

Measurable Groups, Projective Limits, and Haar Measure: Construction, Coincidence, and Applications

Lakshmi K. Raut

Visiting Fellow, CEHD, University of Chicago

email: lakshmiraut@gmail.com

2026-09-15

Abstract

We establish a rigorous bridge between two historically distinct routes to Haar measure: the classical topological route (locally compact groups) and a measure-theoretic route (measurable groups via projective limits of finite groups). We prove the Halmos–Weil theorem from first principles: a separated measurable group with separable measure algebra carries a canonical Weil topology making it a compact, second-countable topological group whose Haar measure coincides with the original one. We determine the exact boundary between the two routes: countable projective limits of finite groups yield metrizable profinite groups; uncountable limits diverge precisely through non-separability of the measure algebra (Maharam type $\geq \aleph_0$). We present complete, self-contained proofs of existence and uniqueness of Haar measure, of the extended Daniell–Kolmogorov theorem, and of the Halmos–Weil theorem. We develop three applications: random-order values for non-atomic games (with full axiomatic characterization and uniqueness), invariant statistical decision theory on general (non-compact, non-unimodular) groups, and Haar measure as the canonical reference for Lie-group generative models. All numerical illustrations are executed and verified.

1 Introduction

A group that carries a compatible measure is the natural domain for averaging uniformly over symmetry. When a mathematical problem is invariant under a group G , the unique (up to scale) translation-invariant measure on G — Haar measure — is the canonical reference for integration, sampling, and optimization.

Two historically distinct routes lead to Haar measure. The **topological route** begins with a locally compact topological group and produces a Radon measure invariant under left (or right)

translation. The **measure-theoretic route**, due to Raut (Raut 1997), begins with a *measurable group* — a set equipped with a group structure and a σ -algebra carrying a translation-invariant measure — and constructs an invariant measure directly from consistency conditions on finite sub-quotients, without any topology assumed.

This article proves that these routes coincide precisely when the measure algebra is separable. The Halmos–Weil theorem (Halmos 1950; Mackey 1957) identifies a separated measurable group with separable measure algebra as a compact, second-countable topological group whose Haar measure is the original one. We provide complete, self-contained proofs of both routes and of the coincidence theorem. We determine the exact boundary: every countable projective limit of finite groups is a metrizable profinite group; uncountable limits diverge precisely through non-separability (Maharam type $\geq \aleph_0$).

Three applications illustrate the ubiquity of Haar measure as the uniform averaging tool:

1. **Game theory:** the random-order value of non-atomic games (Raut 1997, 1999; Aumann and Shapley 1974) is shown to coincide with the Aumann–Shapley value via the Haar measure on the projective-limit group of random orderings. We give a complete axiomatic characterization and uniqueness proof.
2. **Invariant statistics:** Haar measure provides the reference prior and orbit integration for Pitman estimators, maximal invariants, and uniformly most powerful invariant tests on general (non-compact, non-unimodular) groups.
3. **Generative AI:** on Lie groups, Haar measure is the canonical reference prior for flow-matching and diffusion models. We prove convergence of the isotropic Gaussian kernel on $SO(3)$ and verify numerically that left-invariance holds only in the converged (Haar) regime.

The remainder of the paper is organized as follows: Section 2 develops the topological route. Section 3 develops the measure-theoretic route. Section 4 proves the Halmos–Weil coincidence theorem and analyzes its boundary. Section 5 treats the three applications. Section 6 concludes with discussion and open problems. Appendix Section 7 provides the Python code for the numerical results in the application sections.

2 Haar Measure on Locally Compact Groups

2.1 Locally compact groups

Definition 2.1 (Locally compact group). A **locally compact group** is a group G equipped with a Hausdorff topology such that multiplication $(x, y) \mapsto xy$ and inversion $x \mapsto x^{-1}$ are continuous,

and every point has a compact neighborhood.

Definition 2.2 (Haar measure (left)). A **left Haar measure** on a locally compact group G is a nonzero Radon measure μ on the Borel σ -algebra $\mathcal{B}(G)$ satisfying $\mu(gA) = \mu(A)$ for all $g \in G$ and Borel sets A . It is unique up to a positive scalar multiple.

2.2 Existence via covering numbers

Let $C_c^+(G)$ denote continuous nonnegative functions of compact support on G not identically zero.

