# An Equivariance Principle for the Axiomatic Method

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

September 28, 2026

**Abstract**

A $G$-equivariant map $R\colon X \to T$ must send each $x$ into $\mathrm{Fix}_T(\mathrm{Stab}_G(x))$; the set of all such maps bijects naturally with $\prod_{[x]\in X/G} \mathrm{Fix}_T(\mathrm{Stab}_G(x))$. An empty factor is impossibility, a singleton is characterization, a larger factor parametrizes a family. The bijection has been folklore since Eckmann (1954). We build on it a rigidity taxonomy of output spaces and an additive refinement theorem via the coinvariant Grothendieck group $(X^{\mathrm{gp}})_G$, exhibiting the impossibility–characterization dichotomy of the axiomatic corpus as one fixed-set product with three regimes. Five classical results (May's theorem, Nash bargaining, the Shapley value, the ANR impossibility, and the constraint theorem for label-symmetric classification) are derived uniformly, and the recipe is applied cold to a novel problem.

**Keywords.** Axiomatic method; equivariance; stabilizer; fixed-point set; impossibility theorem; characterization theorem; orbit partition; uniform completion; coinvariant Grothendieck group.

**Mathematics Subject Classification (2020).** 91B14; 91B06; 20B30; 05E18.

## 1 Introduction

The axiomatic method proceeds by specifying a class of admissible inputs, imposing axioms on the rule that assigns outputs to inputs, and asking whether such a rule exists and, if so, whether it is unique. Two outcomes are of particular interest. Sometimes the axioms are jointly inconsistent, in which case the conclusion is an *impossibility theorem* in the tradition of Arrow (1963), Gibbard (1973), Satterthwaite (1975). Sometimes the axioms determine a rule up to a scalar or a family, in which case the conclusion is a *characterization theorem* in the tradition of Nash (1950), Shapley (1953). These two outcomes are usually presented as opposites, and indeed the proofs that establish them in the classical literature have little visible structure in common. In fact, whenever the axioms include a symmetry hypothesis, both outcomes admit a partial derivation from the same elementary identification, which we now describe. The thesis of the present paper is that this identification, once developed into a discrete diagnostic framework, exhibits the impossibility-characterization dichotomy as one product-of-fixed-sets bijection with three regimes, and that the framework required to do so is elementary and self-contained.

### 1.1 The Observation

Many axioms in the corpus are symmetry conditions. Anonymity is invariance under permutation of the agents, neutrality is equivariance under permutation of the alternatives, and countless variants of the same idea (label-invariance in classification, unit-invariance in pricing, order-invariance in bargaining) demand that the rule commute with a natural group action. Aggregated together, these axioms demand that the rule be $G$-equivariant for a suitable group $G$ acting on inputs and outputs.

The observation is elementary. If $R\colon X \to T$ is $G$-equivariant, and if $g$ lies in the stabilizer of an input $x$, then

$$g \cdot R(x) = R(g \cdot x) = R(x),$$

so that $R(x)$ is a fixed point of $\mathrm{Stab}_G(x)$ acting on $T$. Hence, an equivariant rule from a $G$-set $X$ to a $G$-set $T$ is not free to take arbitrary values. Its value on the orbit of $x$ is determined by a single choice in $\mathrm{Fix}_T(\mathrm{Stab}_G(x))$, and no other value is available. Emptiness of that fixed set is impossibility. Singularity of that fixed set is uniqueness. Larger cardinality parametrizes a family of admissible rules.

### 1.2 What This Paper Does and Does Not Do

Three clarifications are in order before we proceed.

First, the bijection $\mathrm{Eq}_G(X,T) \cong \prod_{[x]} \mathrm{Fix}_T(\mathrm{Stab}_G(x))$ that sits at the base of the paper is elementary and, in its topological formulation, has been in the literature since at least Eckmann (1954). What is developed here is the framework built on it. We isolate the discrete bijection as a diagnostic principle applicable to any axiomatic problem of the appropriate shape. We prove a rigidity theorem that classifies the output space $T$ by the pair of invariants $(\mathrm{Rig}_G(T), \mathrm{Rig}_G^{=1}(T))$, and compute this pair for the output types most frequent in the axiomatic corpus. We prove an additive refinement theorem that identifies $G$-invariant partial-monoid homomorphisms out of a $G$-monoid $X$ with abelian-group homomorphisms out of the coinvariant Grothendieck group $(X^{\mathrm{gp}})_G$. We compute this coinvariant explicitly for the space of cooperative games, obtaining $\mathcal{G}_{\mathrm{Sym}(\mathcal{N})} \cong \mathbb{R}^n$, and read off from it the per-orbit scalar content of the Shapley formula. Under this framework the impossibility-characterization dichotomy of the axiomatic corpus becomes one product-of-fixed-sets identification. These are the paper's positive claims.

Second, the observation of Section 1.1 does not, on its own, prove any of the substantial theorems in the axiomatic corpus. What it does is isolate the symmetric core of those theorems, so that the remaining work of any given argument becomes visible and, in most cases, small. In every classical instance we examine, the equivariance principle handles the symmetric part in a few lines. The non-symmetry axioms then supply the further content, and are often orthogonal to the symmetric part. Whether the residual work is small or large depends on the theorem. For May's theorem, it is one monotonicity argument. For Nash bargaining, it is the reduction from general problems to symmetric ones by affine rescaling and independence of irrelevant alternatives. For the Shapley value, it is the linearity-plus-null-player determination of a single scalar on each orbit. The paper's claim is not that the equivariance principle proves these theorems by itself. It is that the principle exhibits their symmetric skeleton in a uniform way, that the framework built around it turns this skeleton into an explicit computation, and that the residual non-symmetry axioms then acquire a visible and often orthogonal character.

Third, the paper's target is the discrete case. Continuous refinements, notably the topological impossibility theorem of Chichilnisky (1980) and its cohomological extensions by Baryshnikov (1993) and Weinberger (2004), use continuous versions of the same principle in the setting of preference spaces of positive dimension, and are outside our scope. Quantitative refinements, in which the impossibility is measured rather than absolute, are likewise beyond the scope of this paper. Both are discussed in Section 8.

### 1.3 Contributions

The contributions of the paper are the following.

(1) *A uniform statement of the symmetric core.* We isolate the elementary principle underlying the symmetry hypothesis in a broad class of axiomatic theorems, and package it as a four-step recipe (Section 2.5) applicable to any axiomatic problem of the appropriate shape.

(2) *A rigidity taxonomy of output spaces.* We introduce the pair of invariants $(\mathrm{Rig}_G(T), \mathrm{Rig}_G^{=1}(T))$ that classifies a $G$-set $T$ by the collections of subgroups admitting, respectively, at least one and exactly one fixed element, and we compute this taxonomy for the output types that arise most frequently in the corpus (strict linear orders, partitions, weak orders, probability distributions, real-valued targets). The taxonomy recasts each impossibility or characterization question in the discrete equivariance setting as a containment relation between two collections of subgroups. To our knowledge, this pair of invariants and the containment criterion built from them have not appeared in the axiomatic literature in the present form.

(3) *An additive refinement theorem.* We prove that the set of $G$-invariant partial-monoid homomorphisms out of a $G$-monoid $X$ is in natural bijection with $\mathrm{Hom}_{\mathbf{Ab}}((X^{\mathrm{gp}})_G, M)$ (Theorem 4.3), where $(X^{\mathrm{gp}})_G$ is the coinvariant quotient of the Grothendieck group of $X$. The proof is a composition of two elementary universal properties, and the mathematical substance of the refinement lies in the identification of $(X^{\mathrm{gp}})_G$ as the correct invariant of a partial-monoid-with-$G$-action. This identification sharpens the bare fixed-set count of the equivariance principle whenever the output carries an additive structure and the rule is required to respect it. We then compute the coinvariant $(X^{\mathrm{gp}})_G$ explicitly for $X$ the space of cooperative games under $\mathrm{Sym}(\mathcal{N})$, obtaining $\mathcal{G}_{\mathrm{Sym}(\mathcal{N})} \cong \mathbb{R}^n$, and use this computation to identify the per-orbit scalar content of the Shapley formula, complementing the direct derivation from the equivariance principle in Section 5.3.

(4) *Five classical instances, worked out in full.* We provide complete derivations of five classical results, from the equivariance principle and its additive refinement, namely of May's theorem for two-alternative majority voting, the Nash bargaining solution, the Shapley value, the anonymity-neutrality-resoluteness impossibility for Arrovian aggregation, and the constraint theorem for label-symmetric classification. Nash bargaining is worked through the full reduction from arbitrary problems to symmetric ones by affine rescaling and independence of irrelevant alternatives. The Shapley value is derived on unanimity games from the equivariance principle plus the null-player axiom plus efficiency, and subsequently extended by linearity, with the coinvariant computation of contribution (3) supplying a complementary reading in terms of per-orbit scalar budgets. Label-symmetric classification exhibits

the mechanism outside the classical social-choice and game-theoretic setting, in a machine-learning context in which label-symmetry alone (with no further axiom) yields both a striking constraint theorem (a universal label- symmetric learner cannot predict a label outside the training data whenever at least two labels are absent) and the accompanying characterization bijection.

(5) *A cold application of the recipe.* We illustrate the diagnostic value of the four-step recipe by applying it to a small axiomatic problem not previously in the literature, and we obtain a characterization of the uniform-random-choice selector as the unique anonymous selector on finite menus. The purpose is to exhibit the recipe as a working tool rather than a retrospective organizational device.

(6) *A gallery of further instances.* We treat, more briefly, seven further axiomatic problems drawn from bankruptcy theory, random assignment, behavioral economics, tie-breaking, canonical randomization, pricing, and market design. Each is worked to the point where the equivariance step is completed, with pointers to the sources where the residual non-symmetry arguments are found.

The bijection at the base is elementary and, as noted, has been in the literature since at least Eckmann (1954). The theoretical additions built on it, however, are ours. The rigidity taxonomy of Section 3, the additive refinement theorem of Section 4, the coinvariant computation for symmetric games, the unified reading of the impossibility-characterization dichotomy as one product-of-fixed-sets identification, and the small original characterization of Proposition 6.1, together constitute the paper's positive contribution. The recipe, the five fully worked classical instances, and the gallery of further instances demonstrate that this contribution is a working tool rather than a formalism. A working theorist encountering a new axiomatic problem may, after this paper, check the fixed-set product before proving anything, and either read off the answer or identify precisely which additional structure is doing the remaining work.

### 1.4 Reading Guide

The paper is self-contained. A reader familiar with elementary group actions may read Section 2 in a single sitting, and select worked instances from Sections 5 and 7 at leisure. The refinements in Sections 3 and 4 may be read independently. Several instances in Sections 5 and 7 draw on the author's concurrent working papers (Feys, 2026a,b,c,d,e,f,g,h,i), which are referenced for the residual non-symmetry arguments and, where relevant, for extended derivations; the equivariance step in each such instance is self-contained in the present paper. A reader interested only in the unification result may proceed directly from Section 2 to Section 5, treating the rigidity taxonomy of Section 3 and the additive refinement of Section 4 as reference material to be consulted as needed. Notation is standard. For a group $G$ acting on a set $X$ from the left, we write $\mathrm{Stab}_G(x) = \{g \in G \mid g \cdot x = x\}$ for the stabilizer of $x$, and for a $G$-set $T$ and a subgroup $H \leq G$, we write $\mathrm{Fix}_T(H) = \{t \in T \mid h \cdot t = t \text{ for all } h \in H\}$ for the fixed set of $H$. The set of $G$-orbits of $X$ is written $X/G$, and orbit representatives are denoted by $[x]$. Throughout the paper, $G$ is an arbitrary group, and we work with left actions on sets. Finiteness of $G$ or of $X$ is not assumed, except where explicitly stated.

## 2 The Equivariance Principle

The purpose of this section is to state and prove the central principle in its cleanest form, and to record the two corollaries that furnish the impossibility and characterization sides of the axiomatic dichotomy. The proof is a three-line calculation, and we shall not disguise this brevity. The interest of the statement lies elsewhere. It resides in the range of theorems the resulting bijection organizes, in the taxonomy of output-side rigidity it supports (Section 3), and in the additive refinement it admits (Section 4). The remaining sections of the paper are, in one form or another, uses of this bijection or extensions of it.

### 2.1 Setup

Throughout this section, let $G$ be a group acting from the left on two sets $X$ and $T$. The set $X$ is the set of admissible inputs, and $T$ is the set of admissible outputs. A function $R\colon X \to T$ is *$G$-equivariant* if $R(g \cdot x) = g \cdot R(x)$ for every $g \in G$ and $x \in X$; it is *$G$-invariant* if $R(g \cdot x) = R(x)$ for every $g$ and $x$, which is the special case in which $T$ carries the trivial $G$-action. Write $\mathrm{Eq}_G(X, T)$ for the set of $G$-equivariant maps from $X$ to $T$. Recall that a *fundamental domain* for the $G$-action on $X$ is a subset $\Sigma \subseteq X$ that meets every $G$-orbit in exactly one point. The axiom of choice guarantees existence of such a $\Sigma$ for arbitrary $G$-sets, and no further hypotheses are required.

### 2.2 The Principle

The principle is the following identification.

**Theorem 2.1** (Equivariance Principle)**.** *Let $G$ act on $X$ and $T$, and let $\Sigma \subseteq X$ be a fundamental domain for the $G$-action on $X$. The evaluation map*

$$\mathrm{Eq}_G(X, T) \longrightarrow \prod_{x \in \Sigma} \mathrm{Fix}_T(\mathrm{Stab}_G(x)), \qquad R \longmapsto \big(R(x)\big)_{x \in \Sigma},$$

*is a bijection. In particular, the choice of $\Sigma$ notwithstanding, the cardinality of $\mathrm{Eq}_G(X, T)$ equals $\prod_{[x] \in X/G} |\mathrm{Fix}_T(\mathrm{Stab}_G(x))|$.*

*Proof.* Well-definedness of the evaluation map is the computation

$$h \cdot R(x) = R(h \cdot x) = R(x) \qquad (h \in \mathrm{Stab}_G(x)),$$

which shows $R(x) \in \mathrm{Fix}_T(\mathrm{Stab}_G(x))$ for every $x \in \Sigma$.

Injectivity is immediate. An equivariant $R$ is determined by its values on $\Sigma$, since the identity $R(g \cdot x) = g \cdot R(x)$ propagates the value at $x \in \Sigma$ to the entire orbit of $x$.

Surjectivity is the reconstruction. Given a family $(t_x)_{x \in \Sigma}$ with $t_x \in \mathrm{Fix}_T(\mathrm{Stab}_G(x))$, define $R\colon X \to T$ as follows. For every $y \in X$, there is a unique $x \in \Sigma$ with $y = g \cdot x$ for some $g \in G$, and we set $R(y) \coloneqq g \cdot t_x$. This is well-defined; if also $y = g' \cdot x$, then $g^{-1}g' \in \mathrm{Stab}_G(x)$, so $g^{-1}g' \cdot t_x = t_x$, i.e., $g' \cdot t_x = g \cdot t_x$. We verify equivariance. For any $h \in G$ and $y \in X$ with $y = g \cdot x$ (for $x \in \Sigma$), we have $h \cdot y = (hg) \cdot x$, so $R(h \cdot y) = (hg) \cdot t_x = h \cdot (g \cdot t_x) = h \cdot R(y)$. Finally, the identity $R(x) = e \cdot t_x = t_x$ for $x \in \Sigma$ is immediate. □

### 2.3 Two Corollaries

The following two consequences do the work in every instance we shall examine.

