threshold is *sharp* if θ_n transitions from 0 to 1 over an interval of length $o(p^*(n))$ centered at $p^*(n)$.

Remark 14.2 (Modeling Choices in the Phase Transition Setup)**.** Three modeling choices in the phase transition analysis deserve explicit comment. First, we use the product measure on axiom inclusion (each axiom included independently with probability p), which is the standard setup in random graph and random hypergraph theory. Alternative models (e.g., fixing the number of included axioms to $\lfloor pn \rfloor$ exactly, or using a non-uniform inclusion distribution) would give the same threshold location up to constants, by standard equivalence results for monotone properties. Second, the asymptotic regime $n \to \infty$ requires the domain size N to grow with n (since at most $2^N - 1$ distinct nonempty axioms exist on a domain of size N), and the hypotheses of Theorem 14.3(ii)–(iii) implicitly constrain the relationship between N and n. Third, the threshold location $p^* = \Theta(M_k^{-1/k})$ depends on the count M_k of minimum-size cores, which in turn depends on the specific axiom system; different systems with the same dimension k but different core counts M_k have different thresholds. The hypotheses (a)–(d) of Theorem 14.3(ii) are sufficient conditions for the first- and second-moment methods to succeed; they are not claimed to be necessary, and the theorem is silent about systems that violate them.

Theorem 14.3 (Phase Transition for Impossibility)**.** *Let $\Sigma = \{A_1, A_2, \ldots, A_n\}$ be a collection of n distinct axioms on a finite domain $\mathcal{U}$ with $|\mathcal{U}| = N$. In the asymptotic statements below, we consider a sequence of such systems with $n \to \infty$. The domain size N may grow with n, and must do so for $n > 2^N - 1$, since at most $2^N - 1$ distinct nonempty axioms exist on a domain of size N.*

(i) Existence of a threshold. *The impossibility probability $\theta_n(p)$ has a threshold $p^*(n)$. Indeed, the event $\{\mathcal{M}(\Sigma_p) = \emptyset\}$ is a monotone Boolean function of the inclusion indicators $(X_1, X_2, \ldots, X_n) \in \{0,1\}^n$, and every monotone Boolean function has a threshold (Bollobás and Thomason, 1987).*

(ii) Threshold location. *Suppose the compatibility hypergraph $H(\Sigma)$ satisfies the following conditions, each of which plays a specific role in the proof.*

- *(a) The minimum hyperedge size is $k = \dim(\Sigma)$, with exactly M_k impossibility cores of size k. This identifies the dominant obstruction; the smallest cores are the ones most likely to appear first, and M_k counts how many of them exist.*
- *(b) Any two distinct cores of size k share at most $k-2$ axioms (a* pairwise overlap bound*). This controls the correlations between different cores in the second-moment calculation. Without such a bound, two cores sharing $k-1$ axioms would be almost perfectly correlated, and the second-moment method would fail to produce a matching lower bound.*
- *(c) The compatibility hypergraph $H(\Sigma)$ is vertex-transitive, the number of minimum cores satisfies $M_k = \Theta(n^{k-1})$, and furthermore for every $j \in \{1, 2, \ldots, k-1\}$, the maximum number of size-k cores containing any fixed j-element subset of Σ satisfies $d_j = O(n^{k-1-j})$. Vertex-transitivity ensures that each axiom plays a symmetric role (no axiom is disproportionately involved in cores), and the scaling conditions ensure that the second-moment ratio converges to 1. The bound on d_j is a regularity condition on the core family; it prevents pathological clustering of cores around particular subsets.*
- *(d) The larger cores are subcritical at the threshold, that is, writing M_s for the number of impossibility cores of size s, we have $\sum_{s>k} M_s (p^*(n))^s = o(1)$, where $p^*(n) = M_k^{-1/k}$. This ensures that impossibility at threshold is driven by size-k cores, not by an accumulation of larger impossible sub-systems.*

Then the threshold satisfies $p^*(n) = \Theta(M_k^{-1/k}) = \Theta(n^{-(k-1)/k})$. The first-moment bound (with hypothesis (d)) gives $\theta_n(p) \to 0$ for $p = o(p^*(n))$, and the converse follows from the second-moment method under the overlap and scaling conditions (a)–(c).