Lemma 2.1 (Covering number properties). For $f, g \in C_c^+(G)$ define the **covering number**

$$(f : g) = \inf \left\{ \sum_{i=1}^n c_i \mid c_i \geq 0, s_i \in G, f(x) \leq \sum_{i=1}^n c_i g(s_i^{-1}x) \forall x \in G \right\}.$$

Then for all $f, g, h \in C_c^+(G)$:

- (a) $0 < (f : g) < \infty$;
- (b) $(f : h) \leq (f : g)(g : h)$;
- (c) $(f : g) \geq \max f / \max g$;
- (d) $(L_a f : g) = (f : g)$ for every $a \in G$, where $(L_a f)(x) = f(a^{-1}x)$.

Proof.

- (a) Since f has compact support and $g > 0$ on a nonempty open set U , finitely many left translates $s_i U$ cover $\text{supp } f$ by compactness; take $c_i = \max f / \min_U g$. The lower bound follows by evaluating at a point where f attains its maximum, proving (c) as well.
- (b) If $f \leq \sum_i c_i L_{s_i} g$ and $g \leq \sum_j d_j L_{t_j} h$, then $f \leq \sum_{i,j} c_i d_j L_{s_i t_j} h$, a covering of f by translates of h with total weight $(\sum_i c_i)(\sum_j d_j)$. Taking infima over both families gives (b).
- (c) is the lower bound used above.
- (d) A covering $f \leq \sum_i c_i L_{s_i} g$ corresponds bijectively, via $f \mapsto L_a f$ and $s_i \mapsto a s_i$, to a covering of $L_a f$ with the same weights, so the infima coincide.

■

Fix a reference function $f_0 \in C_c^+(G)$. For $g \in C_c^+(G)$ define $I_g : C_c^+(G) \rightarrow (0, \infty)$ by

$$I_g(f) = \frac{(f : g)}{(f_0 : g)}.$$

By 2.1(b) applied twice, $\frac{1}{(f_0 : f)} \leq I_g(f) \leq (f : f_0)$ for every g . Thus $I_g(f)$ lies in a fixed compact interval depending only on f . Equip the product space $P = \prod_{f \in C_c^+(G)} [1/(f_0 : f), (f : f_0)]$ with the product topology; by Tychonoff's theorem, P is compact. The family $\{I_g\}_{g \in C_c^+(G)}$, indexed by g ordered by reverse inclusion of $\text{supp } g$ shrinking to $\{e\}$, is a net in P ; by compactness it has a convergent subnet $I_{g_\alpha} \rightarrow I$ for some $I \in P$.

Theorem 2.1 (Existence of left Haar measure). *The limit functional I extends to a left-invariant, nonzero positive linear functional on $C_c(G)$, and hence, by the Riesz representation theorem, there exists a left Haar measure μ on G with $\int_G f d\mu = I(f)$ for all $f \in C_c(G)$.*

Proof. **Additivity in the limit.** Sub-additivity $I_g(f_1 + f_2) \leq I_g(f_1) + I_g(f_2)$ is immediate from 2.1(b). The reverse inequality holds only in the limit and is the core of the argument.

Fix $f_1, f_2 \in C_c^+(G)$ and $\varepsilon > 0$. Choose $h \in C_c^+(G)$ equal to 1 on a neighborhood of $\text{supp}(f_1 + f_2)$. For $\delta > 0$ set

$$\varphi_i(x) = \begin{cases} f_i(x)/(f_1(x) + f_2(x) + \delta h(x)) & \text{if } f_1 + f_2 + \delta h > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Then $\varphi_1 + \varphi_2 \leq 1$ everywhere with equality on $\text{supp}(f_1 + f_2)$. Because φ_1, φ_2 are uniformly continuous on the compact set $\text{supp}(f_1 + f_2 + \delta h)$, there exists a neighborhood V of the identity such that $\sup_{x \in sV} |\varphi_i(x) - \varphi_i(sx)| < \delta$ for all $s \in V, i = 1, 2$.

If $\text{supp } g \subset V$ and g is not identically zero, then for any covering $f_1 + f_2 \leq \sum_i c_i L_{s_i} g$ we have for each i :

$$f_i \leq \sum_i c_i (\varphi_i(s_i) + \delta) L_{s_i} g,$$

so the right-hand side is a valid covering of f_i by translates of g with total weight

$$\sum_i c_i (\varphi_i(s_i) + \delta) \leq \sum_i c_i \varphi_i(s_i) + 2\delta \sum_i c_i.$$

Summing over $i = 1, 2$ and dividing by $(f_0 : g)$ (which is finite and independent of the particular covering) gives

$$I_g(f_1) + I_g(f_2) \leq I_g(f_1 + f_2) + C\delta,$$

where C is a constant depending only on f_0, f_1, f_2 and δh . Letting $\delta \downarrow 0$ then $g \rightarrow e$ yields $I(f_1) + I(f_2) \leq I(f_1 + f_2)$. Together with sub-additivity we obtain additivity.