**Corollary 2.2** (Impossibility)**.** *No $G$-equivariant map $R\colon X \to T$ exists if and only if there is some $x \in X$ with $\mathrm{Fix}_T(\mathrm{Stab}_G(x)) = \emptyset$.*

*Proof.* The product in Theorem 2.1 is empty if and only if some factor is empty, and one factor is empty if and only if some orbit representative $x \in \Sigma$ has empty $\mathrm{Fix}_T(\mathrm{Stab}_G(x))$. It remains to verify that this condition is independent of the choice of representative. If $x' = g \cdot x$ lies in the same orbit, then $\mathrm{Stab}_G(x') = g\,\mathrm{Stab}_G(x)\,g^{-1}$, since $h \cdot x' = x'$ if and only if $g^{-1}hg \cdot x = x$. For any $t \in T$ and $h \in \mathrm{Stab}_G(x)$, we have $t \in \mathrm{Fix}_T(\mathrm{Stab}_G(x))$ if and only if $g \cdot t \in \mathrm{Fix}_T(g\,\mathrm{Stab}_G(x)\,g^{-1})$, because $(ghg^{-1}) \cdot (g \cdot t) = g \cdot (h \cdot t)$. Hence, the map $t \mapsto g \cdot t$ restricts to a bijection $\mathrm{Fix}_T(\mathrm{Stab}_G(x)) \to \mathrm{Fix}_T(\mathrm{Stab}_G(x'))$, and in particular the two fixed sets are simultaneously empty or nonempty. □

**Corollary 2.3** (Uniqueness)**.** *A unique $G$-equivariant map $R\colon X \to T$ exists if and only if $|\mathrm{Fix}_T(\mathrm{Stab}_G(x))| = 1$ for every $x \in X$.*

*Proof.* The product in Theorem 2.1 is a singleton if and only if every factor is a singleton. □

The impossibility and uniqueness corollaries are the two ends of a single spectrum, whose middle is populated by settings in which $\mathrm{Fix}_T(\mathrm{Stab}_G(x))$ has more than one element for some orbits, and the family of admissible rules is parametrized by the corresponding product. This intermediate case is precisely what the "characterization up to a family" theorems of the corpus produce. The count is neither zero nor one but a small combinatorial number. Figure 1 shows the three regimes at a glance.

Some axiomatic problems present the output space not as a single $G$-set $T$ but as a $G$-equivariant family of fibers over the input $X$, and the rule is required to select an element of the fiber over each input. The equivariance principle extends verbatim to this setting, and we record the extension for later use.

**Corollary 2.4** (Equivariant Sections)**.** *Let $G$ act on $X$, let $p\colon Y \to X$ be a $G$-equivariant surjection of $G$-sets, and let $\Sigma \subseteq X$ be a fundamental domain for the $G$-action on $X$. The evaluation map*

$$\mathrm{Eq}_G^{\mathrm{sec}}(p) \longrightarrow \prod_{x \in \Sigma} \mathrm{Fix}_{p^{-1}(x)}(\mathrm{Stab}_G(x)), \qquad R \longmapsto \big(R(x)\big)_{x \in \Sigma},$$

*is a bijection, where $\mathrm{Eq}_G^{\mathrm{sec}}(p)$ denotes the set of $G$-equivariant sections of $p$ and $\mathrm{Stab}_G(x)$ acts on the fiber $p^{-1}(x)$ by restriction of the $G$-action on $Y$ (which stabilizes $p^{-1}(x)$ setwise, since $p$ is equivariant).*

*Proof.* The action of $\mathrm{Stab}_G(x)$ on $Y$ preserves $p^{-1}(x)$, because for any $y \in p^{-1}(x)$ and $h \in \mathrm{Stab}_G(x)$, we have $p(h \cdot y) = h \cdot p(y) = h \cdot x = x$, so $h \cdot y \in p^{-1}(x)$. The proof of Theorem 2.1 then carries over with the constraint $R(x) \in p^{-1}(x)$. Well-definedness places $R(x)$ in $\mathrm{Fix}_{p^{-1}(x)}(\mathrm{Stab}_G(x))$; injectivity follows from equivariance; surjectivity uses the same reconstruction $R(g \cdot x) \coloneqq g \cdot t_x$ with $t_x \in \mathrm{Fix}_{p^{-1}(x)}(\mathrm{Stab}_G(x)) \subseteq p^{-1}(x)$, and equivariance of $p$ gives $R(g \cdot x) \in p^{-1}(g \cdot x)$ as required. □

The bundle version reduces to Theorem 2.1 when $p$ is the projection $X \times T \to X$; the fibers are then all equal to $T$, and $\mathrm{Fix}_{p^{-1}(x)}(\mathrm{Stab}_G(x)) = \mathrm{Fix}_T(\mathrm{Stab}_G(x))$. The bundle version shall be used in Section 6 for the uniform-random-choice characterization, where the output space is genuinely fibered over the input.

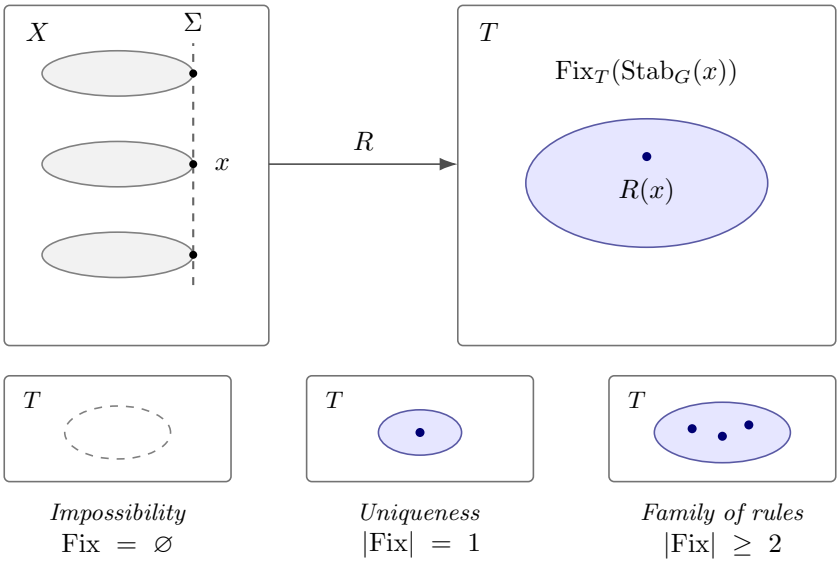



Figure 1: The equivariance principle at a glance. The set $X$ decomposes into $G$-orbits, and a fundamental domain $\Sigma \subseteq X$ selects one representative per orbit. Equivariance forces the value $R(x)$ into the fixed set $\mathrm{Fix}_T(\mathrm{Stab}_G(x)) \subseteq T$. The three panels below correspond to emptiness, singularity, or larger cardinality of that fixed set, yielding impossibility, uniqueness, or a family of admissible rules.

### 2.4 A Two-Line Example: Anonymous Strict Tie-Breaking

To fix ideas before proceeding, we record the shortest interesting instance of the principle. Let $\mathcal{N} = \{1, 2\}$, and let a rule $R\colon \{h\} \to \mathrm{LinOrd}(\mathcal{N})$ assign to a trivial input $h$ a strict order on $\mathcal{N}$. Impose anonymity, namely $\mathrm{Sym}(\mathcal{N})$-equivariance of $R$. The stabilizer of $h$ is $\mathrm{Sym}(\mathcal{N})$ itself, and the fixed set of $\mathrm{Sym}(\mathcal{N})$ acting on $\mathrm{LinOrd}(\mathcal{N})$ is empty (the swap $(1\,2)$ exchanges the two possible orders). By Corollary 2.2, no anonymous strict rule exists. This is the anonymity impossibility of Feys (2026a) in its smallest form, and it is representative of every impossibility we shall encounter.

### 2.5 A Recipe for the Reader

The principle is easy to apply. Given an axiomatic problem in which some rule is claimed to exist, to be unique, or to be characterized up to a family, the following four steps settle the discrete part of the question.

**Step 1.** Identify the input space $X$ and the output space $T$ as sets.

**Step 2.** Identify the symmetry group $G$ whose equivariance is imposed by the axioms, and record the $G$-actions on $X$ and $T$.

**Step 3.** Choose a fundamental domain $\Sigma \subseteq X$ for the $G$-action, and for each representative $x \in \Sigma$ compute the stabilizer $\mathrm{Stab}_G(x) \leq G$.

**Step 4.** For each $x \in \Sigma$ compute the fixed set $\mathrm{Fix}_T(\mathrm{Stab}_G(x)) \subseteq T$.

The conclusion is then read off from the fixed sets. If some $\mathrm{Fix}_T(\mathrm{Stab}_G(x))$ is empty, no $G$-equivariant rule exists (Corollary 2.2). If every $\mathrm{Fix}_T(\mathrm{Stab}_G(x))$ is a singleton, a unique $G$-

equivariant rule exists and its value at each orbit is determined (Corollary 2.3). Otherwise, the $G$-equivariant rules form a family in bijection with the product $\prod_{x \in \Sigma} \mathrm{Fix}_T(\mathrm{Stab}_G(x))$.

The four steps are pure computation, and the hardest of them is Step 4, which reduces to identifying the orbits of the stabilizer's action on the output space. For the two-line example of Section 2.4, the steps are trivial. We have $X = \{h\}$, $T = \mathrm{LinOrd}(\{1, 2\})$, $G = \mathrm{Sym}(\{1, 2\})$; the only representative $h$ has stabilizer $\mathrm{Sym}(\{1, 2\})$; the fixed set of $\mathrm{Sym}(\{1, 2\})$ on $\mathrm{LinOrd}(\{1, 2\})$ is empty. Impossibility follows immediately. The flagship instances of Section 5 and the further instances of Section 7 are worked through the same four steps, with more elaborate input and output spaces.

Two remarks are in order. First, the recipe presumes that the axioms of the problem are, or can be recast as, equivariance under a group action. Nearly every classical symmetry axiom (anonymity, neutrality, unit-invariance, label-invariance, order-invariance, translation-invariance) admits such a recasting. The rare exceptions typically involve monotonicity, continuity, or strategy-proofness, and are discussed in Section 8. Second, when the axioms include algebraic structure on $T$ (an additive or scaling axiom), the recipe still applies but yields a coarser answer than necessary. The sharper answer is furnished by the additive refinement of Section 4.

# 3 Rigidity of the Output

The principle of Section 2 reduces both impossibility and characterization to a single fixed-set computation on the output side, and the question of whether an equivariant rule exists, and whether it is unique, becomes a question about the geometry of the $G$-action on $T$. Which subgroups $H \leq G$ have $\mathrm{Fix}_T(H)$ nonempty, and among those, which have $\mathrm{Fix}_T(H)$ a singleton? This section formalizes the question and works through a taxonomy of output types. The taxonomy is not exhaustive; it is a small library of common cases, sufficient to organize the instances of Sections 5 and 7.

## 3.1 The Rigidity Type

The right invariant of the output space is the following.

**Definition 3.1** (Rigidity Type)**.** Let $T$ be a $G$-set. The *rigidity type* of $T$ is

$$\mathrm{Rig}_G(T) \coloneqq \{H \leq G \mid \mathrm{Fix}_T(H) \neq \emptyset\},$$

regarded as a collection of subgroups of $G$. The *sharp rigidity type* is the subcollection

$$\mathrm{Rig}_G^{=1}(T) \coloneqq \{H \leq G \mid |\mathrm{Fix}_T(H)| = 1\}.$$

The rigidity type is downward-closed in the subgroup lattice. Indeed, if $H \leq K$ and $t \in \mathrm{Fix}_T(K)$, then every $h \in H$ lies in $K$, so $h \cdot t = t$, and therefore $t \in \mathrm{Fix}_T(H)$. In particular, $\mathrm{Fix}_T(K) \subseteq \mathrm{Fix}_T(H)$, so nonemptiness of $\mathrm{Fix}_T(K)$ implies nonemptiness of $\mathrm{Fix}_T(H)$, and $H \in \mathrm{Rig}_G(T)$ whenever $K \in \mathrm{Rig}_G(T)$. Equivalently, the complementary collection of subgroups with empty fixed set is upward-closed, and is exactly the collection of stabilizer subgroups that would produce impossibility.

The corresponding invariant on the input side is the following.

**Definition 3.2** (Stabilizer Spectrum)**.** Let $X$ be a $G$-set. The *stabilizer spectrum* of $X$ is

$$\mathrm{Spec}_G(X) \coloneqq \{\mathrm{Stab}_G(x) \mid x \in X\},$$

regarded as a collection of subgroups of $G$.

The principle of Section 2 translates immediately into the following criterion.

**Proposition 3.3** (Rigidity Criterion)**.** *Let $G$ act on $X$ and $T$. Then*

(i) *a $G$-equivariant map $R\colon X \to T$ exists if and only if* $\mathrm{Spec}_G(X) \subseteq \mathrm{Rig}_G(T)$*;*

(ii) *a unique such map exists if and only if* $\mathrm{Spec}_G(X) \subseteq \mathrm{Rig}_G^{=1}(T)$.

*Proof.* For (i), a $G$-equivariant map exists if and only if no factor in the product of Theorem 2.1 is empty, which holds if and only if $\mathrm{Fix}_T(\mathrm{Stab}_G(x)) \neq \emptyset$ for every $x \in X$. Since the stabilizer depends only on the orbit, this is equivalent to asking that every subgroup in $\mathrm{Spec}_G(X)$ belong to $\mathrm{Rig}_G(T)$, i.e., $\mathrm{Spec}_G(X) \subseteq \mathrm{Rig}_G(T)$. For (ii), a unique $G$-equivariant map exists if and only if every factor in the product is a singleton, which holds if and only if $|\mathrm{Fix}_T(\mathrm{Stab}_G(x))| = 1$ for every $x \in X$, i.e., $\mathrm{Spec}_G(X) \subseteq \mathrm{Rig}_G^{=1}(T)$. □

Note that both $\mathrm{Rig}_G(T)$ and $\mathrm{Spec}_G(X)$ are closed under conjugation in $G$, since $\mathrm{Fix}_T(gHg^{-1}) \cong \mathrm{Fix}_T(H)$ for every $g \in G$ (as shown in the proof of Corollary 2.2), and since $\mathrm{Stab}_G(g \cdot x) = g\,\mathrm{Stab}_G(x)\,g^{-1}$. The containment $\mathrm{Spec}_G(X) \subseteq \mathrm{Rig}_G(T)$ is therefore well-posed up to conjugacy.

The rigidity criterion recasts impossibility and characterization as a matching problem. Is the collection of stabilizer subgroups that arise as inputs contained in the collection of subgroups that admit fixed outputs? A mismatch (a stabilizer with empty fixed set) is an impossibility. A tight match (every stabilizer has a singleton fixed set) is a characterization.

## 3.2 A Taxonomy of Common Output Types

The rigidity type is a function of both $G$ and $T$, and its structure is dictated by the geometry of the $G$-action on $T$. We work through the cases that arise in the corpus. Throughout, $\mathcal{N}$ is a finite set with $|\mathcal{N}| = n$, and $\mathrm{Sym}(\mathcal{N})$ acts on the relevant output space by the induced permutation action.

### 3.2.1 Strict Linear Orders

Let $T = \mathrm{LinOrd}(\mathcal{N})$, the set of strict linear orders on $\mathcal{N}$, with the natural $\mathrm{Sym}(\mathcal{N})$-action. A subgroup $H \leq \mathrm{Sym}(\mathcal{N})$ fixes an order $\prec$ if and only if $H$ acts on $\mathcal{N}$ by order-preserving permutations, i.e., $\sigma(a) \prec \sigma(b)$ whenever $a \prec b$, for every $\sigma \in H$. We claim that the only order-preserving permutation of a finite strictly ordered set is the identity. To see this, let $\sigma$ be an order-preserving permutation of $(\mathcal{N}, \prec)$, and write $a_1 \prec a_2 \prec \cdots \prec a_n$ for the enumeration of $\mathcal{N}$ in increasing order. Then $\sigma(a_1) \prec \sigma(a_2) \prec \cdots \prec \sigma(a_n)$, so $\sigma$ maps the unique $\prec$-minimum $a_1$ to an element $\prec$-below $\sigma(a_2)$; in fact, $\sigma(a_1)$ must equal $a_1$, for if $\sigma(a_1) = a_j$ with $j > 1$, then $\sigma^{-1}(a_1) \prec \sigma^{-1}(a_j) = a_1$, contradicting minimality of $a_1$ (since $\sigma^{-1}$ is also order-preserving). By induction on the elements of $(\mathcal{N}, \prec)$ in increasing order, $\sigma(a_k) = a_k$ for every $k$, so $\sigma = \mathrm{id}$. Hence, $\mathrm{Rig}_{\mathrm{Sym}(\mathcal{N})}(\mathrm{LinOrd}(\mathcal{N}))$ consists of the trivial subgroup alone, i.e., $\mathrm{Rig}_{\mathrm{Sym}(\mathcal{N})}(\mathrm{LinOrd}(\mathcal{N})) = \big\{\{e\}\big\}$. The output space is maximally rigid, and any nontrivial stabilizer in the input yields an impossibility.

This is why the anonymity axiom in a strict-order-valued setting is so often incompatible with any nontrivial symmetric input. Every impossibility of the anonymity-neutrality-resoluteness family, from Arrow to the tie-breaking impossibility of Feys (2026a), is an instance of this rigidity.