(iii) Sharp threshold. *If additionally $H(\Sigma)$ is vertex-transitive and every variable is contained in at least one core of size k (so that no bounded set of variables determines the impossibility event), then the threshold is sharp. This follows from the Friedgut–Kalai sharp threshold theorem (Friedgut and Kalai, 1996), since vertex-transitivity and the condition that every variable participates in a core together imply that the impossibility event cannot be well-approximated by a junta.*

(iv) The critical fraction. *At threshold, the expected number of impossible sub-systems of minimum size is $\Theta(1)$, i.e., a constant expected number of impossibility cores of size k appear exactly at the critical probability. Below the threshold, no impossibility core is likely present; above it, many are.*

Proof. For (i), the event $\{\mathcal{M}(\Sigma_p) = \emptyset\}$ depends on $(X_1, X_2, \ldots, X_n)$, where $X_i = \mathbf{1}[A_i \in \Sigma_p]$. It is monotone. Indeed, if $\mathcal{M}(\Sigma_p) = \emptyset$ and $\Sigma_p \subseteq \Sigma_q$ (i.e., $X_i \le Y_i$ componentwise), then

$\mathcal{M}(\Sigma_q) = \emptyset$ by Theorem 4.1(i). Every monotone Boolean function of n independent Bernoulli variables has a threshold in the sense of Definition 14.1; this is a classical result of Bollobás and Thomason (1987).

For (ii), we begin with the first-moment upper bound. Since $\mathcal{M}(\Sigma_p) = \emptyset$ if and only if some impossibility core (of any size) is contained in Σ_p, the union bound gives

$$\theta_n(p) \;\le\; \sum_{s \ge k} M_s\, p^s \;=\; M_k p^k \;+\; \sum_{s > k} M_s\, p^s.$$

The expected number of impossibility cores of size k present in Σ_p is

$$\mathbb{E}\big[\text{number of size-}k\text{ cores in }\Sigma_p\big] = M_k \cdot p^k.$$

Setting $M_k p^k = \Theta(1)$ gives $p^* = \Theta(M_k^{-1/k})$. Under the scaling $M_k = \Theta(n^{k-1})$, this becomes $p^* = \Theta(n^{-(k-1)/k})$. For $p = o(p^*)$, the leading term satisfies $M_k p^k = o(M_k (p^*)^k) = o(1)$; and since $M_s p^s$ is monotone increasing in p (for $p \in [0,1]$ and $M_s \ge 0$), the tail satisfies

$$\sum_{s > k} M_s\, p^s \;\le\; \sum_{s > k} M_s (p^*)^s \;=\; o(1)$$

by hypothesis (d). Hence, $\theta_n(p) \to 0$ for $p = o(p^*)$, finishing the first-moment direction.