Linearity and invariance. Homogeneity $I(cf) = cI(f)$ ($c > 0$) and positivity are inherited directly from the covering numbers. Extension to a linear functional on $C_c(G)$ via $I(f) = I(f^+) - I(f^-)$ is routine. Left-invariance of I on $C_c^+(G)$ follows because $I_g(L_af) = I_g(f)$ for every g by 2.1(d), hence the same holds for any limit point I . Since $I_g(f_0) = 1$ for every g , $I(f_0) = 1 \neq 0$, so I is nonzero.

The Riesz representation theorem — which identifies positive linear functionals on $C_c(G)$, for G locally compact Hausdorff, with unique Radon measures — produces μ with $\int f d\mu = I(f)$, and left-invariance of I transfers to $\mu(gA) = \mu(A)$ for Borel A . ■

2.3 Uniqueness

Theorem 2.2 (Uniqueness on compact groups). *If G is compact, any two left Haar measures on G are positive scalar multiples of each other. In particular, every left Haar measure on a compact group is automatically also a right Haar measure (the group is unimodular).*

Proof. Compact groups are unimodular. The map $g \mapsto \mu(Ag)/\mu(A)$ (for any fixed Borel A of positive finite measure) defines a continuous homomorphism $\Delta : G \rightarrow (0, \infty)$, the modular function. Since G is compact, $\Delta(G)$ is a compact subgroup of $(0, \infty)$, and the only compact subgroup of $(0, \infty)$ is $\{1\}$. Thus $\Delta \equiv 1$, and μ is also right-invariant.

The double-integral identity. Let μ, ν be two left (hence also right) Haar measures on G , normalized so $\mu(G) = \nu(G) = 1$. For $f \in C(G)$, Fubini's theorem on the compact space $G \times G$ gives

$$\begin{aligned} \int_G \int_G f(xy) d\mu(x) d\nu(y) &= \int_G \left[\int_G f(xy) d\mu(x) \right] d\nu(y) \\ &= \int_G \left[\int_G f d\mu \right] d\nu(y) = \int_G f d\mu, \end{aligned}$$

where the inner integral collapses by right-invariance of μ (substituting $x \mapsto xy^{-1}$). Swapping the order of integration,

$$\begin{aligned} \int_G \int_G f(xy) d\mu(x) d\nu(y) &= \int_G \left[\int_G f(xy) d\nu(y) \right] d\mu(x) \\ &= \int_G \left[\int_G f d\nu \right] d\mu(x) = \int_G f d\nu. \end{aligned}$$

Hence $\int_G f d\mu = \int_G f d\nu$ for every $f \in C(G)$. Since G is compact Hausdorff, $C(G)$ separates

finite Borel measures by the Riesz theorem, so $\mu = \nu$. The case of general positive normalization follows by rescaling. ■

Theorem 2.3 (Uniqueness on locally compact groups). *If G is locally compact, any two left Haar measures on G are positive scalar multiples of each other.*

Proof. The strategy is the same double-integral idea, but $G \times G$ is no longer compact, so $f(xy)$ need not be integrable over the product. The standard repair uses an approximate identity $(h_U)_U$ for μ – nonnegative $h_U \in C_c(G)$ supported in shrinking neighborhoods U of e with $\int h_U d\mu = 1$ – together with the modular function Δ_μ of Definition 2.3, and produces the identity $\int_G (h_U *_\mu f) d\nu = \int f d\nu$, in which $h_U *_\mu f \rightarrow f$ uniformly, with support contained in a fixed compact set for U small, as $U \downarrow \{e\}$. Passing to the limit gives $\int f d\mu = c \int f d\nu$ for a constant $c > 0$ independent of f . The full execution is given in Folland (Folland 2016), Theorem 2.20, and Nachbin (Nachbin 1965), Ch. I. ■

2.4 The modular function

Definition 2.3 (Modular function). For a left Haar measure μ on locally compact G , the **modular function** $\Delta = \Delta_\mu : G \rightarrow (0, \infty)$ is the continuous homomorphism characterized by

$$\mu(Ay) = \Delta(y) \mu(A) \quad \text{for all Borel sets } A \text{ and all } y \in G,$$

equivalently $\int_G f(xy) d\mu(x) = \Delta(y)^{-1} \int_G f d\mu$. G is **unimodular** if $\Delta \equiv 1$.