### 3.2.2 Partitions

Let $T = \mathrm{Part}(\mathcal{N})$, the set of set-partitions of $\mathcal{N}$, with the natural $\mathrm{Sym}(\mathcal{N})$-action. A subgroup $H \leq \mathrm{Sym}(\mathcal{N})$ fixes a partition $\Pi$ if and only if $H$ permutes the blocks of $\Pi$ as a set system, i.e., for every $\sigma \in H$ and every block $B \in \Pi$, the image $\sigma(B)$ is again a block of $\Pi$. Note that this does *not* require each individual block to be $H$-invariant; the group $H$ may fix $\Pi$ by permuting distinct blocks of the same size among themselves. The partition of $\mathcal{N}$ into $H$-orbits, denoted $\Pi_H$, is $H$-fixed by construction, since each $H$-orbit $O$ is $H$-invariant ($h \cdot O = O$ for every $h \in H$), so each orbit maps to itself and in particular $H$ permutes the blocks of $\Pi_H$ (trivially). Moreover, any coarsening of $\Pi_H$ (a partition obtained by merging blocks of $\Pi_H$) is also $H$-fixed, since a union of $H$-invariant sets is $H$-invariant, and $H$ therefore fixes each merged block individually. In particular, the one-block partition $\{\mathcal{N}\}$ is fixed by every subgroup. Hence, $\mathrm{Rig}_{\mathrm{Sym}(\mathcal{N})}(\mathrm{Part}(\mathcal{N}))$ is the entire subgroup lattice of $\mathrm{Sym}(\mathcal{N})$. The output space is fully flexible, and no impossibility can arise from stabilizer emptiness alone.

The sharp rigidity type $\mathrm{Rig}^{=1}$ is more restrictive. The orbit partition $\Pi_H$ is $H$-fixed, but is not in general the unique $H$-fixed partition. On the one hand, every coarsening of $\Pi_H$ (merging blocks of $\Pi_H$) is also $H$-fixed. On the other hand, there may exist $H$-fixed partitions whose blocks are *not* unions of $H$-orbits, arising when $H$ permutes two or more blocks nontrivially; for instance, if $\mathcal{N} = \{1, 2, 3, 4\}$ and $H = \langle (1\ 3)(2\ 4) \rangle$, the partition $\{\{1, 2\}, \{3, 4\}\}$ is $H$-fixed (the generator swaps the two blocks) even though neither block is $H$-invariant. In fact, uniqueness of the $H$-fixed partition holds only under an additional structural assumption on the input, typically compatibility with a background weak order that prevents nontrivial block permutation. This is the point at which the characterization theorems of Feys (2026a) intervene, and the orbit partition emerges as the unique honest partition-valued rule.

### 3.2.3 Weak Orders

Let $T = \mathrm{Wk}(\mathcal{N})$, the set of weak orders on $\mathcal{N}$, with the natural $\mathrm{Sym}(\mathcal{N})$-action. An element of $\mathrm{Wk}(\mathcal{N})$ is a pair $(\Pi, \prec)$ where $\Pi$ is a partition of $\mathcal{N}$ and $\prec$ is a total order on the blocks. A subgroup $H \leq \mathrm{Sym}(\mathcal{N})$ fixes $(\Pi, \prec)$ if and only if $H$ permutes the blocks of $\Pi$ in an order-preserving way with respect to $\prec$. Since an order-preserving permutation of a finite totally ordered set is necessarily the identity, each block must be individually $H$-invariant. In summary, $(\Pi, \prec) \in \mathrm{Fix}_{\mathrm{Wk}(\mathcal{N})}(H)$ if and only if $\Pi$ is a coarsening of the orbit partition $\Pi_H$ and $\prec$ is any total order on the blocks of $\Pi$.

The count now becomes explicit. Let $k$ denote the number of $H$-orbits, and write $\mathcal{O} = \{O_1, O_2, \ldots, O_k\}$ for the orbit set. We have shown that every $H$-fixed weak order $(\Pi, \prec)$ has the property that each block of $\Pi$ is a union of $H$-orbits, hence the underlying partition $\Pi$ is a coarsening of the orbit partition $\Pi_H$. Note that this conclusion is special to weak orders; for bare partitions without a total order on blocks, $H$ may fix a partition by permuting blocks nontrivially, as discussed above. The coarsenings of $\Pi_H$ are in bijection with the set-partitions of $\mathcal{O}$, and there are $B_k$ of them, where $B_k$ is the $k$-th Bell number. For each

such coarsening $\Pi$ with $\ell$ blocks, there are $\ell!$ total orders on the blocks. Summing over all coarsenings, the total number of $H$-fixed weak orders is the number of ordered set-partitions of a $k$-element set, which is the $k$-th Fubini number $a_k$ (also called the ordered Bell number); the sequence begins $a_1 = 1$, $a_2 = 3$, $a_3 = 13$, $a_4 = 75$, $a_5 = 541$.

Two limiting cases are worth naming. For $H$ transitive on $\mathcal{N}$ (so $k = 1$), the unique $H$-fixed weak order is the trivial one in which all elements are tied. Hence, $\mathrm{Rig}^{=1}_{\mathrm{Sym}(\mathcal{N})}(\mathrm{Wk}(\mathcal{N}))$ contains every transitive subgroup, and characterization theorems for weak-order-valued rules under a transitive symmetry are automatic. For $H = \{e\}$, every weak order is fixed and no rigidity holds.

Weak orders interpolate between partitions and probability distributions in rigidity. The fixed set $\mathrm{Fix}_{\mathrm{Wk}(\mathcal{N})}(H)$ has $a_k$ elements (the $k$-th Fubini number), which exceeds the Bell number $B_k$ for $k \geq 2$. Note, however, that $B_k$ counts only those $H$-fixed partitions whose blocks are unions of $H$-orbits. The total number of $H$-fixed partitions may be larger when $H$ admits nontrivial block permutations, as the example in the partition case above illustrates, and may in fact exceed $a_k$. In particular, weak orders are not uniformly "more flexible" than partitions; the comparison $a_k \geq B_k$ relates weak orders to orbit-coarsenings, not to all $H$-fixed partitions. The sharp rigidity type $\mathrm{Rig}^{=1}$ for weak orders coincides with that of the simplex. Uniqueness holds precisely under transitivity of $H$. Weak orders arise in the axiomatic theory of ranking rules that admit ties, and in the theory of preference aggregation with indifferences. A characterization theorem in this setting typically selects one of the $a_k$ admissible weak orders by an additional monotonicity or Pareto axiom.

### 3.2.4 Probability Distributions

Let $T = \Delta(\mathcal{N})$, the simplex of probability distributions on $\mathcal{N}$, with the $\mathrm{Sym}(\mathcal{N})$-action by coordinate permutation. A subgroup $H \leq \mathrm{Sym}(\mathcal{N})$ fixes a distribution $p \in \Delta(\mathcal{N})$ if and only if $p(\sigma(i)) = p(i)$ for every $\sigma \in H$ and $i \in \mathcal{N}$, which holds if and only if $p$ is constant on each $H$-orbit of $\mathcal{N}$. To see the "if" direction, suppose $p$ is constant on $H$-orbits and let $\sigma \in H$; then for every $i \in \mathcal{N}$, the element $\sigma(i)$ lies in the same $H$-orbit as $i$, so $p(\sigma(i)) = p(i)$, and $\sigma$ fixes $p$. For the "only if" direction, suppose $\sigma$ fixes $p$ for every $\sigma \in H$; then for any $i, j$ in the same $H$-orbit, there exists $\sigma \in H$ with $\sigma(i) = j$, so $p(j) = p(\sigma(i)) = p(i)$. The $H$-fixed distributions are therefore precisely those that assign a constant value to each $H$-orbit $O_1, O_2, \ldots, O_k$, say $p(i) = c_\ell$ for $i \in O_\ell$, subject to $\sum_{\ell=1}^{k} |O_\ell| \cdot c_\ell = 1$ and $c_\ell \geq 0$. This is a $(k-1)$-dimensional simplex, and is in particular nonempty for every $H$. Hence, $\mathrm{Rig}_{\mathrm{Sym}(\mathcal{N})}(\Delta(\mathcal{N}))$ is the entire subgroup lattice.

For the full subgroup $H = \mathrm{Sym}(\mathcal{N})$, the action on $\mathcal{N}$ is transitive, so there is a single orbit $O_1 = \mathcal{N}$, and the constraint $|O_1| \cdot c_1 = 1$ forces $c_1 = 1/n$. The unique $\mathrm{Sym}(\mathcal{N})$-fixed distribution is therefore the uniform distribution $p(i) = 1/n$ for every $i$. More generally, any transitive subgroup $H \leq \mathrm{Sym}(\mathcal{N})$ has a single orbit on $\mathcal{N}$, so the same constraint forces the unique $H$-fixed distribution to be uniform. For any subgroup $H$ with $k \geq 2$ orbits on $\mathcal{N}$, the fixed set is a simplex of dimension $k - 1 \geq 1$, and in particular contains more than one element. Hence, $\mathrm{Rig}^{=1}_{\mathrm{Sym}(\mathcal{N})}(\Delta(\mathcal{N}))$ is precisely the collection of transitive subgroups of $\mathrm{Sym}(\mathcal{N})$, and uniqueness holds if and only if the stabilizer acts transitively on $\mathcal{N}$. This is the mechanism behind the uniform-completion rule of Feys (2026b), and behind the anonymity characterization of the uniform reference in maximum-entropy pricing (Feys, 2026c).

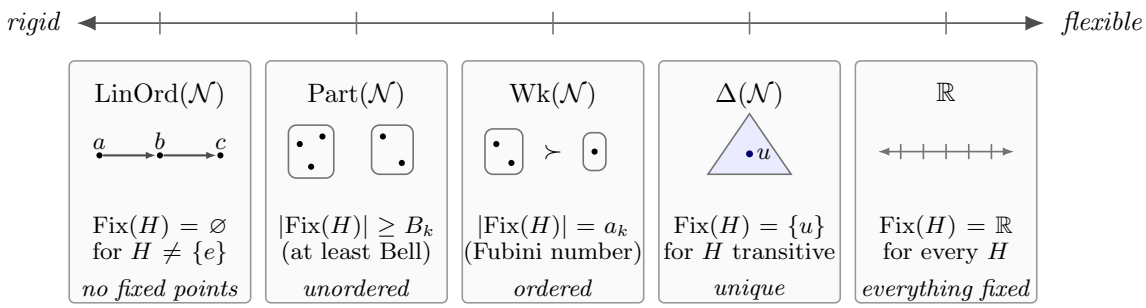



Figure 2: The rigidity spectrum of common output types under a $\mathrm{Sym}(\mathcal{N})$-action. From left to right, output rigidity decreases and the fixed sets grow. The count $|\mathrm{Fix}(H)|$ is written in terms of $k$, the number of $H$-orbits on $\mathcal{N}$. For weak orders the count $a_k$ is the $k$-th Fubini number (ordered Bell number); for partitions the count is at least $B_k$ (the $k$-th Bell number, counting coarsenings of the orbit partition) and may be larger when $H$ admits nontrivial block permutations. Every impossibility theorem for a strict-order-valued rule and every characterization theorem for a distribution-valued, weak-order-valued, or partition-valued rule can be located on this axis.

#### 3.2.5 The Real Line and Vector Spaces

Let $T = \mathbb{R}$ with the trivial $G$-action, or more generally an abelian group with trivial $G$-action. Every subgroup $H \leq G$ has $\mathrm{Fix}_T(H) = T$ (since $h \cdot t = t$ for every $h \in H$ and $t \in T$, by definition of the trivial action), so $\mathrm{Rig}_G(T)$ is the entire subgroup lattice, and $\mathrm{Rig}_G^{=1}(T)$ is empty (unless $T$ itself is a singleton). The output is fully flexible in the equivariance sense, and impossibility cannot arise from the discrete principle alone. Characterizations of $\mathbb{R}$-valued rules therefore require additional structural axioms. The typical additional axiom is additivity or linearity, and the additive refinement of Section 4 shall handle exactly this case.

#### 3.2.6 The Symmetry Gap

The five cases just examined trace out a spectrum, from the maximally rigid (strict linear orders) to the maximally flexible ($\mathbb{R}$ with trivial action), summarized in Figure 2. Every impossibility theorem is a symptom of the mismatch between the input's symmetry (large stabilizers) and the output's rigidity (small fixed sets); every characterization theorem is a symptom of a perfect match (every stabilizer meets a singleton fixed set). The rigidity spectrum is not a fine-grained classification, but a diagnostic device. Given an axiomatic problem, one first identifies the output type, locates it on the spectrum, and reads off the range of possibilities.

## 4 The Additive Refinement

The rigidity taxonomy of Section 3 handles the case of a bare $G$-action on the output. Many characterization theorems in the corpus, however, impose an additional algebraic axiom on the rule: additivity, linearity, or scaling. The Shapley value, the maximum-entropy pricing kernel, the weighted Aitchison centroid, and any other rule whose output space is a group or a vector space and whose axioms include a homomorphism condition, all fit this

pattern. For such rules, the bare fixed-set count of Theorem 2.1 is coarser than necessary. This section introduces the additive refinement, which handles the case of a $G$-invariant partial-monoid homomorphism from a $G$-monoid to an abelian group, and shows that the answer is a single hom-set out of a specific abelian group determined by the input.

### 4.1 The Grothendieck Group and the Coinvariants

Two elementary constructions from algebra are the building blocks.

Let $(X, \otimes, e)$ be a commutative cancellative partial monoid, with partial operation $\otimes$ defined on a subset $D \subseteq X \times X$ and identity element $e$.[1] The *Grothendieck group* of $X$, denoted $X^{\mathrm{gp}}$, is the free abelian group on the elements of $X$ modulo the relations $[e] = 0$ and $[x \otimes y] = [x] + [y]$ for every $(x, y) \in D$.[2] The canonical map $\iota\colon X \to X^{\mathrm{gp}}$, $x \mapsto [x]$, is a partial-monoid homomorphism (it takes $e$ to 0 and $x \otimes y$ to $[x] + [y]$ whenever defined), and it satisfies the following universal property.

**Lemma 4.1** (Universal Property of the Grothendieck Group)**.** *For any abelian group $M$, the map*

$$\mathrm{Hom}_{\mathbf{Ab}}(X^{\mathrm{gp}}, M) \longrightarrow \{\textit{partial-monoid homomorphisms } X \to M\}, \qquad \varphi \longmapsto \varphi \circ \iota,$$

*is a bijection.*

*Proof.* Every partial-monoid homomorphism $f\colon X \to M$ induces a map on the generators of the free abelian group on $X$, and this map respects the defining relations $[e] = 0$ and $[x \otimes y] = [x] + [y]$ because $f(e) = 0$ and $f(x \otimes y) = f(x) + f(y)$. It therefore factors through the quotient $X^{\mathrm{gp}}$, yielding a unique group homomorphism $\hat{f}\colon X^{\mathrm{gp}} \to M$ with $\hat{f} \circ \iota = f$. Conversely, restriction to $\iota(X)$ recovers the original map. □

Assume now that a group $G$ acts on $X$ by monoid automorphisms; that is, $G$ acts on the underlying set, preserves the domain $D$ (if $(x, y) \in D$, then $(g \cdot x, g \cdot y) \in D$), commutes with the operation ($g \cdot (x \otimes y) = (g \cdot x) \otimes (g \cdot y)$), and fixes the identity ($g \cdot e = e$). We claim that $G$ acts on $X^{\mathrm{gp}}$ by group automorphisms, via $g \cdot [x] = [g \cdot x]$. To verify well-definedness, we check that this assignment respects the defining relations of $X^{\mathrm{gp}}$. The relation $[e] = 0$ is preserved because $g \cdot [e] = [g \cdot e] = [e] = 0$, using $g \cdot e = e$. The relation $[x \otimes y] = [x] + [y]$ is preserved because $g \cdot [x \otimes y] = [g \cdot (x \otimes y)] = [(g \cdot x) \otimes (g \cdot y)] = [g \cdot x] + [g \cdot y] = g \cdot [x] + g \cdot [y]$, using the monoid-automorphism property. Hence, $g \cdot (-)$ extends to a well-defined group endomorphism of $X^{\mathrm{gp}}$, and the map $g \mapsto (a \mapsto g \cdot a)$ is a group homomorphism $G \to \mathrm{Aut}(X^{\mathrm{gp}})$, since $(gh) \cdot [x] = [gh \cdot x] = g \cdot [h \cdot x] = g \cdot (h \cdot [x])$. The *coinvariants* are

$$(X^{\mathrm{gp}})_G \coloneqq X^{\mathrm{gp}} \big/ \big\langle g \cdot a - a \mid g \in G,\, a \in X^{\mathrm{gp}} \big\rangle,$$

with canonical projection $\pi_G\colon X^{\mathrm{gp}} \to (X^{\mathrm{gp}})_G$, and the coinvariant construction has the following universal property.