For the converse direction ($\theta_n(p) \to 1$ for $p = \omega(p^*)$), it suffices to show that the event $\{Z \ge 1\}$ occurs with probability tending to 1, since $\{Z \ge 1\} \subseteq \{\mathcal{M}(\Sigma_p) = \emptyset\}$. We apply the second-moment method to $Z = \sum_C \mathbf{1}[C \subseteq \Sigma_p]$, the number of size-$k$ cores contained in Σ_p, where the sum runs over the M_k size-k cores. We have $\mathbb{E}[Z] = M_k p^k$. For the second moment, $\mathbb{E}[Z^2] = \sum_{C,C'} p^{|C \cup C'|}$. The diagonal terms contribute $M_k p^k = \mathbb{E}[Z]$. For off-diagonal terms, the pairwise overlap bound (b) ensures that any two distinct cores C, C' share at most $j \le k-2$ axioms, giving $|C \cup C'| = 2k - j \ge k+2$. For a fixed core C and a fixed j-element subset of C, let d_j denote the maximum number of size-k cores containing that subset. By hypothesis (c), $d_j = O(n^{k-1-j})$ for every $j \in \{1, 2, \ldots, k-1\}$. Note that the case $j = 1$ is a direct consequence of vertex-transitivity, since each of the $\Theta(n)$ axioms participates in exactly $M_k \cdot k/n = \Theta(n^{k-2})$ cores by double counting. The analogous bound for $j \ge 2$ is imposed as an additional regularity condition in (c), because vertex-transitivity alone controls average degrees over j-subsets but not their maxima. Grouping off-diagonal ordered pairs by overlap size $j = |C \cap C'|$ gives $\mathbb{E}[Z^2] \le \mathbb{E}[Z] + \sum_{j=0}^{k-2} M_k \binom{k}{j} d_j\, p^{2k-j}$. For pairs with overlap $j = 0$ (disjoint cores), the corresponding indicator random variables depend on disjoint sets of Bernoulli inputs and are therefore independent, so their covariance vanishes and they contribute nothing to $\mathrm{Var}(Z) = \mathbb{E}[Z^2] - \mathbb{E}[Z]^2$. For the remaining terms with overlap $j \ge 1$, dividing by $\mathbb{E}[Z]^2 = M_k^2 p^{2k}$ gives contribution $\binom{k}{j} d_j\, p^{-j} / M_k = O(n^{-j})\, p^{-j}$. For $p = \omega(p^*)$ with $p^* = \Theta(n^{-(k-1)/k})$, we have $p^{-j} = o(n^{j(k-1)/k})$, so each term is $o(n^{j(k-1)/k - j})$; note that $j(k-1)/k - j = -j/k$, so this is $o(n^{-j/k}) \to 0$. Together with the diagonal contribution $\mathbb{E}[Z]/\mathbb{E}[Z]^2 = 1/\mathbb{E}[Z] \to 0$, we obtain $\mathbb{E}[Z^2]/\mathbb{E}[Z]^2 \to 1$. The Paley–Zygmund inequality then gives $\Pr[Z > 0] \ge \mathbb{E}[Z]^2/\mathbb{E}[Z^2] \to 1$, so $\theta_n(p) \ge \Pr[Z > 0] \to 1$.

For (iii), we apply the Friedgut–Kalai theorem (Friedgut and Kalai, 1996), which asserts the following. Let $f\colon \{0,1\}^n \to \{0,1\}$ be a monotone Boolean function, and let $I_i(f) = \Pr[f(X) \ne f(X^{\oplus i})]$ denote the influence of the i-th variable, where $X^{\oplus i}$ differs

from X only in coordinate i. If the threshold of f is not sharp, then there exists a variable i with $I_i(f) \geq c/\log n$ for some absolute constant $c > 0$, i.e., the function is well-approximated by a junta (a function depending on a bounded number of variables). In our setting, $f(X_1, X_2, \ldots, X_n) = \mathbf{1}[\mathcal{M}(\Sigma_p) = \emptyset]$. Vertex-transitivity of $H(\Sigma)$ implies that the automorphism group of the impossibility event acts transitively on the n coordinates, so all influences are equal: $I_i(f) = I_j(f)$ for all i, j. The condition that every variable belongs to at least one core of size k ensures that $I_i(f) > 0$ for every i (each variable participates in some minimal impossible sub-system and therefore has positive influence). Since all influences are equal and positive, their common value is $I(f)/n$, where $I(f) = \sum_i I_i(f)$ is the total influence. By Margulis's lemma (or the Russo formula), $I(f) = \theta_n'(p)$, the derivative of the impossibility probability. If the threshold were not sharp, the junta condition would require $I(f)/n \geq c/\log n$, giving $I(f) \geq cn/\log n$. However, vertex-transitivity and the equal-influence property force $I(f)/n = I_i(f)$ for each i, and the Friedgut–Kalai theorem then implies that the transition width is $O(p^*/\log n)$, which is $o(p^*)$. Hence, the threshold is sharp.