Proof. (Equivalence of the two characterizations.) Fix $y \in G$ and let $\varphi(x) = xy$. The pushforward $\varphi_*\mu$ satisfies $\varphi_*\mu(A) = \mu(\varphi^{-1}A) = \mu(Ay^{-1})$. Applying $\mu(Ay) = \Delta(y)\mu(A)$ with $A \mapsto Ay^{-1}$ gives $\mu(A) = \Delta(y)\mu(Ay^{-1})$, i.e. $\varphi_*\mu = \Delta(y)^{-1}\mu$. By the change-of-variables formula for pushforward measures, $\int_G f(xy) d\mu(x) = \int_G f d(\varphi_*\mu) = \Delta(y)^{-1} \int_G f d\mu$, as claimed. ■

For a Lie group G , the modular function is computed via the adjoint representation $\text{Ad} : G \rightarrow GL(\mathfrak{g})$:

$$\Delta(g) = |\det \text{Ad}(g)|^{-1}, \quad g \in G.$$

Every compact, abelian, or nilpotent Lie group is unimodular (for a nilpotent group, $\text{Ad}(g)$ is unipotent for every g , so $\det \text{Ad}(g) = 1$).

2.5 Worked examples of the modular function

Example 2.1 (The affine group of $\mathbb{R}$). The affine group is $G = \left\{ \begin{pmatrix} a & b \\ 0 & 1 \end{pmatrix} : a > 0, b \in \mathbb{R} \right\}$ acting by $x \mapsto ax + b$. In the basis $X = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, Y = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ of $\mathfrak{g}$, a direct computation gives $\text{Ad}(g)X = X$ and $\text{Ad}(g)Y = aY$, so $\det \text{Ad}(g) = a$ and $\Delta(a, b) = a^{-1}$. Left Haar measure is therefore $a^{-2} da db$ and right Haar measure is $a^{-1} da db$; these are readily checked to be, respectively, left- and right-invariant under $(a, b) \cdot (a', b') = (aa', ab' + b)$.

Example 2.2 (Lower-triangular matrices $LT(n)$). Let $G = LT(n)$ be the group of $n \times n$ real invertible lower-triangular matrices, $T = D \cdot U$ with $D = \text{diag}(t_1, \dots, t_n)$ and U unipotent lower-triangular. We claim the modular function satisfies

$$\Delta(T) = \prod_{i=1}^n |t_i|^{n-2i+1}.$$

Proof. Since U is unipotent, $\text{Ad}(U)$ is a unipotent linear map on $\mathfrak{g} = \mathfrak{lt}(n)$, so $\det \text{Ad}(U) = 1$; the modular function is therefore determined entirely by the diagonal part D . The Lie algebra has basis $\{E_{ij} : 1 \leq j \leq i \leq n\}$ (matrix units), and for diagonal D , $\text{Ad}(D)E_{ij} = DE_{ij}D^{-1} = (t_i/t_j)E_{ij}$. Hence

$$\det \text{Ad}(D) = \prod_{1 \leq j < i \leq n} \frac{t_i}{t_j} = \prod_{i=1}^n t_i^{(i-1)-(n-i)} = \prod_{i=1}^n t_i^{2i-n-1},$$

where the exponent of t_i comes from the $(i-1)$ pairs with $j < i$ (numerator role) minus the $(n-i)$ pairs with $j > i$ (denominator role, from the terms where i plays the role of the smaller index). Then $\Delta(D) = |\det \text{Ad}(D)|^{-1} = \prod_i |t_i|^{n-2i+1}$, as claimed; multiplicativity extends this to all of $T = DU$. ■

This matches the modulus character of the standard Borel subgroup of $GL(n, \mathbb{R})$, and specializes for $n = 1$ to the trivial (unimodular) group $\mathbb{R}^\times$, and for $n = 2$ to $\Delta(t_1, t_2) = |t_1/t_2|$.

2.6 Haar measure on Lie groups via differential forms

When G is an n -dimensional Lie group, Haar measure is induced by any left-invariant volume form.