[1] Cancellativity is not used in the universal-property arguments below; it is retained in the hypothesis because it is the standard condition under which the canonical map $\iota\colon X \to X^{\mathrm{gp}}$ is injective, and every partial monoid arising in our applications satisfies it. The refinement remains true for non-cancellative $X$, though the canonical map may then fail to be injective.

[2] The bracket notation $[x]$ for the class of $x$ in $X^{\mathrm{gp}}$ is standard in $K$-theory and should not be confused with the notation $[x]$ for the orbit of $x$ used in Section 2; context always disambiguates.

**Lemma 4.2** (Universal Property of Coinvariants)**.** *Let $A$ be a $\mathbb{Z}G$-module and $M$ an abelian group with the trivial $G$-action. Every $G$-equivariant homomorphism $\varphi\colon A \to M$ factors uniquely through the coinvariant quotient. That is, there is a natural bijection*

$$\mathrm{Hom}_G(A, M) \;\cong\; \mathrm{Hom}_{\mathbf{Ab}}(A_G, M).$$

*Proof.* Write $N = \langle g \cdot a - a \mid g \in G,\, a \in A\rangle$ for the subgroup of $A$ generated by all elements $g \cdot a - a$, so that $A_G = A/N$ and $\pi_G\colon A \to A_G$ is the canonical projection.

*Forward.* Let $\varphi\colon A \to M$ be a $G$-equivariant homomorphism. Since $M$ carries the trivial $G$-action, equivariance means $\varphi(g \cdot a) = \varphi(a)$ for every $g \in G$ and $a \in A$. It therefore holds that $\varphi(g \cdot a - a) = \varphi(g \cdot a) - \varphi(a) = 0$, so $N \subseteq \ker\varphi$. By the first isomorphism theorem for abelian groups, $\varphi$ factors through $A/N = A_G$; there exists a unique homomorphism $\bar{\varphi}\colon A_G \to M$ with $\bar{\varphi} \circ \pi_G = \varphi$.

*Backward.* Given a homomorphism $\psi\colon A_G \to M$, define $\varphi_\psi \coloneqq \psi \circ \pi_G\colon A \to M$. This is a group homomorphism as a composition of group homomorphisms. It is $G$-equivariant because $\varphi_\psi(g \cdot a) = \psi(\pi_G(g \cdot a)) = \psi(\pi_G(a)) = \varphi_\psi(a)$, where the second equality uses the defining property $\pi_G(g \cdot a) = \pi_G(a)$ of the coinvariant projection.

*Mutual inverses.* Starting from $\varphi$, the forward construction yields $\bar{\varphi}$ with $\bar{\varphi} \circ \pi_G = \varphi$; applying the backward construction to $\bar{\varphi}$ recovers $\bar{\varphi} \circ \pi_G = \varphi$. Starting from $\psi$, the backward construction yields $\varphi_\psi = \psi \circ \pi_G$; applying the forward construction factors $\varphi_\psi$ through $A_G$, and the unique factoring homomorphism is $\psi$ (since $\psi \circ \pi_G = \varphi_\psi$ and factorization through $\pi_G$ is unique by surjectivity of $\pi_G$). Hence, the correspondence $\varphi \leftrightarrow \bar{\varphi}$ is a bijection $\mathrm{Hom}_G(A, M) \cong \mathrm{Hom}_{\mathbf{Ab}}(A_G, M)$. □

## 4.2 The Additive Refinement

Combining the two universal properties yields the following clean identification.

**Theorem 4.3** (Additive Refinement of the Equivariance Principle)**.** *Let $(X, \otimes, e)$ be a commutative cancellative partial monoid on which a group $G$ acts by monoid automorphisms. Let $M$ be an abelian group with the trivial $G$-action. The set of $G$-invariant partial-monoid homomorphisms $R\colon X \to M$ is in natural bijection with* $\mathrm{Hom}_{\mathbf{Ab}}((X^{\mathrm{gp}})_G, M)$.

*Proof.* The two directions are as follows.

*Forward.* Let $R\colon X \to M$ be $G$-invariant, satisfying $R(e) = 0$ and $R(x \otimes y) = R(x) + R(y)$ whenever $(x, y) \in D$. By Lemma 4.1, there exists a unique $\hat{R}\colon X^{\mathrm{gp}} \to M$ with $\hat{R}([x]) = R(x)$ for every $x$. The invariance $R(g \cdot x) = R(x)$ implies $\hat{R}([g \cdot x]) = \hat{R}([x])$, i.e., $\hat{R}(g \cdot [x] - [x]) = 0$. Hence, $\hat{R}$ vanishes on the subgroup $N = \langle g \cdot a - a \mid g \in G,\, a \in X^{\mathrm{gp}}\rangle$, so $N \subseteq \ker\hat{R}$. By the first isomorphism theorem, $\hat{R}$ factors uniquely through $X^{\mathrm{gp}}/N = (X^{\mathrm{gp}})_G$; there exists a unique $\bar{R}\colon (X^{\mathrm{gp}})_G \to M$ with $\hat{R} = \bar{R} \circ \pi_G$.

*Backward.* Given $\varphi\colon (X^{\mathrm{gp}})_G \to M$, set $R_\varphi(x) \coloneqq \varphi(\pi_G([x]))$. We check the four required properties. Well-definedness is instant; for each $x \in X$, the element $\pi_G([x]) \in (X^{\mathrm{gp}})_G$ is determined by $x$, and $\varphi$ is a function. For $G$-invariance, observe that we have

$$R_\varphi(g \cdot x) = \varphi(\pi_G([g \cdot x])) = \varphi(\pi_G(g \cdot [x])) = \varphi(\pi_G([x])) = R_\varphi(x),$$

where the second equality uses functoriality of the Grothendieck construction and the third uses the defining property of coinvariants. The identity axiom holds trivially, because

$R_\varphi(e) = \varphi(\pi_G([e])) = \varphi(\pi_G(0)) = \varphi(0) = 0$. Additivity holds because, for $(x, y) \in D$, we have

$$\begin{aligned} R_\varphi(x \otimes y) &= \varphi\big(\pi_G([x \otimes y])\big) \\ &= \varphi\big(\pi_G([x] + [y])\big) \\ &= \varphi\big(\pi_G([x]) + \pi_G([y])\big) \\ &= \varphi\big(\pi_G([x])\big) + \varphi\big(\pi_G([y])\big) \\ &= R_\varphi(x) + R_\varphi(y). \end{aligned}$$

*Mutual inverses.* Starting from $R$, the forward construction yields $\bar{R}\colon (X^{\mathrm{gp}})_G \to M$ with $\bar{R}(\pi_G([x])) = R(x)$ for every $x \in X$. Applying the backward construction to $\bar{R}$ gives $R_{\bar{R}}(x) = \bar{R}(\pi_G([x])) = R(x)$, so $R_{\bar{R}} = R$. Starting from $\varphi$, the backward construction yields $R_\varphi$ with $R_\varphi(x) = \varphi(\pi_G([x]))$. We now apply the forward construction to $R_\varphi$. The universal property of $X^{\mathrm{gp}}$ yields $\widehat{R_\varphi}([x]) = R_\varphi(x) = \varphi(\pi_G([x]))$, and factoring through $\pi_G$ yields $\overline{R_\varphi}$ with $\overline{R_\varphi}(\pi_G([x])) = \varphi(\pi_G([x]))$. Since $\{\pi_G([x]) \mid x \in X\}$ generates $(X^{\mathrm{gp}})_G$ (the generators $[x]$ of $X^{\mathrm{gp}}$ map onto generators of the quotient), we conclude $\overline{R_\varphi} = \varphi$. □

The additive refinement sharpens the bare equivariance principle in the following sense. An equivariant rule that is additionally an additive homomorphism is not merely constrained to fixed sets; it is forced to be a homomorphism out of a specific abelian group determined by the input. Impossibility becomes triviality of a fixed set (as before), and characterization becomes rigidity of the coinvariant Grothendieck group $(X^{\mathrm{gp}})_G$. When $(X^{\mathrm{gp}})_G$ is generated freely by a single element, the additive equivariant rule is determined by a single scalar; when $(X^{\mathrm{gp}})_G$ is a torsion group, the additive equivariant rule may be forced to be zero.

Theorem 4.3 is the result that the earlier draft of the present work proved by categorical means, under the name Master Theorem, using an adjunction between a category of axiomatic fibrations and the category of abelian groups. Nothing in the earlier proof required this framing. The two composed universal properties do the entire work, and the factorization is illustrated in Figure 3. The refinement's mathematical content lies not in the difficulty of its proof, but in the identification of $(X^{\mathrm{gp}})_G$ as the correct invariant of a partial-monoid-with-$G$-action. This identification is what turns a bare hom-count into an explicit computation.

**Remark 4.4.** Theorem 4.3 is stated for abelian groups, i.e., $\mathbb{Z}$-modules. When $X$ is additionally an $R$-module (for a commutative ring $R$) and $G$ acts by $R$-module automorphisms, the coinvariant relations $g \cdot a - a$ are $R$-linear in $a$, so $(X^{\mathrm{gp}})_G$ inherits an $R$-module structure and the bijection of Theorem 4.3 restricts to $\mathrm{Hom}_R((X^{\mathrm{gp}})_G, M)$ by the identical proof. We shall use this $R$-linear version in the Shapley application of Section 4.3, where $R = \mathbb{R}$.

### 4.3 The Coinvariant Grothendieck Group of Symmetric Games

To exhibit the refinement in action, we compute $(X^{\mathrm{gp}})_G$ in the case that motivates the Shapley formula. The computation is short and yields a clean answer.

Fix a finite set of players $\mathcal{N}$ with $|\mathcal{N}| = n$, and let $\mathcal{G}$ denote the real vector space of cooperative games $v\colon 2^{\mathcal{N}} \to \mathbb{R}$ with $v(\emptyset) = 0$. Then $\dim_{\mathbb{R}} \mathcal{G} = 2^n - 1$, and a distinguished basis is the family of *unanimity games* $\{u_S\}_{\emptyset \neq S \subseteq \mathcal{N}}$, defined by

$$u_S(T) = \begin{cases} 1 & \text{if } S \subseteq T, \\ 0 & \text{otherwise.} \end{cases}$$

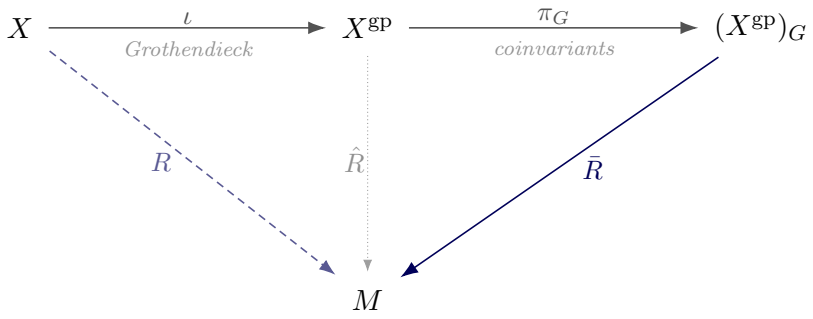



Figure 3: The two-step factorization at the heart of Theorem 4.3. A $G$-invariant partial-monoid homomorphism $R\colon X \to M$ (dashed) factors uniquely through the Grothendieck group $X^{\mathrm{gp}}$ and then through the coinvariant quotient $(X^{\mathrm{gp}})_G$, yielding a unique abelian-group homomorphism $\bar{R}\colon (X^{\mathrm{gp}})_G \to M$ (solid). The left triangle commutes by the universal property of $X^{\mathrm{gp}}$ (Lemma 4.1); the right triangle commutes by the universal property of coinvariants (Lemma 4.2). The bijection $R \leftrightarrow \bar{R}$ is the content of Theorem 4.3.

Every game $v \in \mathcal{G}$ admits the Möbius expansion $v = \sum_{\emptyset \neq S \subseteq \mathcal{N}} c_S(v)\, u_S$, where the coefficients are recovered by $c_S(v) = \sum_{T \subseteq S} (-1)^{|S|-|T|} v(T)$.

The symmetric group $\mathrm{Sym}(\mathcal{N})$ acts on $\mathcal{G}$ by $(\sigma \cdot v)(T) = v(\sigma^{-1}(T))$. On the unanimity basis, this action reads $\sigma \cdot u_S = u_{\sigma(S)}$, since $(\sigma \cdot u_S)(T) = u_S(\sigma^{-1}(T)) = 1$ if and only if $S \subseteq \sigma^{-1}(T)$, i.e., $\sigma(S) \subseteq T$. Hence, the $\mathrm{Sym}(\mathcal{N})$-orbits on the unanimity basis are indexed by the size $|S| \in \{1, 2, \ldots, n\}$, and there are exactly $n$ such orbits.

We regard $\mathcal{G}$ as a partial monoid under addition of games (the operation is total, defined on all of $\mathcal{G} \times \mathcal{G}$, and cancellative since $\mathcal{G}$ is a vector space; the "partial" qualifier in Theorem 4.3 is not restrictive here). Since $\mathcal{G}$ is already an abelian group (indeed a vector space), the Grothendieck construction yields $\mathcal{G}^{\mathrm{gp}} = \mathcal{G}$, with the identity map as the canonical inclusion. The coinvariants are

$$\mathcal{G}_{\mathrm{Sym}(\mathcal{N})} = \mathcal{G} \big/ \big\langle \sigma \cdot v - v \mid \sigma \in \mathrm{Sym}(\mathcal{N}),\, v \in \mathcal{G} \big\rangle.$$

We claim that $\mathcal{G}_{\mathrm{Sym}(\mathcal{N})}$ is a free $\mathbb{R}$-module of rank $n$, with basis the classes $\big\{[u_{\{1,2,\ldots,k\}}]\big\}_{k=1}^{n}$. Indeed, the unanimity games form a basis of $\mathcal{G}$, and the $\mathrm{Sym}(\mathcal{N})$-action permutes this basis with $n$ orbits (one per size $k$). In any free module modulo the coinvariant relations of a basis-permuting group action, the quotient is free on the orbits, since the relations $[b] - [g \cdot b]$ identify exactly the basis elements within each orbit and no relation links distinct orbits. Concretely, the relations $\sigma \cdot u_S - u_S = u_{\sigma(S)} - u_S$ identify $[u_{\sigma(S)}]$ with $[u_S]$ in the quotient, and since any two subsets of the same size are related by some permutation, all classes $[u_S]$ with $|S| = k$ coincide. No further identifications hold, because the relations do not link distinct sizes.

By the $\mathbb{R}$-linear version of Theorem 4.3 (see Remark 4.4), the set of $\mathrm{Sym}(\mathcal{N})$-invariant $\mathbb{R}$-linear maps $\mathcal{G} \to \mathbb{R}$ is in bijection with $\mathrm{Hom}_{\mathbb{R}}(\mathcal{G}_{\mathrm{Sym}(\mathcal{N})}, \mathbb{R}) \cong \mathbb{R}^n$, where the $\mathbb{R}^n$ is coordinated by the classes $\big\{[u_{\{1,2,\ldots,k\}}]\big\}_{k=1}^{n}$. Explicitly, an $\mathrm{Sym}(\mathcal{N})$-invariant additive scalar $\Phi\colon \mathcal{G} \to \mathbb{R}$ is determined precisely by the $n$ real numbers $\Phi(u_S)$ for $|S| = k \in \{1, 2, \ldots, n\}$, and any such choice extends by $\Phi(v) = \sum_S c_S(v)\, \Phi(u_S)$.

This computation determines the space of $\mathrm{Sym}(\mathcal{N})$-invariant additive scalar functionals on $\mathcal{G}$, and it supplies the per-orbit scalar content that underlies the Shapley formula. The role

of $\mathbb{R}^n$ is to record, for each orbit size $k$, the "per-orbit efficiency budget" available to a symmetric additive value. Efficiency (the axiom $\sum_i R_i(v) = v(\mathcal{N})$) is itself an invariant scalar functional, and it is picked out inside $\mathbb{R}^n$ as the vector $(1, 1, \ldots, 1)$; the equivariance principle applied to the vector-valued rule $R\colon \mathcal{G} \to \mathbb{R}^{\mathcal{N}}$ then distributes the budget within each orbit. The direct derivation, which uses the equivariance principle on unanimity games together with the null-player axiom, efficiency, and linearity, is carried out in Section 5.3, and the coinvariant computation of this subsection is invoked there as a complementary reading rather than as the primary tool.