For (iv), this follows directly from the definition of p^* and the first-moment calculation in (ii). Below the threshold, the expected count $M_k p^k \to 0$, and hypothesis (d) controls the tail from larger cores, so the probability of any core existing vanishes. Above the threshold, $M_k p^k \to \infty$, and the second-moment argument gives $\Pr[Z > 0] \to 1$. □

Remark 14.4 (When (d) Holds Automatically)**.** Hypothesis (d) of Theorem 14.3(ii) controls the contribution of impossibility cores of size larger than k and is needed to identify the appearance threshold of size-k cores with the impossibility threshold. It holds automatically in two natural settings. First, if every impossibility core of Σ has the same size k, then $M_s = 0$ for all $s > k$, and the tail sum is identically zero. Second, if the counts M_s satisfy $M_s = O(n^{s-1})$ for every $s > k$, then a direct calculation gives $M_s(p^*(n))^s = O(n^{-1/k})$ for $s = k + 1$, and the bound decreases geometrically in s, so the tail sum is $o(1)$ whenever $k \geq 2$. More generally, hypothesis (d) is a mild subcriticality condition that fails only when the larger cores are pathologically numerous relative to the leading rate.

Remark 14.5 (Interpretation: The Knife-Edge of Possibility)**.** Theorem 14.3 says that for large axiom systems, there is no gradual transition from possibility to impossibility. As axioms are added one by one (or as the inclusion probability p increases), the system lives in the possibility regime for a long time, and then at the critical fraction p^* collapses abruptly into impossibility. This is the axiomatic analogue of a phase transition in statistical physics. The system is "possible" almost surely below p^* and "impossible" almost surely above it, with the transition occurring over a vanishingly narrow window.

The dependence of the threshold on the impossibility dimension is striking. Systems with higher-dimensional impossibility (requiring more axioms to generate a contradiction) undergo their phase transition at a lower p^* when the number of cores grows polynomially in n. Intuitively, if you need many axioms to create a contradiction, each individual axiom contributes less pressure, and the critical point is reached sooner (at a lower inclusion probability).

Observe that $p^*(n) \to 0$ as $n \to \infty$ for any fixed $k \geq 2$ and polynomially growing M_k. This has a noteworthy implication. In a large universe of axioms, *any fixed positive inclusion*

probability eventually makes the system impossible almost surely. There is no safe positive density; possibility is fragile in large systems.

Example 14.6 (Running Example: Phase Transition)**.** Returning to Example 3.1, the only impossibility core is $\{A_1, A_2\}$, so $\dim(\Sigma) = k = 2$ and $M_k = 1$. The random sub-system Σ_p includes each of the four axioms independently with probability p. The system is impossible if and only if both A_1 and A_2 are included, which occurs with probability p^2. Here, the "threshold" is $p^* = \Theta(M_k^{-1/k}) = \Theta(1)$, consistent with Theorem 14.3(ii). The transition is not sharp, as the system is too small for the asymptotics to take effect, but the structure is visible. Below $p \approx 0$, impossibility is unlikely; near $p = 1$, it is almost certain.

Example 14.7 (The Canonical Family)**.** To verify that the hypotheses of Theorem 14.3 are satisfiable in a genuinely asymptotic setting, consider the following concrete parametric family. Let $\mathcal{U} = \{u_1, u_2, \ldots, u_N\}$ and fix an integer ℓ with $1 \le \ell < N$. For each ℓ-element subset $J \subseteq [N]$, define the axiom $A_J = \mathcal{U} \setminus \{u_j \mid j \in J\}$, i.e., the axiom satisfied by all objects except those indexed by J. Since $\ell < N$, each A_J is nonempty. The full axiom system is $\Sigma = \{A_J \mid J \in \binom{[N]}{\ell}\}$, with $n = \binom{N}{\ell}$ axioms. A sub-system $\{A_{J_1}, A_{J_2}, \ldots, A_{J_m}\}$ is impossible if and only if $\bigcup_{i=1}^{m} J_i = [N]$, i.e., the ℓ-sets cover all of $[N]$. Consequently, an impossibility core is a minimal covering of $[N]$ by ℓ-sets, and the impossibility dimension of Σ satisfies $\dim(\Sigma) = \lceil N/\ell \rceil$. In the special case $\ell = 1$, each axiom takes the form $A_i = \mathcal{U} \setminus \{u_i\}$, and the unique impossibility core is Σ itself with $\dim(\Sigma) = N$. The compatibility hypergraph is vertex-transitive (the symmetric group S_N acts transitively on the ℓ-subsets of $[N]$), and any two distinct minimum-size cores share at most $\dim(\Sigma) - 2$ axioms whenever $\dim(\Sigma) \ge 2$. By Theorem 14.3(ii)–(iii), with $k = \dim(\Sigma)$, the threshold is $p^*(n) = \Theta(M_k^{-1/k})$ and is sharp under the hypotheses of the theorem.