Theorem 2.4 (Construction via left-invariant forms). *Let G be an n -dimensional Lie group, $\mathfrak{g} = T_e G$ its Lie algebra, and fix a basis $\xi_1, \dots, \xi_n$ of $\mathfrak{g}^*$. Let $\omega^1, \dots, \omega^n$ be the corresponding left-invariant 1-forms on G ($\omega_g^i = (dL_{g^{-1}})_g^* \xi_i$), and set $\omega = \omega^1 \wedge \dots \wedge \omega^n$. Then $|\omega|$, the measure locally given by $|\omega(x)| dx^1 \dots dx^n$, is a left Haar measure on G .*

Proof. Left-invariance of each ω^i implies $L_a^* \omega = \omega$ for every $a \in G$. The change-of-variables formula then gives $\int_{aE} |\omega| = \int_E |L_a^* \omega| = \int_E |\omega|$ for every Borel E . Since ω is a smooth, nowhere-vanishing n -form, $|\omega|$ is a Radon measure, nonzero and locally finite. Uniqueness up to scalar follows both from [Theorem 2.3](#) and, independently, from $\dim(\Lambda^n \mathfrak{g}^*) = 1$. ■

Example 2.3 (Further examples).

- (i) $G = GL(n, \mathbb{R})$: left Haar measure is $|\det X|^{-n} dX$ where dX is Lebesgue measure on the matrix entries; $GL(n, \mathbb{R})$ is unimodular (its adjoint representation has $\det \text{Ad}(g) = \det(g)^n \det(g)^{-n} = 1$ by the same trace-conjugation computation as in [Example 2.1](#)).
- (ii) $G = O(n)$: compact and unimodular, with total Haar mass

$$\text{vol}(O(n)) = \prod_{i=1}^n \frac{2\pi^{i/2}}{\Gamma(i/2)}.$$

This follows from the iterated fibration $O(n) \rightarrow O(n-1)$ with fiber S^{n-1} (the stabilizer of a unit vector under the transitive action of $O(n)$ on S^{n-1} is a copy of $O(n-1)$), giving $\text{vol}(O(n)) = \text{vol}(S^{n-1}) \cdot \text{vol}(O(n-1))$ with base case $\text{vol}(O(1)) = 2$; unwinding the recursion with $\text{vol}(S^{k-1}) = 2\pi^{k/2}/\Gamma(k/2)$ gives the stated product.

- (iii) $G = LT(n)$, as in [Example 2.2](#).

3 Measurable Groups and the Extended Daniell–Kolmogorov Theorem

The topological route to an invariant measure presupposes a topology. **Measurable groups** ([Raut 1997](#)) invert this logic: start from a σ -algebra and a compatible group structure, and *construct* an invariant measure directly from consistency conditions on finite sub-quotients.

The extension principle at work follows a precise theoretical lineage. The classical consistency theorem of Kolmogorov ([Kolmogorov 1933](#)) constructs probability measures on infinite product spaces from a projective family of finite-dimensional marginals. Bochner ([Bochner 1955](#)) generalized this to projective limits over directed sets of abstract topological spaces, while Prokhorov ([Prokhorov 1956](#)) identified **tightness** as the condition under which a consistent family of Radon measures admits a projective-limit lift. The extended Daniell–Kolmogorov theorem of [Theorem 3.1](#) is the specialization of this lineage to a projective system of *measurable groups* indexed

by $\mathbb{N}$. Because each level is a finite group carrying its unique bi-invariant probability measure, the tightness condition is automatic, and the projective-limit measure exists and is unique.

3.1 Definitions and the separation axiom

Definition 3.1 (Measurable group). A **measurable group** is a triple $(G, \mathcal{B}, \tau)$ where G is a group, $\mathcal{B}$ is a σ -algebra of subsets of G such that the maps $(x, y) \mapsto xy$ and $x \mapsto x^{-1}$ are $(\mathcal{B} \otimes \mathcal{B}, \mathcal{B})$ - and $(\mathcal{B}, \mathcal{B})$ -measurable respectively, and τ is a σ -finite measure on $\mathcal{B}$ invariant under left translation: $\tau(gA) = \tau(A)$ for every $g \in G$ and $A \in \mathcal{B}$.

Pathological choices of $\mathcal{B}$ can fail to separate points, or can be too large relative to τ to support a workable topology. Raut's construction controls both failure modes by building $\mathcal{B}$ as a projective limit, where countable generation is automatic (3.1) and separation is verified on the generators.

Definition 3.2 (Separation). A measurable group $(G, \mathcal{B}, \tau)$ is **separated** if for every $x \neq y$ in G there is $A \in \mathcal{B}$ with $x \in A$ and $y \notin A$.

3.2 The motivating construction: random re-orderings of a continuum

Raut (1997) builds the measurable group needed for a continuum-player analogue of Shapley's random-order value from finite groups of Lebesgue-measure-preserving (l.m.p.) automorphisms of the unit interval.