**Remark 4.5.** The identification $\mathcal{G}_{\mathrm{Sym}(\mathcal{N})} \cong \mathbb{R}^n$ has a coordinate-free description. Any $\mathrm{Sym}(\mathcal{N})$-invariant additive map $\mathcal{G} \to \mathbb{R}$ factors through the map $v \mapsto (v_k)_{k=1}^n$, where $v_k \coloneqq \sum_{|S|=k} c_S(v)$ is the total Möbius mass of size $k$. This is the coordinate change from the unanimity basis to the "size-graded" basis, which is what the coinvariant projection $\pi_G$ singles out. Concretely, two games $v$ and $w$ have the same image in $\mathcal{G}_{\mathrm{Sym}(\mathcal{N})}$ if and only if they carry the same total Möbius mass at every coalition size, regardless of how that mass is distributed among coalitions of that size.

### 4.4 On the Content of the Refinement

The examples of Sections 5 and 7 in which the additive refinement is used are the Shapley value (worked in full in Section 5.3 from the equivariance principle, with the coinvariant computation above supplying a complementary reading), the maximum-entropy pricing kernel (treated as a pointer in Section 7.4), and the weighted Aitchison centroid for automated market makers (also a pointer in Section 7.4). In each of these cases, the equivariance principle by itself yields the correct value at fully symmetric inputs and a family of admissible values at less symmetric ones; the additive refinement then reduces the family to a hom-count out of an explicitly computable abelian group. The passage from a potentially large product of fixed sets to a single hom-set is the sharpening that additivity buys, and it is this sharpening that makes the coinvariant Grothendieck group the natural invariant of the problem.

## 5 Five Flagship Classical Instances

We now work five canonical instances in full, from the equivariance principle and its refinements. The purpose of this section is to exhibit that the mechanism identified in Section 2 handles the symmetric core of each argument uniformly, and that the residual work (the non-symmetry axioms) has, in each case, a clean and small character. In each instance, the equivariance step occupies a few lines, and the remainder of the proof is visibly orthogonal to it. The five instances are May's theorem for two-alternative majority voting, the Nash bargaining solution, the Shapley value, the anonymity-neutrality-resoluteness impossibility for Arrovian aggregation, and the constraint theorem for label-symmetric classification. The last of these lies in statistical learning theory rather than in social choice or game theory, and is included to exhibit that the mechanism operates in a distinctly different domain with the same clean division of labor. Table 1 summarizes the five instances in the format of the recipe of Section 2.5.

| Instance | Input $X$ | Output $T$ | Group $G$ | Key fixed set | Result |
|---|---|---|---|---|---|
| May | $A^{\mathcal{N}}$, $\lvert A\rvert=2$ | $A$ | $\mathrm{Sym}(\mathcal{N})\times\mathrm{Sym}(A)$ | $\{+,-\}^{(n+1)/2}$ | char.$^*$ |
| Nash | $\mathcal{S}$ | $\mathbb{R}^2$ | $\mathrm{Sym}(\{1,2\})$ | diagonal $\{(x,x)\}$ | char.$^*$ |
| Shapley | games $2^{\mathcal{N}}\to\mathbb{R}$ | $\mathbb{R}^{\mathcal{N}}$ | $\mathrm{Sym}(\mathcal{N})$ | diagonal (symm. games) | char.$^*$ |
| ANR | $\mathrm{LinOrd}(A)^{\mathcal{N}}$ | $\mathrm{LinOrd}(A)$ | $\mathrm{Sym}(\mathcal{N})\times\mathrm{Sym}(A)$ | $\varnothing$ | impossible |
| Learner | $\mathcal{D}$ (datasets) | $Y^{\mathcal{X}}$ | $\mathrm{Sym}(Y)$ | $Y_D^{\mathcal{X}}$ per orbit | char./constr. |

Table 1: The five flagship classical results, decomposed through the recipe laid out in Section 2.5. The "key fixed set" column reports the product $\prod_{[x]}\mathrm{Fix}_T(\mathrm{Stab}_G(x))$ at a distinguished orbit representative (or a single factor, where the product is trivial). Entries marked char.$^*$ denote characterization on the symmetric core, with a small additional non-equivariance axiom (positive responsiveness for May, IIA and affine rescaling for Nash, linearity plus the null-player axiom for Shapley) resolving the residual freedom. The Learner entry, marked char./constr., needs no residual axiom. Label-symmetry alone yields both the parametrizing bijection and the constraint that $L(D)(x)$ lies in the label set $Y_D$ of the training data. Each entry is derived in the corresponding subsection below.

### 5.1 May's Theorem for Two-Alternative Majority Voting

May (1952) characterizes simple majority voting on two alternatives by three axioms: anonymity, neutrality, and positive responsiveness. The theorem is the earliest and cleanest instance of the characterization side of the axiomatic method, and it is a canonical example of the equivariance principle in its "family of rules" regime.

Fix a finite set of voters $\mathcal{N}$ with $|\mathcal{N}| = n$ odd, and let $A = \{+,-\}$ denote the two alternatives in question. A profile is an element of $A^{\mathcal{N}}$, and a rule is a map $f\colon A^{\mathcal{N}}\to A$. Set $G = \mathrm{Sym}(\mathcal{N})\times\mathrm{Sym}(A)$ acting on $A^{\mathcal{N}}$ by $(\sigma,\tau)\cdot\pi = \tau\circ\pi\circ\sigma^{-1}$, and on $A$ by $(\sigma,\tau)\cdot a = \tau(a)$. Anonymity is the $\mathrm{Sym}(\mathcal{N})\times\{e\}$-invariance of $f$, and neutrality is the $\{e\}\times\mathrm{Sym}(A)$-covariance of $f$; jointly, they say that $f$ is $G$-equivariant.

The orbits of $G$ on $A^{\mathcal{N}}$ are indexed by the smaller of the two tallies $\min(n_+,n_-)$; there are $(n-1)/2+1$ orbits. To see this, observe that two profiles lie in the same $G$-orbit if and only if one can be obtained from the other by permuting voters (which preserves tallies) and possibly swapping the two alternatives (which interchanges $n_+$ and $n_-$). Therefore, the orbits are determined by the unordered pair $\{n_+,n_-\}$, which is equivalent to $\min(n_+,n_-)$. For an orbit representative $\pi$ with $n_+ > n_-$, the stabilizer $\mathrm{Stab}_G(\pi)$ consists of pairs $(\sigma,\tau)$ that preserve $\pi$. The alternative-swap $\tau = (+\,-)$ preserves $\pi$ only if it is accompanied by a voter permutation carrying $+$-voters to $-$-voters, which requires $n_+ = n_-$; since $n$ is odd this never happens, and $\mathrm{Stab}_G(\pi) = \big(\mathrm{Sym}(\pi^{-1}(+))\times\mathrm{Sym}(\pi^{-1}(-))\big)\times\{e\}$, viewed as a subgroup of $\mathrm{Sym}(\mathcal{N})\times\mathrm{Sym}(A)$. This stabilizer acts on $A$ through the $\mathrm{Sym}(A)$-component, which is trivial, so $\mathrm{Fix}_A(\mathrm{Stab}_G(\pi)) = A = \{+,-\}$. By Theorem 2.1, the set of anonymous-neutral rules on $A^{\mathcal{N}}$ is in bijection with

$$\prod_{[\pi]\in A^{\mathcal{N}}/G}\mathrm{Fix}_A(\mathrm{Stab}_G(\pi)) \;\cong\; \{+,-\}^{(n+1)/2}.$$

There are accordingly $2^{(n+1)/2}$ anonymous-neutral rules, one for each independent binary choice at each orbit. Note that because the orbits are already orbits under the full group $G = \mathrm{Sym}(\mathcal{N})\times\mathrm{Sym}(A)$, the alternative-swap is already absorbed into the orbit structure, and the choices at distinct orbits are genuinely independent. Concretely, an anonymous-

neutral rule is determined by choosing, for each value $k \in \{0, 1, \ldots, (n-1)/2\}$, whether a profile with $\min(n_+, n_-) = k$ maps to the majority alternative or the minority alternative. The two extreme members of this family are majority (always follow the more-common alternative) and anti-majority (always follow the less-common alternative); for $n \geq 3$ there are also "mixed" rules that follow the majority at some tallies and the minority at others. For example, at $n = 3$ the family has $2^2 = 4$ members, namely the two extreme rules and two mixed rules; at $n = 5$ it has $2^3 = 8$, namely the two extreme rules and six mixed rules.

May's positive-responsiveness axiom, which asserts that flipping a single vote from $-$ to $+$ when the outcome is already $+$ cannot change the outcome to $-$, eliminates every rule except majority. Indeed, positive responsiveness forces $f$ to be monotone in the tally $n_+ - n_-$; if $f(\pi) = +$ at a profile with tally difference $d$, then $f(\pi') = +$ at every profile with tally difference $d' > d$. The only monotone element of the family $\{+, -\}^{(n+1)/2}$ is the one that assigns $+$ whenever $n_+ > n_-$, which is majority. The theorem is thus a two-step reading of the principle. The equivariance product gives a family of $2^{(n+1)/2}$ rules, and an additional non-equivariance axiom (a monotonicity condition) selects the unique one. The general pattern, that equivariance narrows the rule to a parametrized family and a further non-equivariance axiom selects a unique member, is the shape taken by many of the results below.

### 5.2 Nash Bargaining Uniqueness

Nash (1950) considers two-player bargaining problems, namely convex compact subsets $S \subseteq \mathbb{R}^2_{\geq 0}$ containing the origin as a disagreement point and having a nonempty intersection with the open positive orthant. The set of admissible problems is denoted $\mathcal{S}$. A bargaining rule is a function $R\colon \mathcal{S} \to \mathbb{R}^2$ with $R(S) \in S$ for every $S$. Nash imposes four axioms.

(N1) *Pareto efficiency:* $R(S)$ lies on the Pareto frontier of $S$, i.e., there is no $y \in S$ with $y \geq R(S)$ componentwise and $y \neq R(S)$.

(N2) *Symmetry:* If $S$ is symmetric under the coordinate swap $\tau\colon (x_1, x_2) \mapsto (x_2, x_1)$, that is, $\tau(S) = S$, then $R(S)$ lies on the diagonal $\{(x, x) \mid x \in \mathbb{R}\}$.

(N3) *Independence of irrelevant alternatives (IIA):* If $S \subseteq S'$ and $R(S') \in S$, then it must hold that $R(S) = R(S')$.

(N4) *Invariance under affine rescaling:* For scalars $\alpha_1, \alpha_2 > 0$, if $A(x) = (\alpha_1 x_1, \alpha_2 x_2)$ and $A(S) \in \mathcal{S}$, then $R(A(S)) = A(R(S))$.

Nash's conclusion is that $R(S) = \arg\max_{x \in S} x_1 x_2$, the unique maximizer of the Nash product on $S$. We derive this from the equivariance principle in three steps.

**Step 1: the symmetric case.** Take $G = \{e, \tau\}$, the two-element group acting on $\mathbb{R}^2$ by the coordinate swap. For a symmetric problem $S$, the stabilizer of $S$ under $G$ is the full group, and by the equivariance principle $R(S)$ lies in $\mathrm{Fix}_{\mathbb{R}^2}(\tau) = \{(x, x) \mid x \in \mathbb{R}\}$, the diagonal. Pareto efficiency further restricts $R(S)$ to the Pareto frontier of $S$. The intersection of the Pareto frontier of a symmetric convex compact set with the diagonal is nonempty and, in the presence of a positive-orthant point, a single point. Indeed, along the diagonal, the map $x \mapsto (x, x)$ is a homeomorphism from $\mathbb{R}$ to the diagonal. Its preimage of $S$ is a closed interval $[a, b] \subseteq \mathbb{R}$ (nonempty by symmetry and the positive-orthant assumption).

The Pareto frontier of $S$ intersected with the diagonal consists of those $(x, x)$ that are not dominated by any other point of $S$, and the largest such $x$ is $b$ itself (since $(b, b) + \epsilon \cdot (1, 1) \notin S$ for small $\epsilon > 0$ by definition of $b$). Any $(x, x)$ with $x < b$ is dominated by $(b, b) \in S$, hence not Pareto-efficient. Therefore, the intersection consists of the single point $(b, b)$, which we denote $\nu(S)$. The symmetric case gives $R(S) = \nu(S)$.

**Step 2: the Nash product on general problems.** Let $S \in \mathcal{S}$ be arbitrary. The Nash product $\phi(x) = x_1 x_2$ is a strictly log-concave function on the open positive orthant (since $\log x_1 + \log x_2$ is a sum of strictly concave functions). Restricted to the compact set $S$ (which is a subset of $\mathbb{R}^2_{\geq 0}$ by hypothesis), the function $\phi$ attains its maximum. The maximum cannot occur on the boundary of $\mathbb{R}^2_{\geq 0}$, since $\phi = 0$ there while $\phi > 0$ at some interior point of $S \cap \mathbb{R}^2_{>0}$ (which is nonempty by hypothesis). Hence, the maximizer $x^*(S)$ lies in the open positive orthant. Uniqueness follows from strict log-concavity: any two distinct maximizers would give a strictly higher value at their midpoint by strict concavity of $\log \phi$, contradicting maximality.

**Step 3: reduction to the symmetric case.** We now show that $R(S) = x^*(S)$. Since $x^*(S) = (x_1^*, x_2^*)$ has both coordinates strictly positive, the scalars $\alpha_i = 1/x_i^*$ are well-defined and positive. Let $A(x) = (\alpha_1 x_1, \alpha_2 x_2)$ and $S' = A(S)$. Under this rescaling, the Nash-product maximizer becomes $A(x^*(S)) = (1, 1)$; that is, $x^*(S') = (1, 1)$. Now consider the symmetric "box"

$$B = \big\{(y_1, y_2) \in \mathbb{R}^2_{\geq 0} \mid y_1 + y_2 \leq 2\big\}.$$

We claim $S' \subseteq B$. Note that $S' = A(S) \subseteq \mathbb{R}^2_{\geq 0}$ since $A$ preserves the non-negative orthant. Take any $y = (y_1, y_2) \in S'$ with $y_1, y_2 > 0$; the case where some $y_i = 0$ follows because $B$ is closed and every such $y$ is the limit of convex combinations $\lambda y + (1 - \lambda)(1, 1) \in S' \cap \mathbb{R}^2_{>0}$ as $\lambda \to 1$, each of which satisfies $y_1 + y_2 \leq 2$. The point $(1, 1)$ maximizes the concave function $\log \phi = \log y_1 + \log y_2$ over the convex set $S' \cap \mathbb{R}^2_{>0}$, and the maximizer lies in the interior of the domain of $\log \phi$. By the first-order necessary condition for the maximum of a differentiable concave function over a convex set, the gradient inequality $\langle \nabla \log \phi(x^*),\, y - x^* \rangle \leq 0$ holds for every $y \in S' \cap \mathbb{R}^2_{>0}$. At $x^* = (1, 1)$, the gradient is $\nabla \log \phi = (1, 1)$, so the inequality reads $(y_1 - 1) + (y_2 - 1) \leq 0$, i.e., $y_1 + y_2 \leq 2$. Hence, $y \in B$, and $S' \subseteq B$ as claimed. The set $B$ is symmetric under the coordinate swap, and by Step 1 applied to $B$, we have $R(B) = \nu(B) = (1, 1)$ (the unique Pareto-diagonal point of $B$). Since $S' \subseteq B$ and $R(B) = (1, 1) \in S'$ (because $(1, 1) = x^*(S')$ is a point of $S'$ by construction), IIA yields

$$R(S') = R(B) = (1, 1).$$

By affine invariance, $R(S') = R(A(S)) = A(R(S))$, so $A(R(S)) = (1, 1)$, and inverting gives $R(S) = A^{-1}(1, 1) = x^*(S)$. This completes the derivation of Nash's characterization.

Note the division of labor. The equivariance principle enters at exactly one place, Step 1, where it forces the value at a symmetric input onto the diagonal. The remainder of the argument is the interplay of affine invariance (which selects the rescaling in which the reduction works) and IIA (which transports the value from the enclosing symmetric problem to the original). This division is representative of the pattern the equivariance principle exposes throughout the corpus.