Example 14.8 (A Concrete Numerical Illustration of the Threshold)**.** We proceed to compute the threshold $p^*(n)$ explicitly for the canonical family of Example 14.7 at $\ell = 2$ (pairwise-exclusion axioms) on a domain of size N. Here, each axiom $A_{\{i,j\}} = \mathcal{U} \setminus \{u_i, u_j\}$ excludes exactly two elements, and the total number of axioms is $n = \binom{N}{2} = N(N-1)/2$. A sub-system $\{A_{J_1}, A_{J_2}, \ldots, A_{J_m}\}$ is impossible if and only if the pairs $J_1, J_2, \ldots, J_m \subseteq [N]$ jointly cover all of $[N]$, so the minimum-size impossibility cores are the minimum edge covers of the complete graph K_N. When N is even, each minimum edge cover is a perfect matching of K_N, of size $k = N/2$, and the count is $M_k = (N-1)(N-3)\cdots 3 \cdot 1 = (N-1)!!$. By Theorem 14.3(ii), the threshold satisfies

$$p^*(n) \;=\; \Theta\Big(M_k^{-1/k}\Big) \;=\; \Theta\Big(\big((N-1)!!\big)^{-2/N}\Big).$$

For $N = 10$, we have $k = 5$, $M_k = 9!! = 945$, and $p^*(n) \approx 945^{-1/5} \approx 0.254$. For $N = 20$, we have $k = 10$, $M_k = 19!! = 654{,}729{,}075$, and $p^*(n) \approx (654{,}729{,}075)^{-1/10} \approx 0.131$. For $N = 50$, we have $k = 25$, $\ln M_k \approx 73.15$, and $p^*(n) \approx \exp(-73.15/25) \approx 0.054$. The threshold decreases as N grows, and the phase transition becomes sharper. Around p^*, the impossibility probability rises from close to zero to close to one, and the width of the transition window relative to p^* shrinks like $O(1/\log n)$ by the Friedgut–Kalai sharpness bound. This illustrates the "knife-edge" phenomenon. For $N = 50$, an inclusion probability of 0.04 leaves the system asymptotically possible, while an inclusion probability of 0.08 makes it asymptotically impossible, a factor-2 change in p crossing the entire transition.

Remark 14.9 (No Safe Positive Density)**.** The phase transition result has an immediate qualitative consequence. Let $\mathcal{U}$ be a finite domain and let $(A_1, A_2, \ldots)$ be an infinite sequence of axioms drawn from $\mathcal{P}(\mathcal{U}) \setminus \{\emptyset\}$, in which every axiom of a finite impossible system $\{B_1, B_2, \ldots, B_m\}$ occurs with infinite multiplicity. Then, for any fixed $p > 0$, the random sub-system Σ_p obtained by including each A_i independently with probability p satisfies $\Pr[\mathcal{M}(\Sigma_p) = \emptyset] \to 1$ as $n \to \infty$. Indeed, each B_j occurs with multiplicity $c_j(n) \to \infty$, so the probability that B_j is excluded from Σ_p is $(1-p)^{c_j(n)} \to 0$; a union bound over the m distinct axioms gives the claim. In short, no positive inclusion density is safe against impossibility in large systems.

Remark 14.10 (The Uniqueness Analogue)**.** There is a parallel threshold phenomenon for the uniqueness regime. Define $\xi_n(p) = \Pr[|\mathcal{M}(\Sigma_p)| = 1]$. This function need not be monotone (adding axioms can push a singleton model class into impossibility), but it does have a "window" of p values around which uniqueness is most likely. By Theorem 14.3, the system is almost surely possible for $p \ll p^*(n)$ and almost surely impossible for $p \gg p^*(n)$. The uniqueness regime occupies at most a window of width $O(p^*(n))$ around the threshold, which tends to zero. The conjecture that $\xi_n(p^*(n)) \to 0$ as well, i.e., that at the threshold the system almost surely skips uniqueness with a first-order transition, remains open.