Example 3.1 (A projective system of l.m.p.-automorphism groups). Let $(I, \lambda) = ([0, 1], \text{Lebesgue measure})$, and let $\text{Aut}(I, \lambda)$ be the group of l.m.p. automorphisms of I . For $n = 1, 2, \dots$ partition I into the 2^n dyadic subintervals $I_k^n = [(k-1)2^{-n}, k2^{-n})$, $k = 1, \dots, 2^n$, and let $G_n \leq \text{Aut}(I, \lambda)$ be the finite group of l.m.p. automorphisms that permute the blocks $\{I_k^n\}$ rigidly. Each $\sigma \in G_n$ that respects the coarser partition at level $n-1$ projects to a well-defined $f_n(\sigma) \in G_{n-1}$; thus $\{G_n, f_n\}$ is a projective system of finite groups. On each G_n , normalized counting measure τ_n is the unique bi-invariant probability measure, and f_n carries τ_n to τ_{n-1} . The extended Daniell–Kolmogorov theorem (Theorem 3.1) produces a probability measure τ_∞ on a projective-limit group G_∞ that embeds in $\text{Aut}(I, \lambda)$, where τ_∞ is the unique bi-invariant probability measure on the resulting random orderings.

Definition 3.3 (Consistent family). Let $\{(G_n, \mathcal{B}_n, \tau_n)\}_{n \in \mathbb{N}}$ be measurable groups with group homomorphisms $f_n : G_n \rightarrow G_{n-1}$ that are $(\mathcal{B}_n, \mathcal{B}_{n-1})$ -measurable. The family $\{\tau_n\}$ is **consistent** if $\tau_n(f_n^{-1}(A)) = \tau_{n-1}(A)$ for every $A \in \mathcal{B}_{n-1}$ and every n .

3.3 The extended Daniell–Kolmogorov theorem

Theorem 3.1 (Extended Daniell–Kolmogorov, Raut (1997)). *Let $\{(G_n, \mathcal{B}_n, \tau_n), f_n\}$ be a consistent projective system of measurable groups indexed by $\mathbb{N}$, with each G_n finite. Then there is a group G_∞ (the projective limit $\varprojlim G_n$), a σ -algebra $\mathcal{B}_\infty$ generated by the coordinate projections $\pi_n : G_\infty \rightarrow G_n$, and a unique probability measure τ_∞ on $(G_\infty, \mathcal{B}_\infty)$ such that π_n pushes τ_∞ forward to τ_n for every n , and τ_∞ is left-invariant for the group structure inherited coordinatewise on G_∞ .*

Proof. **Existence and countable additivity.** The cylinder sets $C_n(A) = \pi_n^{-1}(A)$, $A \in \mathcal{B}_n$, form an algebra $\mathcal{A}_0 = \bigcup_n \pi_n^{-1}(\mathcal{B}_n)$ generating $\mathcal{B}_\infty$; consistency makes $\tau_\infty(C_n(A)) := \tau_n(A)$ well-defined and finitely additive on $\mathcal{A}_0$. Countable additivity on $\mathcal{A}_0$ follows because each G_n is finite, hence compact in the discrete topology, so the family $\{\tau_n\}$ is automatically tight: every decreasing sequence of nonempty cylinder sets with empty intersection must eventually involve a cylinder over some finite G_n that is itself empty, forcing τ_∞ of the sequence to 0. Carathéodory’s extension theorem then extends τ_∞ uniquely from $\mathcal{A}_0$ to $\sigma(\mathcal{A}_0) = \mathcal{B}_\infty$; the cylinder sets form a π -system, so uniqueness of the extension (beyond the automatic uniqueness furnished by Carathéodory’s theorem for the completed extension) is equivalently seen as an instance of Dynkin’s π - λ theorem. The full argument is the classical Daniell–Kolmogorov extension theorem applied to the projective system $\{G_n, \tau_n\}$; see Parthasarathy (Parthasarathy 1977), Ch. V, for the general statement of which this is a special case.