### 5.3 The Shapley Value

Shapley (1953) considers cooperative games $v\colon 2^{\mathcal{N}} \to \mathbb{R}$ with $v(\emptyset) = 0$, and asks for a value $R\colon \mathcal{G} \to \mathbb{R}^{\mathcal{N}}$ that allocates the worth $v(\mathcal{N})$ among the $n$ players. The axioms are efficiency, symmetry, linearity, and the null-player property.

(S1) *Efficiency:* $\sum_{i\in\mathcal{N}} R_i(v) = v(\mathcal{N})$.

(S2) *Symmetry:* $R$ is $\mathrm{Sym}(\mathcal{N})$-equivariant, i.e., $R_{\sigma(i)}(\sigma \cdot v) = R_i(v)$ for every $\sigma \in \mathrm{Sym}(\mathcal{N})$, $i \in \mathcal{N}$, $v \in \mathcal{G}$.

(S3) *Linearity:* $R(v + w) = R(v) + R(w)$ and $R(cv) = cR(v)$ for every $v, w \in \mathcal{G}$ and $c \in \mathbb{R}$.

(S4) *Null player:*[3] If $i \in \mathcal{N}$ is a null player for $v$ (that is, $v(T \cup \{i\}) = v(T)$ for every $T \subseteq \mathcal{N} \setminus \{i\}$), then $R_i(v) = 0$.

We shall derive Shapley's formula from these axioms via the equivariance principle, applied on the unanimity basis of Section 4.3. The derivation has three steps.

**Step 1: the equivariance principle on unanimity games.** Fix a nonempty subset $S \subseteq \mathcal{N}$ with $|S| = k \geq 1$, and consider the unanimity game $u_S$. The stabilizer of $u_S$ under $\mathrm{Sym}(\mathcal{N})$ is the subgroup of permutations that preserve $S$ setwise; that is,

$$\mathrm{Stab}_{\mathrm{Sym}(\mathcal{N})}(u_S) = \mathrm{Sym}(S) \times \mathrm{Sym}(\mathcal{N} \setminus S),$$

where $\mathrm{Sym}(S)$ acts on $S$ and $\mathrm{Sym}(\mathcal{N} \setminus S)$ acts on $\mathcal{N} \setminus S$, and both act on $\mathbb{R}^{\mathcal{N}}$ by coordinate permutation. By the equivariance principle, $R(u_S)$ lies in $\mathrm{Fix}_{\mathbb{R}^{\mathcal{N}}}(\mathrm{Sym}(S) \times \mathrm{Sym}(\mathcal{N} \setminus S))$. A vector $(x_i)_{i\in\mathcal{N}}$ is fixed by this subgroup if and only if $x_i$ is constant on $S$ and constant on $\mathcal{N} \setminus S$; that is, $R(u_S) = a \cdot \mathbf{1}_S + b \cdot \mathbf{1}_{\mathcal{N}\setminus S}$ for some real numbers $a$ and $b$, where $\mathbf{1}_S$ denotes the indicator vector of $S$ in $\mathbb{R}^{\mathcal{N}}$.

**Step 2: determining $a$ and $b$ by the null-player axiom and efficiency.** The null-player axiom fixes $b$. For any $i \notin S$, the player $i$ is a null player for $u_S$; indeed, for any $T \subseteq \mathcal{N} \setminus \{i\}$, we have $u_S(T \cup \{i\}) = 1$ if and only if $S \subseteq T \cup \{i\}$, which since $i \notin S$ is equivalent to $S \subseteq T$, i.e., $u_S(T)$. Hence, $R_i(u_S) = 0$ for every $i \notin S$, so $b = 0$. Efficiency then fixes $a$. The sum $\sum_i R_i(u_S) = k \cdot a$ equals $u_S(\mathcal{N}) = 1$, so $a = 1/k$.

We conclude

$$R(u_S) = \frac{1}{|S|}\mathbf{1}_S \qquad (\emptyset \neq S \subseteq \mathcal{N}).$$

**Step 3: extension to arbitrary games by linearity.** Every game $v \in \mathcal{G}$ admits the Möbius expansion

$$v = \sum_{\emptyset\neq S\subseteq\mathcal{N}} c_S(v)\, u_S, \qquad c_S(v) = \sum_{T\subseteq S} (-1)^{|S|-|T|} v(T),$$

and by linearity,

$$R(v) = \sum_{\emptyset\neq S\subseteq\mathcal{N}} c_S(v)\, R(u_S) = \sum_{\emptyset\neq S\subseteq\mathcal{N}} \frac{c_S(v)}{|S|}\, \mathbf{1}_S.$$

[3] What we call a *null player* (zero marginal contribution everywhere) is sometimes called a *dummy player* in the literature. The two formulations yield the same characterization of the Shapley value in the presence of linearity and efficiency.

Extracting the $i$-th coordinate yields

$$R_i(v) = \sum_{S \ni i} \frac{c_S(v)}{|S|} = \sum_{S \ni i} \frac{1}{|S|} \sum_{T \subseteq S} (-1)^{|S|-|T|} v(T),$$

which is one standard form of the Shapley formula. Reorganizing the sum (a straightforward exchange of summations, standard in the cooperative-game literature) gives the more familiar marginal-contributions form

$$R_i(v) = \sum_{T \subseteq \mathcal{N} \setminus \{i\}} \frac{|T|!(n-|T|-1)!}{n!} \big[ v(T \cup \{i\}) - v(T) \big].$$

Two comments on the structure of the argument are in order. First, the equivariance principle enters at exactly one place, Step 1, where it forces $R(u_S)$ into the two-parameter family $a\mathbf{1}_S + b\mathbf{1}_{\mathcal{N}\setminus S}$. The remainder is the null-player axiom (which sets $b = 0$), efficiency (which sets $a = 1/|S|$), and linearity (which extends from the unanimity basis to all games). The division of labor is again clean.

Second, the coinvariant computation of Section 4.3 gives a complementary view of what is going on. The additive refinement identifies the space of $\mathrm{Sym}(\mathcal{N})$-invariant additive scalar functionals on $\mathcal{G}$ with $\mathbb{R}^n$, coordinated by orbit size $k$. Efficiency selects the specific functional $v \mapsto v(\mathcal{N})$, which corresponds to the vector $(1, 1, \ldots, 1) \in \mathbb{R}^n$ (one unit of "budget" per orbit). The equivariance principle then distributes each orbit's budget equally among its members (the $1/|S|$ factor), and the null-player axiom restricts the split to the members of $S$. Efficiency plays the role of "how much to spend per orbit," equivariance plays the role of "how to split within an orbit," and the null-player axiom plays the role of "whom to exclude from the split." Together they determine $R(u_S)$ uniquely, and linearity extends to all games.

### 5.4 Label-Symmetric Classification

The three preceding instances all draw from social choice, bargaining, and cooperative game theory, three closely related corners of theoretical economics. To exhibit the equivariance principle at work in a genuinely different domain, we now treat an instance from statistical learning theory at the same level of detail. The result is a sharp constraint on universal label-symmetric learners, and its derivation is a direct application of the recipe of Section 2.5.

**Setup.** Fix a finite set $Y$ of labels with $|Y| = k \geq 2$. Let also $\mathcal{X}$ be a nonempty covariate space. Following Feys (2026e), a *dataset* on $Y$ is defined as a finite indexed family $D = ((x_1, y_1), (x_2, y_2), \ldots, (x_N, y_N)) \in (\mathcal{X} \times Y)^N$ for some $N \geq 1$; write $\mathcal{D}$ for the set of all such datasets. The *label set of* $D$ is $Y_D \coloneqq \{y_1, y_2, \ldots, y_N\} \subseteq Y$, i.e., the set of labels that appear in $D$. A *learner* is a function $L$ assigning to each dataset $D \in \mathcal{D}$ a prediction function $L(D) \colon \mathcal{X} \to Y$; the learner is *universal* if it is defined on every $D \in \mathcal{D}$.

The group $\mathrm{Sym}(Y)$ acts on $Y$ by permutation, on $\mathcal{X} \times Y$ by $\sigma \cdot (x, y) = (x, \sigma(y))$, on datasets by $\sigma \cdot D = ((x_1, \sigma(y_1)), (x_2, \sigma(y_2)), \ldots, (x_N, \sigma(y_N)))$, and on prediction functions by $(\sigma \cdot f)(x) = \sigma(f(x))$. A learner $L$ is *label-symmetric* if it is $\mathrm{Sym}(Y)$-equivariant, that is, if $L(\sigma \cdot D) = \sigma \cdot L(D)$ for every $\sigma \in \mathrm{Sym}(Y)$ and every $D \in \mathcal{D}$. The intuitive content of label-symmetry is that the learner treats the labels as unnamed slots. Renaming the labels renames the predictions correspondingly, and no label is intrinsically preferred.

**The constraint theorem.** The equivariance principle yields the following theorem, established in Feys (2026e).

**Theorem 5.1** (Constraint Theorem)**.** *Let $L$ be a universal label-symmetric learner. For every dataset $D \in \mathcal{D}$ with $|Y_D^c| \geq 2$ (i.e., at least two labels absent from $D$) and every covariate $x \in \mathcal{X}$, we have $L(D)(x) \in Y_D$. In particular, the learner cannot predict a label outside the training data whenever at least two labels are absent.*

*Proof.* We apply the recipe of Section 2.5, arguing at a fixed but arbitrary covariate $x \in \mathcal{X}$; the componentwise $\mathrm{Sym}(Y)$-action on $Y^{\mathcal{X}}$ permits this reduction.

*Step 1: input and output spaces.* The input space is $\mathcal{D}$, and the output space, for the value at $x$, is $Y$.

*Step 2: symmetry group and stabilizer.* The group is $\mathrm{Sym}(Y)$. For a dataset $D$ with label set $Y_D$, we have

$$\mathrm{Stab}_{\mathrm{Sym}(Y)}(D) = \mathrm{Sym}(Y_D^c),$$

that is, the stabilizer is the subgroup of permutations of $Y$ that fix each label in $Y_D$ pointwise (equivalently, permute only labels in $Y_D^c$). Indeed, a permutation $\sigma$ fixes $D$ if and only if $\sigma(y_i) = y_i$ for every $i \in \{1, 2, \ldots, N\}$, which is equivalent to $\sigma$ fixing every element of $Y_D$ pointwise.

*Step 3: the fixed set on the output.* We claim that $\mathrm{Fix}_Y(\mathrm{Sym}(Y_D^c)) = Y_D$. For the inclusion $Y_D \subseteq \mathrm{Fix}_Y(\mathrm{Sym}(Y_D^c))$, every $y \in Y_D$ is fixed by every $\sigma \in \mathrm{Sym}(Y_D^c)$, since $\mathrm{Sym}(Y_D^c)$ permutes only labels in $Y_D^c$ and fixes $Y_D$ pointwise by definition. For the reverse inclusion $\mathrm{Fix}_Y(\mathrm{Sym}(Y_D^c)) \subseteq Y_D$, suppose $y \in Y_D^c$. Since $|Y_D^c| \geq 2$ by hypothesis, there exists $y' \in Y_D^c$ with $y' \neq y$, and the transposition $(y\ y') \in \mathrm{Sym}(Y_D^c)$ sends $y$ to $y' \neq y$, so $y$ is not fixed and therefore $y \notin \mathrm{Fix}_Y(\mathrm{Sym}(Y_D^c))$.

*Step 4: reading off the conclusion.* By Theorem 2.1 applied to the value at $x$, we have $L(D)(x) \in \mathrm{Fix}_Y(\mathrm{Stab}_{\mathrm{Sym}(Y)}(D)) = \mathrm{Fix}_Y(\mathrm{Sym}(Y_D^c)) = Y_D$. Since $x \in \mathcal{X}$ was arbitrary, we conclude that $L(D)(x) \in Y_D$ for every covariate $x$. □

The equivariance principle enters at exactly one place, Step 3, where it identifies the fixed set as $Y_D$. The remainder of the argument is the mechanical setup of the recipe. Note that no non-symmetry axiom is required.

**The characterization bijection.** The constraint theorem is the negative side of the story, and it says what a universal label-symmetric learner cannot do. The positive side is the natural bijection produced directly by Theorem 2.1, which parametrizes all universal label-symmetric learners. Indeed, fix a fundamental domain $\Sigma \subseteq \mathcal{D}$ for the $\mathrm{Sym}(Y)$-action on $\mathcal{D}$; the orbits of $\mathrm{Sym}(Y)$ on $\mathcal{D}$ are indexed by the isomorphism type of the labeled dataset (equivalently, by the pattern of covariates coupled with the size of $Y_D$). For a representative $D \in \Sigma$ with label set $Y_D$, the stabilizer is $\mathrm{Sym}(Y_D^c)$, and the fixed set on $Y^{\mathcal{X}}$ is $Y_D^{\mathcal{X}}$, since the $\mathrm{Sym}(Y)$-action on $Y^{\mathcal{X}}$ is pointwise in $x$ and the fixed set at each $x$ is $Y_D$ independently. Theorem 2.1 then yields

$$\{\text{universal label-symmetric learners}\} \longleftrightarrow \prod_{[D] \in \mathcal{D}/\mathrm{Sym}(Y)} Y_D^{\mathcal{X}}.$$

When $\mathcal{X}$ is finite and the enumeration is worked through, this reduces to a closed-form count of universal label-symmetric learners on $N$-point datasets involving Bell and Stirling numbers; the details are carried out in Feys (2026e).

**On the division of labor.** The label-symmetric classification instance is unusually clean. Label-symmetry alone, with no monotonicity, no efficiency, no linearity, and no null-player axiom, yields both the constraint theorem (Theorem 5.1) and the characterization bijection above. It is the extreme case of the pattern in which the equivariance principle does all of the work. The ANR impossibility of Section 5.5, treated next, has the same character on the impossibility side. May, Nash, and Shapley each require one further non-symmetry axiom to close the argument (positive responsiveness, IIA together with affine rescaling, and linearity together with the null-player axiom and efficiency, respectively). The cross-domain point is that the same recipe, unchanged in shape and unchanged in content, applies as smoothly in a machine-learning setting as it does in classical social choice.

### 5.5 The ANR Impossibility

Arrow (1963) considers preference-aggregation rules $R\colon \mathrm{LinOrd}(A)^{\mathcal{N}} \to \mathrm{LinOrd}(A)$, and imposes independence of irrelevant alternatives, unanimity, and non-dictatorship. In the anonymity-neutrality-resoluteness variant, treated in the tradition of Moulin (1983) and Bubboloni and Gori (2014) and further developed in Feys (2026g), the axioms are replaced by anonymity ($\mathrm{Sym}(\mathcal{N})$-invariance of the rule under voter permutation), neutrality ($\mathrm{Sym}(A)$-equivariance under alternative permutation), and resoluteness (values in $\mathrm{LinOrd}(A)$ rather than in weak orders).

Take $|A| = m$ and $|\mathcal{N}| = n$ with $m \mid n$. Write $A = \{a_0, a_1, \dots, a_{m-1}\}$, fix the cyclic permutation $c = (a_0\ a_1\ \cdots\ a_{m-1}) \in \mathrm{Sym}(A)$ of order $m$, and fix the reference order $a_0 <_A a_1 <_A \cdots <_A a_{m-1}$. Partition $\mathcal{N}$ into $q := n/m$ groups $B_1, B_2, \dots, B_q$ of size $m$. Within each group $B_s$, enumerate the voters as $B_s = \{v_0^s, v_1^s, \dots, v_{m-1}^s\}$, and define the Condorcet profile $x_\Delta \in \mathrm{LinOrd}(A)^{\mathcal{N}}$ by assigning to voter $v_j^s$ the order $c^j(<_A)$, the $j$-th cyclic shift of the reference order. We shall show that the stabilizer of $x_\Delta$ under $G = \mathrm{Sym}(\mathcal{N}) \times \mathrm{Sym}(A)$ projects onto the cyclic subgroup $C_m = \langle c \rangle \le \mathrm{Sym}(A)$.