15 Scope and Limitations of the Finite-Domain Assumption

The framework developed in this paper rests on the assumption that the domain $\mathcal{U}$ is finite. We discuss the scope of this assumption, its consequences, and the prospects for extension.

The finite-domain setting yields clean results. The Impossibility Core Existence Theorem (Theorem 7.1) follows immediately from the finiteness of $\mathcal{U}$, chains terminate by cardinality, and the cooperative game (Σ, v_μ) is well-defined for any probability measure on $\mathcal{U}$. In fact, the Galois connection, the four-regime taxonomy, and the cooperative game structure all generalize to infinite domains under mild measurability conditions. Specifically, if $\mathcal{U}$ is equipped with a σ-algebra $\mathcal{F}$ and a probability measure μ, and every axiom $A \in \Sigma$ is $\mathcal{F}$-measurable, then $v_\mu(S) = 1 - \mu(\mathcal{M}(S))$ remains a well-defined, normalized, monotone, submodular set function, and all the results of Section 8 (including the Shapley value, the interaction theorem, and the characterization of axiom games) carry over without change.

The results that depend essentially on finiteness are the compactness theorem and the phase transition analysis. On an infinite domain, a system may be impossible without possessing any finite impossibility core, as the counterexample $\mathcal{U} = \mathbb{N}$, $A_n = \mathbb{N} \setminus \{n\}$ in Section 7 demonstrates. Extending the compactness result to infinite domains would require additional topological or algebraic compactness assumptions on the domain, analogous to the compactness theorem in first-order logic. The phase transition results of Section 14 are asymptotic statements about growing families of axioms on a fixed finite domain and do not directly apply when $\mathcal{U}$ is infinite.

The classical theorems of Section 13 bring out a further limitation. Arrow's theorem, Nash's bargaining theorem, and the Shapley value axiomatization all involve infinite domains in their standard formulations, namely the set of all social welfare functions, the set of all bargaining problems on the plane, and the space of all value operators on n-player cooperative games. The framework of this paper does not directly apply to those infinite

settings as theorems. We address this by two complementary strategies: the catalogue of Section 12 realizes the framework on small finite domains where every model-class size is derived by enumeration, capturing the qualitative pattern of the classical theorems in fully derivable form, while Section 13 discusses the parallels as heuristic hypotheses. A related limitation is that domain restriction, one of the three rescue types of Definition 9.1, is vacuous in the extensional setup of this paper (see Remark 9.7); extending the framework to accommodate intensional axioms that change meaning under domain restriction is an open problem. We emphasize, however, that the limitation concerns domain restriction specifically, not axiom weakening. Replacing an axiom A by a strictly larger set $A' \supsetneq A$ is an extensional operation that works within the present framework and is precisely what the Rescue Theorem (Theorem 9.4(ii)) characterizes. What lies outside the framework is the intensional phenomenon by which restricting the domain changes the very meaning of each axiom (for example, the set of social welfare functions satisfying IIA changes when the set of alternatives changes), yielding genuinely different subsets of the restricted domain.

For the extrapolation to infinite domains specifically, the Axiom Dominance Theorem (Theorem 8.12) provides a partial remedy. If a canonical measure μ^* on the true infinite Arrow domain could be identified and shown to satisfy the dominance chain $A_{\mathrm{IIA}} \succeq_{\mu^*} A_{\mathrm{P}} \succeq_{\mu^*} A_{\mathrm{ND}}$, then Theorem 8.12(iii) would give the qualitative ranking $\phi_{A_{\mathrm{IIA}}} > \phi_{A_{\mathrm{P}}} > \phi_{A_{\mathrm{ND}}}$ automatically, without any specific numerical computation. The analogous statement for Nash would give a ranking of the four axioms whenever a corresponding dominance chain holds under a canonical Nash measure. Whether such canonical measures exist and whether they satisfy the dominance chains is an open question. The framework's contribution to the classical setting is thus twofold. It provides a principled, canonical method for converting any measure into a unique attribution satisfying efficiency, symmetry, linearity, and the null-axiom property, and it certifies the qualitative ranking for all measures in a transparently defined class. The identification of a canonical measure on the true infinite domain remains open.

16 Conclusion

We have developed the abstract mathematics of the axiomatic method at three levels, namely qualitative, quantitative, and probabilistic.