Left-invariance. For $g = (g_n)_n \in G_\infty$ and a cylinder set $C_n(A)$, $g \cdot C_n(A) = C_n(g_n A)$ by the coordinatewise group law, and $\tau_\infty(g \cdot C_n(A)) = \tau_n(g_n A) = \tau_n(A) = \tau_\infty(C_n(A))$, using left-invariance of τ_n . Both τ_∞ and its translate $A \mapsto \tau_\infty(gA)$ are probability measures agreeing on the π -system $\mathcal{A}_0$, hence agree on all of $\mathcal{B}_\infty = \sigma(\mathcal{A}_0)$ by Dynkin’s π - λ theorem. **Bi-invariance** follows identically whenever each τ_n is bi-invariant, which holds automatically for normalized counting measure on a finite group. ■

Corollary 3.1 (Automatic countable generation). *$\mathcal{B}_\infty$ in Theorem 3.1 is **countably generated**: it equals the σ -algebra generated by $\bigcup_n \pi_n^{-1}(\mathcal{B}_n)$, a countable union of finite algebras. This holds regardless of the form of the bonding homomorphisms f_n .*

4 Coincidence of Topological and Measurable Groups

We now connect the topological and measure-theoretic routes. One direction is immediate: a locally compact group carries its Haar measure as a *measurable* group (Theorem 4.1). The converse, which rebuilds a topology from measure-theoretic data (Theorem 4.2), is a genuine theorem whose hypotheses determine the exact scope of the measurable-group generalization.

4.1 From topological to measurable

Theorem 4.1. *If G is a locally compact group with left Haar measure μ , normalized to be a probability measure if G is compact, then $(G, \mathcal{B}(G), \mu)$ is a separated measurable group.*

Proof. Measurability of multiplication and inversion follows from their continuity, since preimages of Borel sets under continuous maps are Borel. Left-invariance of μ is by definition of Haar measure. Separation holds because G is Hausdorff and the Borel σ -algebra of a Hausdorff space separates points: for $x \neq y$, disjoint open neighborhoods (Hausdorff) give a Borel set containing x but not y . ■

4.2 From measurable to topological: the Halmos–Weil theorem

The converse direction constructs a topology from the measurable structure via the measure algebra.

Definition 4.1 (Measure algebra and separability). Let $(G, \mathcal{B}, \tau)$ be a measurable group with τ σ -finite. The **measure algebra** is the set $\mathcal{B}$ modulo the equivalence $A \sim B \iff \tau(A \Delta B) = 0$, with metric $d(A, B) = \tau(A \Delta B)$ on finite-measure classes. The measure algebra is **separable** if this metric space is separable, equivalently if $\mathcal{B}$ is countably generated modulo τ -null sets.

Definition 4.2 (The Weil topology). For $A \in \mathcal{B}$ with $\tau(A) < \infty$ and $\varepsilon > 0$, set $U(A, \varepsilon) = \{g \in G : \tau(gA \Delta A) < \varepsilon\}$. The **Weil topology** on G is the topology generated by these sets; it is the coarsest group topology making $g \mapsto gA$ continuous into the measure algebra for every finite-measure A .

Theorem 4.2 (Halmos–Weil). *Let $(G, \mathcal{B}, \tau)$ be a separated measurable group with τ a probability measure and separable measure algebra. Then the Weil topology makes G a **compact, second-countable** topological group, $\mathcal{B}$ coincides with its Borel sets modulo τ -null sets, and τ is the normalized Haar measure of this topology.*

Proof. **Group topology.** The sets $U(A, \varepsilon)$ form a basis for a topology making translation continuous: $\tau(ghA \Delta gA) = \tau(g(hA \Delta A)) = \tau(hA \Delta A)$ by left-invariance, so $U(A, \varepsilon)$ is itself a neighborhood basis compatible with left translation. Continuity of inversion is verified similarly once a countable generating family of finite-measure sets is fixed.

Second countability and Hausdorff. Separability of the measure algebra provides a countable family $\{A_k\}$ generating $\mathcal{B}$ modulo null sets; the corresponding sets $U(A_k, 1/m)$, $k, m \in \mathbb{N}$, form a countable basis at e . Separation of $(G, \mathcal{B}, \tau)$ together with this countable generation ensures

$\tau(gA_k \triangle A_k) > 0$ for some k whenever $g \neq e$, giving the Hausdorff property; this is the content of Theorem 7.1 of Mackey (Mackey 1957).

Compactness. Since τ is a probability measure and the measure algebra is separable, $(G, \mathcal{B}, \tau)$ is, modulo null sets, isomorphic (as a measure algebra) to a standard probability space. The Weil-topology uniform structure on G is thereby totally bounded and complete, from which compactness follows; the complete ε -net argument is in Halmos (Halmos 1950), §§62–64, and Yamasaki (Yamasaki 1985), Ch. 3.