A pair $(\sigma, \tau) \in \mathrm{Sym}(\mathcal{N}) \times \mathrm{Sym}(A)$ stabilizes $x_\Delta$ if and only if $\tau \circ x_\Delta(\sigma^{-1}(v)) = x_\Delta(v)$ for every voter $v \in \mathcal{N}$, i.e., the profile obtained by permuting voters by $\sigma$ and relabeling alternatives by $\tau$ equals the original. We claim that for each $k \in \{0, 1, \dots, m-1\}$, the pair $(\sigma_k, c^k)$ stabilizes $x_\Delta$, where $\sigma_k$ is the voter permutation that cyclically shifts each group by $k$ positions; that is, $\sigma_k(v_j^s) = v_{j+k \bmod m}^s$ for every $s$ and $j$. Indeed, voter $\sigma_k^{-1}(v_j^s) = v_{j-k \bmod m}^s$ carries the order $c^{j-k \bmod m}(<_A)$, and applying $c^k$ yields $c^k \circ c^{j-k}(<_A) = c^j(<_A) = x_\Delta(v_j^s)$, which is the required identity. Hence, the projection of $\mathrm{Stab}_G(x_\Delta)$ onto $\mathrm{Sym}(A)$ contains the full cyclic subgroup $C_m = \langle c \rangle$. For any $(\sigma, \tau) \in \mathrm{Stab}_G(x_\Delta)$, equivariance of $R$ gives $\tau \cdot R(x_\Delta) = R((\sigma, \tau) \cdot x_\Delta) = R(x_\Delta)$, so $R(x_\Delta)$ is fixed by every $\tau$ in the $\mathrm{Sym}(A)$-projection of $\mathrm{Stab}_G(x_\Delta)$, and in particular by every element of $C_m$. The action of $C_m$ on $\mathrm{LinOrd}(A)$ is free, since no nontrivial cyclic shift of the enumeration $a_0 <_A \cdots <_A a_{m-1}$ preserves the order, by the analysis of Section 3.2. Hence, $\mathrm{Fix}_{\mathrm{LinOrd}(A)}(C_m) = \emptyset$, and impossibility follows from Corollary 2.2.

**Remark 5.2.** The divisibility hypothesis $m \mid n$ is a sufficient condition for the construction of the Condorcet profile $x_\Delta$. The ANR impossibility holds more broadly (for instance, for any $m \ge 3$ and $n \ge 2$ under appropriate axiom sets; see Moulin (1983) and Bubboloni and Gori (2014)). The point of the derivation above is not to recover the sharpest known sufficient condition, but to exhibit the equivariance mechanism cleanly at a single well-chosen profile: one whose stabilizer projects onto a cyclic subgroup acting freely on $\mathrm{LinOrd}(A)$, making the impossibility immediate.

The topological impossibility theorem of Chichilnisky (1980) and its refinements by Baryshnikov (1993) and Weinberger (2004) treat the continuous version of the same mechanism. The discrete-continuous split is precisely the split between Section 2 and the topological framework, and we return to it in Section 8.

# 6 A Cold Application of the Recipe

The five instances of Section 5 are retrospective applications of the equivariance principle to theorems whose statements were known in advance. To exhibit the recipe as a diagnostic tool rather than a retrospective organizer, we now apply it to a small axiomatic problem whose answer we do not presuppose, and derive a characterization from scratch.

## 6.1 The Uniform-Random-Choice Selector

The problem is the following. Fix a countably infinite universe $\mathcal{N}$ of items (agents, alternatives, options; the interpretation is immaterial), and let $\mathcal{F}$ denote the set of finite nonempty subsets of $\mathcal{N}$. A *selector* is a function $\rho\colon \mathcal{F} \to \bigsqcup_{S\in\mathcal{F}} \Delta(S)$ that assigns to each finite "menu" $S \in \mathcal{F}$ a probability distribution $\rho(S)$ supported on $S$. The symmetric group $\mathrm{Sym}(\mathcal{N})$ acts on $\mathcal{F}$ by $\sigma \cdot S = \sigma(S)$ and on distributions by pushforward, so $(\sigma \cdot \mu)(i) = \mu(\sigma^{-1}(i))$ for $\mu \in \Delta(S)$ and $i \in \sigma(S)$. The single axiom we impose is anonymity.

*Anonymity:* For every $\sigma \in \mathrm{Sym}(\mathcal{N})$ and every $S \in \mathcal{F}$, we have $\rho(\sigma \cdot S) = \sigma \cdot \rho(S)$.

The question is which selectors satisfy anonymity, and whether there is a unique one.

We apply the four-step recipe of Section 2.5, in the bundled form furnished by Corollary 2.4, since the output space is genuinely fibered over the input.

**Step 1: input and output spaces.** The input space is $X = \mathcal{F}$. The output space is the total space $Y = \bigsqcup_{S\in\mathcal{F}} \Delta(S)$ of the bundle $p\colon Y \to X$ whose fiber over $S$ is $p^{-1}(S) = \Delta(S)$, and a selector is a section of $p$.

**Step 2: symmetry group and actions.** The group is $G = \mathrm{Sym}(\mathcal{N})$, acting on $\mathcal{F}$ by $\sigma \cdot S = \sigma(S)$ and on $Y$ by the pushforward action just described, which sends a distribution $\mu \in \Delta(S)$ to $\sigma \cdot \mu \in \Delta(\sigma(S))$. The projection $p$ is $G$-equivariant by construction, since $p(\sigma \cdot \mu) = \sigma(S) = \sigma \cdot p(\mu)$ for $\mu \in \Delta(S)$.

**Step 3: fundamental domain and stabilizers.** A fundamental domain for the $\mathrm{Sym}(\mathcal{N})$-action on $\mathcal{F}$ is obtained by choosing, for each integer $k \geq 1$, the canonical representative $S_k = \{1, 2, \ldots, k\}$. Indeed, the orbits of $\mathrm{Sym}(\mathcal{N})$ on $\mathcal{F}$ are indexed by cardinality $k = |S|$, since any bijection between two $k$-element subsets of $\mathcal{N}$ extends to a permutation of $\mathcal{N}$. The stabilizer of $S_k$ is

$$\mathrm{Stab}_{\mathrm{Sym}(\mathcal{N})}(S_k) = \mathrm{Sym}(S_k) \times \mathrm{Sym}(\mathcal{N} \setminus S_k),$$

the subgroup of permutations that preserve $S_k$ setwise.

**Step 4: fixed sets on the output.** By Corollary 2.4, the relevant fixed set at $S_k$ is $\mathrm{Fix}_{\Delta(S_k)}(\mathrm{Sym}(S_k) \times \mathrm{Sym}(\mathcal{N} \setminus S_k))$. The second factor $\mathrm{Sym}(\mathcal{N} \setminus S_k)$ acts trivially on $\Delta(S_k)$, since it moves only elements outside $S_k$ and $\Delta(S_k)$ is supported on $S_k$. Hence, the

fixed set reduces to $\mathrm{Fix}_{\Delta(S_k)}(\mathrm{Sym}(S_k))$, the set of $\mathrm{Sym}(S_k)$-invariant distributions on $S_k$. By the taxonomy of Section 3.2, the action of $\mathrm{Sym}(S_k)$ on $S_k$ is transitive, so the unique $\mathrm{Sym}(S_k)$-invariant distribution on $S_k$ is the uniform distribution $U_{S_k}(i) = 1/k$ for every $i \in S_k$.

**Reading off the conclusion.** Every fixed set is a singleton, so Corollary 2.4 yields a unique anonymous selector, namely the uniform-random-choice rule

$$\rho(S)(i) = \frac{1}{|S|} \qquad (i \in S).$$

We state the conclusion as a proposition.

**Proposition 6.1** (Uniform-Random-Choice Characterization)**.** *The unique anonymous selector $\rho\colon \mathcal{F} \to \bigsqcup_{S\in\mathcal{F}} \Delta(S)$ is the uniform-random-choice rule that assigns to every finite nonempty $S \in \mathcal{F}$ the uniform distribution $U_S$ on $S$.*

*Proof.* Immediate from the four-step application of the recipe above and Corollary 2.4. □

The interest of the proposition is not the result (uniform-random-choice is folklore) but the derivation. No axiom beyond anonymity was imposed. The recipe delivered the answer in one line per step, without any prior knowledge of the target. Given a new axiomatic problem, the same recipe applies without adjustment.

**Remark 6.2.** The reader may object that the problem was chosen to make the recipe look effective. This objection is partially just: the input and output spaces above are unusually clean, and the group action is transparent. A less transparent problem might require more work at Step 4. The claim, however, is not that Step 4 is always trivial; it is that the recipe reduces every discrete equivariance question to Step 4. For problems where Step 4 is genuinely hard, the recipe at least isolates the difficulty and prevents it from being confused with the symmetric setup, which is often the source of much confusion in the corpus.

# 7 Further Instances

The following section works, more briefly, seven further instances of the equivariance principle. The purpose is to make plain, by cumulative example, that the mechanism identified in Section 2 operates uniformly across a wide range of axiomatic problems. Three instances are worked from the principle directly (bankruptcy, random assignment, and the nudge trilemma), and four are treated as pointers to companion work in which the details are found (the orbit-partition tie-breaker, the uniform-completion rule, the maximum-entropy pricing kernel, and the fairness impossibility for automated market makers). The pointer treatment reflects the fact that these four are the topic of separate working papers in which the residual non-symmetry work is carried out in full.

## 7.1 Bankruptcy Problems and Equal Treatment of Equals

Consider a bankruptcy problem in the sense of O'Neill (1982) and Aumann and Maschler (1985), namely a triple $(E, \mathcal{N}, c)$ with $E \geq 0$ an estate, $\mathcal{N}$ a finite set of claimants with $|\mathcal{N}| = n$, and $c \in \mathbb{R}^{\mathcal{N}}_{\geq 0}$ a claim vector satisfying $\sum_{i\in\mathcal{N}} c_i \geq E$. A division rule assigns to each such problem an allocation $R(E, \mathcal{N}, c) \in \mathbb{R}^{\mathcal{N}}_{\geq 0}$ satisfying the efficiency condition

$\sum_{i\in\mathcal{N}} R_i(E,\mathcal{N},c) = E$ and the individual-bounds condition $0 \le R_i \le c_i$ for every $i$. The axiom of equal treatment of equals demands $\mathrm{Sym}(\mathcal{N})$-equivariance of $R$ under joint permutation of the claimants and their claims, and is universally imposed across the family of rules studied in the literature, including the proportional rule, the constrained equal-awards rule, the constrained equal-losses rule, and the Talmud rule of Aumann and Maschler (1985).

Take a symmetric problem, that is, one with $c_1 = c_2 = \cdots = c_n$. Its stabilizer under $\mathrm{Sym}(\mathcal{N})$ is the full symmetric group, and the fixed set of $\mathrm{Sym}(\mathcal{N})$ acting on $\mathbb{R}^{\mathcal{N}}$ by coordinate permutation is the diagonal $\{(t,t,\ldots,t) \mid t \in \mathbb{R}\}$, as computed in Section 5.3. Efficiency collapses the diagonal to the single point $(E/n, E/n, \ldots, E/n)$, and by Corollary 2.3 equal division is the unique rule on symmetric problems. No further axiom is required for this conclusion. The stronger characterizations of specific rules (the Talmud rule by consistency in the sense of Aumann and Maschler (1985), the constrained equal-awards rule by claim-truncation invariance in the tradition of Thomson (2003)) select a particular extension of equal division to asymmetric problems, and each such characterization proceeds by a further $\mathrm{Sym}(\mathcal{N}')$-equivariance argument at a partially-symmetric input.

The instance illustrates a general pattern. An axiom of the form "equal treatment of equals" or "symmetric agents receive symmetric outcomes" is a $\mathrm{Sym}(\mathcal{N})$-equivariance requirement in disguise, and the equivariance principle immediately reduces the analysis at any symmetric input to a fixed-set computation on the output side. Efficiency, additivity, or consistency then supplies the further axioms that resolve the remaining parametric freedom at asymmetric inputs.

### 7.2 Random Assignment at Symmetric Preference Profiles

Consider the house allocation problem of Hylland and Zeckhauser (1979) in the framework of Bogomolnaia and Moulin (2001), namely a finite set of agents $\mathcal{N}$ with $|\mathcal{N}| = n$, a finite set of objects $\mathcal{O}$ with $|\mathcal{O}| = n$, and a profile $\succ = (\succ_i)_{i\in\mathcal{N}}$ of strict preferences of the agents over the objects. A random assignment rule outputs a bistochastic matrix $P = P(\succ) \in [0,1]^{\mathcal{N}\times\mathcal{O}}$ with $\sum_j P_{ij} = 1$ for every agent $i$ and $\sum_i P_{ij} = 1$ for every object $j$; the entry $P_{ij}$ is interpreted as the probability that agent $i$ receives object $j$. Two well-studied rules on this domain are random serial dictatorship and the probabilistic serial rule of Bogomolnaia and Moulin (2001).

Impose equal treatment of equals in its natural form, namely $\mathrm{Sym}(\mathcal{N})$-equivariance of the rule under joint permutation of agents and their preference labels. That is, if $\sigma \in \mathrm{Sym}(\mathcal{N})$ and $\sigma\cdot\succ$ denotes the profile with $(\sigma\cdot\succ)_i = \succ_{\sigma^{-1}(i)}$, then $P(\sigma\cdot\succ)_{ij} = P(\succ)_{\sigma^{-1}(i),j}$ for every $i,j$. Consider a fully symmetric profile in which every agent has the same preference $\succ^*$ over $\mathcal{O}$. The stabilizer of such a profile under $\mathrm{Sym}(\mathcal{N})$ is the full symmetric group, since permuting agents leaves the joint profile invariant when all agents have identical preferences. The fixed set of $\mathrm{Sym}(\mathcal{N})$ acting on the space of bistochastic matrices by row permutation consists of the matrices whose rows are all equal; a matrix $P$ with $P_{ij} = P_{\sigma^{-1}(i),j}$ for every $\sigma \in \mathrm{Sym}(\mathcal{N})$ has $P_{ij}$ independent of $i$, hence all rows equal. The column-sum-one constraint then forces each row of $P$ to be the uniform distribution $(1/n, 1/n, \ldots, 1/n)$ on $\mathcal{O}$. By Corollary 2.3, uniform random assignment is the unique equal-treatment-of-equals rule on the fully symmetric profile.

The same argument, restricted to partial symmetries of the profile, delivers partial forcing at less symmetric inputs. At a profile in which two agents share the same preferences but a third differs, the residual stabilizer $\mathrm{Sym}(\{1,2\}) \times \{e\}$ forces the two symmetric

rows of $P$ to coincide, though the third row is not determined by equivariance alone. Further characterizations of specific rules (random serial dictatorship by strategy-proofness, the probabilistic serial rule by ordinal efficiency in the sense of Bogomolnaia and Moulin (2001)) require additional axioms and further equivariance arguments at generic profiles, but the discrete backbone at each symmetric profile is the fixed-set computation just carried out.

### 7.3 The Nudge Trilemma

Feys (2026f) considers frame-selection rules for behavioral nudges, taking a binding choice situation $S_0$ to a frame $g(S_0) = (\varphi, r)$ consisting of a channel $\varphi$ and a reason $r$ drawn from a binary reason space $\{r_0, r_1\}$. The axiom of absence of reference (AR) imposes $C_2$-equivariance under the swap $r_0 \leftrightarrow r_1$, with the swap acting trivially on the channel and on $S_0$.

Since $S_0$ is $C_2$-fixed, the stabilizer of $S_0$ is the full group $C_2$, and by the equivariance principle $g(S_0)$ must be $C_2$-fixed. The reason component $\{r_0, r_1\}$ carries a free $C_2$-action (the nontrivial element sends $r_0 \mapsto r_1$ and $r_1 \mapsto r_0$), and this action has no fixed point, since neither $r_0$ nor $r_1$ is mapped to itself by the swap. The fixed set of $C_2$ on the full frame space is therefore empty (a fixed point $(\varphi, r)$ would require $r$ to be fixed by the swap, which is impossible), and no such $g$ exists by Corollary 2.2. The impossibility here is precisely the two-line example of Section 2.4 recast into the behavioral setting. The nudge trilemma is an instance of the equivariance principle at its smallest.

### 7.4 Pointers to Companion Work

The following four instances are the subject of separate working papers, in which both the equivariance step and the residual non-symmetry arguments are carried out in full. We record here only the equivariance step for each, and refer to the companion papers for the remainder. The four are grouped together to make explicit that they follow one and the same pattern: at a distinguished symmetric input, the equivariance principle collapses the output into a small fixed set (a uniform distribution, an orbit partition, or the empty set); the residual axioms then act on the remaining freedom.