At the qualitative level, the model-class map $\mathcal{M}$ is antitone, forms part of a Galois connection with the theory map Th, and drives a chain of strictly decreasing model classes terminating in one of the four regimes of Definition 4.2. Uniqueness is characterized by Theorem 6.1, impossibility by Theorem 7.1, and independence by Theorem 7.3.

At the quantitative level, the axiom game (Σ, v_μ) is a cooperative game, and the axiom Shapley value $\phi(\Sigma, \mu)$ is the unique, canonical, ordering-independent attribution of collapse credit (Theorem 8.4). Submodularity of v_μ is equivalent to diminishing marginal returns (Theorem 8.8), and the Axiom Dominance Theorem (Theorem 8.12) shows the qualitative Shapley ranking is robust over all measures respecting a given restrictiveness ordering, placing the examples of Section 12 on firm structural foundation. Theorem 8.17 characterizes axiom games as the set functions whose complement $1 - v$ has nonnegative Möbius coefficients, strictly stronger than submodularity for three or more axioms. The axiom Shapley value is the centroid of the base polytope (Proposition 8.14). The Rescue Theorem (Theorem 9.4) identifies the optimal single-axiom drop via the full-coalition marginal and distinguishes

canonical from minimum-cost weakening. Remark 9.5 and Example 9.10 show the drop criterion differs from the Shapley ranking. Impossibility depth $d(\Sigma)$ equals the transversal number $\tau(\mathcal{C}(\Sigma))$ of the impossibility cores (Theorem 9.8).

At the structural and probabilistic level, the impossibility dimension $\dim(\Sigma)$ (Section 10) is a combinatorial invariant measuring the minimum size of an impossibility core, and the Dimension-Witness Duality (Theorem 10.6) shows it equals the minimum domain size needed to certify independence. The Feasibility Complex Theorem (Theorem 11.3) shows that the collection of possible sub-systems forms a simplicial complex; for an impossibility core, this complex is homeomorphic to a sphere. The Alexander Duality Theorem (Theorem 11.11) identifies the Alexander dual of $K(\Sigma)$ with the co-core complex $K^{\vee}(\Sigma)$, whose maximal faces are the complements of the impossibility cores, and shows that $\widetilde{H}_i(K(\Sigma); \urcorner) \cong \widetilde{H}_{n-i-3}(K^{\vee}(\Sigma); \urcorner)$. The Contractibility Theorem (Theorem 11.7) and the cone-point acyclicity criterion (Corollary 11.12(ii)) both follow as corollaries. The Phase Transition Theorem (Theorem 14.3) shows that the transition from possibility to impossibility in random axiom systems is abrupt, governed by a sharp threshold $p^*(n) = \Theta(M_k^{-1/k})$ that depends on the number and size of impossibility cores.

Several directions remain open. Perhaps the most pressing concerns the intensional phenomenon of domain restriction (Remark 9.7). In the standard formulations of Arrow's or Nash's theorem, restricting the domain changes the *meaning* of each axiom, yielding genuinely different subsets of the restricted domain; extending the framework to accommodate this intensional operation is a natural next step. A second direction concerns the uniqueness window in phase transitions (Section 14). The conjecture that $\xi_n(p^*(n)) \to 0$, stating that uniqueness is almost surely bypassed at the critical probability, would imply a first-order transition. A third direction is to extend the Betti-number theory of the feasibility complex beyond the Alexander duality of Theorem 11.11. For non-core impossible systems with many overlapping cores, the higher Betti numbers $\beta_i(K(\Sigma))$ reduce via the co-core complex to a combinatorial question about core overlaps, and computing these for the classical examples may reveal new structural invariants. A fourth direction is the entropy lower bound question. For a fixed impossibility core of dimension k, the infimum of $H(\Sigma, \mu)$ over all measures μ giving each core axiom positive Shapley weight is zero in general (by degenerating μ to concentrate blame on one axiom), but under natural regularity conditions on μ a positive lower bound may hold. A fifth direction for future inquiry concerns the approximation ratio of the Shapley-greedy for the transversal problem. A fifth open question pertains to establishing tight bounds on the approximation ratio achieved by the Shapley-greedy algorithm in the context of the transversal problem, namely, quantifying the extent to which the greedy solution deviates from the optimal solution in the worst case. For submodular coverage problems, the greedy achieves a $(1-1/e)$ approximation. Whether an analogous bound holds for the rescue problem is open. A sixth direction, discussed in Sections 13 and 15, concerns the extension of the framework to the classical infinite-domain theorems of Arrow, Nash, and the Shapley value axiomatization. The Axiom Dominance Theorem (Theorem 8.12) shows that the qualitative Shapley ranking is invariant over the class of measures respecting a given dominance chain, so the identification of a canonical measure on the true infinite Arrow domain (respectively, Nash domain) satisfying the standard dominance chain would elevate the qualitative rankings from heuristic hypotheses to theorems about the infinite setting itself. This program appears difficult and remains open.

References

Kenneth J. Arrow. *Social Choice and Individual Values*. Wiley, New York, 1951; 2nd ed., Yale University Press, New Haven, 1963.

Yuliy M. Baryshnikov. Unifying Impossibility Theorems: A Topological Approach. *Advances in Applied Mathematics*, 14(4):404–415, 1993.

Julian H. Blau. A Direct Proof of Arrow's Theorem. *Econometrica*, 40(1):61–67, 1972.

Béla Bollobás and Andrew Thomason. Threshold functions. *Combinatorica*, 7(1):35–38, 1987.

Donald E. Campbell and Jerry S. Kelly. Impossibility theorems in the Arrovian framework. In Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, pages 35–94. Elsevier, Amsterdam, 2002.

Jack Edmonds. Submodular functions, matroids, and certain polyhedra. In R. Guy, H. Hanani, N. Sauer, and J. Schönheim, editors, *Combinatorial Structures and Their Applications*, pages 69–87. Gordon and Breach, New York, 1970.

Peter C. Fishburn. *Utility Theory for Decision Making*. Wiley, New York, 1970.

Ehud Friedgut and Gil Kalai. Every monotone graph property has a sharp threshold. *Proceedings of the American Mathematical Society*, 124(10):2993–3002, 1996.

Satoru Fujishige. *Submodular Functions and Optimization*. Annals of Discrete Mathematics 58. Elsevier, Amsterdam, 2nd edition, 2005.

Bernhard Ganter and Rudolf Wille. *Formal Concept Analysis: Mathematical Foundations*. Springer, Berlin, 1999.

Michel Grabisch and Marc Roubens. An axiomatic approach to the concept of interaction among players in cooperative games. *International Journal of Game Theory*, 28(4):547–565, 1999.

Dmitry N. Kozlov. *Combinatorial Algebraic Topology*. Algorithms and Computation in Mathematics 21. Springer, Berlin, 2008.

László Lovász. Submodular functions and convexity. In A. Bachem, M. Grötschel, and B. Korte, editors, *Mathematical Programming: The State of the Art*, pages 235–257. Springer, Berlin, 1983.

John F. Nash. The Bargaining Problem. *Econometrica*, 18(2):155–162, 1950.

Alvin E. Roth, editor. *The Shapley Value: Essays in Honor of Lloyd S. Shapley*. Cambridge University Press, Cambridge, 1988.

Donald G. Saari. *Disposing Dictators, Demystifying Voting Paradoxes: Social Choice Analysis*. Cambridge University Press, Cambridge, 2008.

Amartya K. Sen. *Collective Choice and Social Welfare*. Holden-Day, San Francisco, 1970.

Glenn Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, NJ, 1976.

Lloyd S. Shapley. A Value for n-Person Games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume II, Annals of Mathematics Studies 28, pages 307–317. Princeton University Press, Princeton, NJ, 1953.

Shmuel Weinberger. On the topological social choice model. *Journal of Economic Theory*, 115(2):377–384, 2004.

Rudolf Wille. Restructuring Lattice Theory: An Approach Based on Hierarchies of Concepts. In I. Rival, editor, *Ordered Sets*, NATO Advanced Study Institutes Series 83, pages 445–470. Reidel, Dordrecht, 1982.

Eyal Winter. The Shapley Value. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 3, chapter 53, pages 2025–2054. Elsevier, Amsterdam, 2002.