**The orbit partition as the unique honest tie-breaker.** Feys (2026a) considers partition-valued tie-breakers on inputs consisting of a weak order $\succeq$ on $\mathcal{N}$ and an auxiliary information item $h$ drawn from a $\mathrm{Sym}(\mathcal{N})$-set. For any input $(h, \succeq)$, the joint stabilizer $H = \mathrm{Stab}_{\mathrm{Sym}(\mathcal{N})}(h, \succeq)$ acts on $\mathcal{N}$ with a canonical orbit partition $\Omega(h, \succeq)$, which is $H$-fixed by construction. The equivariance principle forces the output into the $H$-fixed, $\succeq$-consistent partitions of $\mathcal{N}$. The residual work, which shows that $\Omega(h, \succeq)$ is the unique such partition under a rank-respecting invariance condition, is Theorem B of Feys (2026a).

**The uniform-completion rule.** Feys (2026b) considers probabilistic tie-breakers, in which the rule outputs a probability distribution over strict linear orders refining a given weak order. A symmetry axiom on the completion group $\Gamma(\Omega, \succeq)$ forces the rule to be $\Gamma$-invariant. The completion group acts freely and transitively on the finite set of admissible completions, and by Section 3.2 the unique $\Gamma$-invariant distribution on a transitive $\Gamma$-orbit is uniform. The uniform-completion rule is therefore forced by Corollary 2.3. This is the observation that underlies Laplace's principle of insufficient reason, sortition, random serial

dictatorship in matching, and the equal-a-priori postulate of statistical mechanics; each is an instance of "the unique symmetric distribution on a transitive orbit is uniform."

**The maximum-entropy pricing kernel.** Feys (2026c) considers pricing-measure selectors on a finite state space $\Omega$, taking a market $\mathcal{X}$ to a positive probability measure $\rho(\mathcal{X}) \in \Delta(\Omega)$ consistent with $\mathcal{X}$. Anonymity is $\mathrm{Sym}(\Omega)$-equivariance, and independent composition is a partial-monoid additivity condition on markets. On the trivial market $\mathcal{X}_0$, the stabilizer under $\mathrm{Sym}(\Omega)$ is the full group, and by Section 3.2 the unique $\mathrm{Sym}(\Omega)$-fixed element of $\Delta(\Omega)$ is the uniform distribution, forcing $\rho(\mathcal{X}_0) = U_\Omega$. Independent composition, in the additive-refinement setting of Section 4, propagates this identification along partial-monoid factorization. Continuity closes the argument on nontrivial markets, and the unique compatible rule is the maximum-entropy pricing measure under the uniform reference.

**Fairness and strategy-proofness in automated market makers.** Feys (2026d) considers aggregation rules for automated market makers on the weighted-product family, taking a profile of liquidity-provider preferred weight vectors $(\alpha^*_\ell)_{\ell=1}^m$ on the asset simplex $\Delta(A)$ to an aggregate weight vector on $\Delta(A)$, with $|A| = n$. Under the Arrovian fairness cluster, the aggregate at the fully symmetric profile is forced to be the unique $\mathrm{Sym}(m) \times \mathrm{Sym}(A)$-fixed element of $\Delta(A)$, namely the uniform distribution. A one-parameter perturbation of provider 1's report toward a distinguished asset breaks the $\mathrm{Sym}(m)$-symmetry but preserves a residual $\mathrm{Sym}(m-1) \times \mathrm{Sym}(n-1)$-symmetry (fixing the perturbed provider and the distinguished asset), and at this perturbed input the fairness cluster forces the aggregate to be the weighted Aitchison centroid. Strategy-proofness against a further perturbation is incompatible with the resulting geometric-mean form for $n \geq 3$ and $m \geq 3$, and the impossibility is an instance of Corollary 2.2 at the perturbed stabilizer. At $n = 2$, the fixed set of the residual stabilizer on the two-simplex is a single point, and a fair strategy-proof rule exists. The boundary case is exactly the case in which the sharp rigidity of Section 3.2 is met.

# 8 Where the Principle Does Not Apply

The equivariance principle governs impossibilities and characterizations whose obstruction is concentrated at a stabilizer-fixed input, and whose axioms are, at least at the relevant input, expressible as $G$-equivariance for some group action. Both of these hypotheses admit natural failures. The scope of the principle is therefore delimited, and we spell out the boundary explicitly.

## 8.1 Topological Obstructions

The first family of results outside the discrete framework is the class of topological impossibility theorems in the tradition of Chichilnisky (1980) and Baryshnikov (1993). Chichilnisky's theorem asserts that no continuous, anonymous, unanimous aggregation rule exists on the space of continuous preferences over a positive-dimensional alternative space. The obstruction here is topological rather than combinatorial. The failure is not concentrated at a single symmetric profile, but is distributed over a topological space of profiles, and is detected by a nontrivial homotopy or cohomology class of the aggregation map. Baryshnikov's cohomological refinement makes the obstruction explicit as a nontrivial

class in the singular cohomology of a configuration space, and Weinberger (2004) systematizes the equivariant-topology framework. The equivariance principle of Section 2 captures the discrete shadow of this obstruction, namely the impossibility for finite alternative sets, but the continuous version requires substantial additional apparatus that we do not treat here.

### 8.2 Quantitative and Approximate Equivariance

The second family is the class of quantitative impossibility theorems, in which the impossibility is not that no rule exists but that no rule performs well against a specified benchmark. The Robust Halfspace Characterization of Feys (2026i) and the Gibbard-Satterthwaite impossibility for pricing kernels of Feys (2026h) are of this type. The mechanism in these theorems is still equivariance, but the fixed-set argument yields only a weaker conclusion (the rule is close to a fixed point, not equal to it), and additional analytic machinery (Fourier analysis on the discrete cube, entropic estimates on the simplex) is needed to close the gap. The equivariance principle is the discrete backbone of these quantitative results, but it is not the entire result.

### 8.3 Non-Symmetry Axioms

The scope of the principle is delimited on the axiom side as well. Pareto efficiency, positive responsiveness, incentive compatibility, and the ordinary form of monotonicity are constraints on the values of the rule that are not expressible as invariance under any natural group action. Independence of irrelevant alternatives, in the form used by Arrow (1963), admits a partial symmetry reading (invariance under changes to the profile that leave the restriction to a pair of alternatives fixed) but is not strictly a group-equivariance condition.

When such axioms are present, the equivariance principle typically supplies a first-pass reduction of the following form. The rule is constrained to lie in a small family parametrized by the fixed-set product, and the non-symmetry axioms then select one element of the family or produce a further impossibility. This is the pattern observed in the May instance of Section 5.1, in which equivariance narrows the admissible rules to a family of $2^{(n+1)/2}$ anonymous-neutral rules, and positive responsiveness selects majority as the unique member. The Nash instance of Section 5.2 works similarly. Equivariance forces the value onto the diagonal for symmetric problems, and IIA together with affine rescaling reduces the general case to the symmetric one. The Shapley instance of Section 5.3 works similarly again. Equivariance forces the value onto a two-parameter family on unanimity games, and the null-player axiom plus efficiency fix the two parameters. The equivariance principle is, in this sense, a piece of an argument rather than a complete argument, and the non-symmetry axioms are the piece that fills the remainder.

### 8.4 Domain Restrictions and Partial Group Actions

The principle presumes that the input space $X$ carries a full $G$-action, meaning that every element of $G$ takes $X$ to $X$. In many settings the natural symmetry group acts only partially on the domain of interest: single-peaked preference domains are not stable under arbitrary permutations of the alternatives, and restricted domains arising from network structure, typed agents, or combinatorial constraints raise the same issue. The correct framework in such cases replaces the equivariance principle with an analogue for partial group actions or groupoids, and we do not treat this generalization here; it belongs to a natural sequel.

### 8.5 Interaction of Multiple Group Actions

Finally, when several distinct symmetries are imposed and they do not naturally combine into a single group action, the equivariance principle applies to the largest single group action available, and the residual axioms operate as auxiliary constraints. Anonymity of voters and neutrality of alternatives, for instance, combine multiplicatively as $\mathrm{Sym}(\mathcal{N}) \times \mathrm{Sym}(A)$; the interaction of anonymity with independence of irrelevant alternatives is not of this form. The Arrow-ANR analysis of Section 5.5 illustrates this compromise. The equivariance handles anonymity and neutrality jointly and produces the impossibility, whereas the classical Arrow theorem in its independence-of-irrelevant-alternatives formulation requires additional non-symmetry apparatus.

### 8.6 Between the Regimes

Between these several families and the setting of Section 2 lies a middle ground in which the equivariance principle applies, but with additional algebraic structure on the fixed sets that is not visible in the bare setting. This is where the additive refinement of Section 4 intervenes. When the output space carries an algebraic structure (an abelian group, a simplex, a vector space) and the rule is required to respect it, the fixed-set product of Theorem 2.1 refines to a homomorphism count out of the coinvariant Grothendieck group. The three families (discrete equivariance, quantitative equivariance, topological equivariance) form a natural hierarchy of impossibility mechanisms, with the discrete version of Section 2 as its combinatorial base. The value of the discrete principle comes precisely from its position at the base. Any statement of a discrete impossibility or characterization theorem can be checked against the principle in a few lines, and the principle serves as a first line of organization before the heavier machinery is brought to bear.

## 9 Concluding Remarks

The equivariance principle is elementary and its proof is short. Its interest lies not in the difficulty of its proof but in the range of theorems it organizes and in the diagnostic recipe it makes available.

The interest is, in this sense, structural. The bare bijection of Theorem 2.1 restates a fact well-known in equivariant topology and does not, by itself, add a new theorem to the classical axiomatic corpus. The framework built around it, however, does add several. The rigidity taxonomy of Section 3 classifies the output-side geometry that governs each impossibility and characterization. The additive refinement of Section 4 identifies $(X^{\mathrm{gp}})_G$ as the correct invariant of a partial-monoid-with-$G$-action and yields the hom-set classification of Theorem 4.3. The coinvariant computation of Section 4.3 determines the space of symmetric additive scalar functionals on cooperative games. The uniform-random-choice characterization of Proposition 6.1, though modest, is a genuine consequence of the recipe applied to a problem not previously in the literature. Encompassing these individual results, the unification of the impossibility side and the characterization side of the corpus as one product-of-fixed-sets identification, together with the taxonomy explaining when each regime prevails, is the paper's most durable structural contribution.

We have exhibited five flagship classical instances (May's theorem, Nash bargaining, the Shapley value, the ANR impossibility, and the constraint theorem for label-symmetric classification), each worked in full from the equivariance principle and its additive refinement,

together with a cold application of the recipe to a small novel problem and a gallery of seven further instances. In every case, the equivariance step is short and the residual non-symmetry work is either small (positive responsiveness for May, the null-player axiom plus efficiency for Shapley) or a clean rescaling-plus-IIA closure (Nash bargaining) or absent altogether (the ANR impossibility and the label-symmetric constraint theorem, in each of which equivariance alone suffices). The refinements developed here (the rigidity taxonomy of Section 3 and the additive refinement of Section 4) turn the bare principle into a working tool. The additive refinement has been exercised explicitly in Section 4.3, where the coinvariant Grothendieck group of symmetric cooperative games is computed and supplies the per-orbit scalar content of the Shapley formula, complementing the equivariance derivation of Section 5.3.

Two conceptual points deserve emphasis. First, the impossibility-characterization dichotomy is not two theorems but one. Both are readings of the same fixed-set product, distinguished only by the cardinality of the factors. Empty factors are impossibilities, singleton factors are characterizations, and larger factors are parametrized families. The three regimes shade smoothly into each other, and the passage from one to another is controlled by the geometry of the output space.

Second, the rigidity taxonomy of Section 3 explains why particular output spaces exhibit particular behaviors under equivariance. Strict linear orders are maximally rigid, and any nontrivial stabilizer on the input side produces an impossibility. Partitions are fully flexible under the fixed-set relation, but exhibit sharp rigidity (in the sense $|\mathrm{Fix}| = 1$) only under a background weak order. Weak orders sit between partitions and probability distributions. They admit $a_k$ fixed elements ($a_k \geq B_k$, comparing Fubini numbers to orbit-coarsenings rather than to all $H$-fixed partitions), and exhibit sharp rigidity precisely under transitivity of the stabilizer. Probability distributions on a set carrying a transitive group action have a unique symmetric member, namely the uniform distribution, and every impossibility-characterization theorem for such distributions turns on this fact. Real-valued targets under a trivial action are the most flexible, and characterizations there require the additive refinement.

Several directions extend the framework naturally. On the analytic side, the topological refinements of Chichilnisky (1980), Baryshnikov (1993), and Weinberger (2004) extend the mechanism to positive-dimensional input spaces, and the quantitative refinements of Feys (2026i) and Feys (2026h) extend it to approximate equivariance. On the algebraic side, a systematic classification of the coinvariant Grothendieck groups $(X^{\mathrm{gp}})_G$ for standard cooperative-game, pricing, and market-design structures would turn the additive refinement into a lookup table. On the combinatorial side, an extension of the rigidity taxonomy to lattice-valued and metric-space-valued targets, and an analogue of the principle for partial group actions on restricted preference domains, both appear tractable. Taken together, these directions suggest that the equivariance mechanism, though elementary, admits a genuine research program.

A methodological word is in order. When the substantive content of a result is elementary, heavy categorical machinery tends to hide the mechanism rather than reveal it. The additive refinement of Section 4 is a case in point. It falls out from two composed universal properties (Lemmas 4.1 and 4.2), and the correct invariant of the problem is the coinvariant Grothendieck group $(X^{\mathrm{gp}})_G$. The identification of this invariant is the mathematical substance of the refinement, and it is best presented at the elementary level at which it lives. The present paper is written accordingly.

## Declarations

**Conflict of interest.** The author declares no conflicts of interest.

**Funding.** No funding was received for this work.

**Data availability.** No empirical data were used.

## References

Kenneth J. Arrow. *Social Choice and Individual Values*. Cowles Foundation Monograph, second edition. Yale University Press, 1963.

Robert J. Aumann and Michael Maschler. Game theoretic analysis of a bankruptcy problem from the Talmud. *Journal of Economic Theory*, 36(2):195–213, 1985.

Yuliy M. Baryshnikov. Unifying impossibility theorems: A topological approach. *Advances in Applied Mathematics*, 14(4):404–415, 1993.

Anna Bogomolnaia and Hervé Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.

Daniela Bubboloni and Michele Gori. Anonymous and neutral majority rules. *Social Choice and Welfare*, 43(2):377–401, 2014.

Graciela Chichilnisky. Social choice and the topology of spaces of preferences. *Advances in Mathematics*, 37(2):165–176, 1980.

Beno Eckmann. Räume mit Mittelbildungen. *Commentarii Mathematici Helvetici*, 28(1):329–340, 1954.

Frank M. V. Feys. An Axiomatic Theory of Tie-Breaking: Impossibility, Characterization, and Decomposition. Working paper, 2026.

Frank M. V. Feys. The Uniform Completion Theorem: Canonical Randomization in Abstract Tie-Breaking. Working paper, 2026.

Frank M. V. Feys. An Axiomatic Theory of Pricing Kernel Selection. Working paper, 2026.

Frank M. V. Feys. Fairness and Strategy-Proofness in Automated Market Makers. Working paper, 2026.

Frank M. V. Feys. A Fixpoint Theory of Label-Symmetric Classification. Working paper, 2026.

Frank M. V. Feys. Nudging without Reasons: A Trilemma and the Cost of Verification. Working paper, 2026.

Frank M. V. Feys. The ANR Impossibility for Rankings as a Corollary of Abstract Tie-Breaking. Working paper, 2026.

Frank M. V. Feys. Belief Aggregation for Pricing Kernels: Arrovian and Strategic Impossibilities. Working paper, 2026.

Frank M. V. Feys. Boolean Functions and Systemic Risk: A Fourier-Analytic Framework. Working paper, 2026.

Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.

Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political Economy*, 87(2):293–314, 1979.

Kenneth O. May. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica*, 20(4):680–684, 1952.

Hervé Moulin. *The Strategy of Social Choice*. Advanced Textbooks in Economics, volume 18. North-Holland, 1983.

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

Barry O'Neill. A problem of rights arbitration from the Talmud. *Mathematical Social Sciences*, 2(4):345–371, 1982.

Mark A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

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, pages 307–317. Princeton University Press, 1953.

William Thomson. Axiomatic and game-theoretic analysis of bankruptcy and taxation problems: A survey. *Mathematical Social Sciences*, 45(3):249–297, 2003.

Charles A. Weibel. *An Introduction to Homological Algebra*. Cambridge Studies in Advanced Mathematics, volume 38. Cambridge University Press, 1994.

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