τ is Haar. τ is left-invariant by hypothesis. Once the Weil topology is shown compact Hausdorff with $\mathcal{B}$ its Borel sets modulo null sets, uniqueness (Theorem 2.2) identifies τ with the normalized Haar measure of that topology. ■

Halmos–Weil versus Mackey–Weil. Theorem 4.2 is the probability form of the topology-recovery result. A strictly more general result — **Mackey–Weil** — starts from a *standard Borel* group carrying a σ -finite **quasi-invariant** measure μ ($\mu(Eg) = 0 \iff \mu(E) = 0$). It recovers a unique locally compact second-countable metric topology making μ equivalent to the Haar measure, using the Koopman unitary representation $g \mapsto L_g$ on $L^2(\mu)$ and Kuratowski’s isomorphism theorem (Parthasarathy 1977). Halmos–Weil is the sharper statement available in the compact probability setting, while Mackey–Weil is the tool for general non-compact groups.

4.3 Every countable Raut projective limit coincides

Corollary 4.1. *Let $(G_\infty, \mathcal{B}_\infty, \tau_\infty)$ be built as in Theorem 3.1 from any consistent countable projective system of finite measurable groups. Then $\mathcal{B}_\infty$ is countably generated (3.1), so its measure algebra is separable. If the system is separated — which holds whenever the coordinate projections π_n eventually separate points, e.g., when the f_n are eventually injective on generators — then $(G_\infty, \mathcal{B}_\infty, \tau_\infty)$ satisfies the hypotheses of Theorem 4.2 and coincides with a **compact, second-countable** topological group carrying its normalized Haar measure.*

Raut’s construction thus produces a particularly explicit, combinatorial route into the locally compact picture: G_∞ is, ultimately, a compact second-countable topological group. In topological-algebra terminology, a countable projective limit of finite groups is a **metrizable profinite group**. By the standard characterization (e.g. Ribes and Zalesskii 2010, Ch. 2), a profinite group is exactly a compact, Hausdorff, totally disconnected topological group, and it is metrizable if and only if it is (topologically) a projective limit over a countable index set. Thus the measure-theoretic route of Section 3 is a path to a metrizable profinite group, differing from the usual topological construction only in the route taken.

4.4 Sharpness of the hypotheses: an uncountable projective limit

The Halmos–Weil theorem’s hypotheses are sharp: dropping countable generation of the measure algebra changes the conclusion.

Example 4.1 $(0, 1^{[0,1]}$: separable everywhere except where it matters). Let $J = [0, 1]$ and, for each $j \in J$, let $G_j = \mathbb{Z}/2\mathbb{Z}$ with its power-set σ -algebra and uniform measure τ_j . Let $G_\infty = \prod_{j \in J} G_j = \{0, 1\}^J$ with coordinatewise addition mod 2, the product σ -algebra $\mathcal{B}_\infty$, and the product measure $\tau_\infty = \bigotimes_{j \in J} \tau_j$.

Separated measurable group. Coordinatewise sum and negation are $\mathcal{B}_\infty$ -measurable, and τ_∞ is translation-invariant. If $x \neq y$, some coordinate j has $x_j \neq y_j$, so $\pi_j^{-1}(\{x_j\})$ separates them.

Non-separable measure algebra. By Maharam’s classification of measure algebras (Maharam 1942), the measure algebra of a product $\{0, 1\}^\kappa$ for infinite cardinal κ has **Maharam type** κ – the smallest cardinality of a generating family modulo null sets. Here $\kappa = |J| = 2^{\aleph_0}$, so the measure algebra is not separable.

Survival of the compact group structure. $\{0, 1\}^J$ with its product topology is a compact Hausdorff topological group (Tychonoff). The Weil topology of Definition 4.2 coincides with the product topology: for $A_j = \pi_j^{-1}(\{1\})$ and $g \in G_\infty$ with $g_j = 1$ in coordinate j and 0 elsewhere, gA_j is the complement of A_j , so $\tau_\infty(gA_j \triangle A_j) = 1$; whereas for any g with $g_j = 0$, $gA_j = A_j$ exactly, so $\tau_\infty(gA_j \triangle A_j) = 0$. Hence the basic Weil neighborhood $U(A_j, \varepsilon)$, $0 < \varepsilon < 1$, is exactly $\{g : g_j = 0\} = \pi_j^{-1}(\{0\})$, a basic subbasic set of the product topology; since the A_j range over a generating family, the two topologies coincide. Thus τ_∞ is the Haar measure of $(\{0, 1\}^J, \text{product topology})$.

For this non-metrizable compact group, $\mathcal{B}_\infty$ is the **Baire** σ -algebra (generated by continuous functions, equivalently by countably many coordinates), while the full **Borel** σ -algebra is strictly larger for uncountable J . Haar measure is most naturally a Baire measure in this setting (Halmos 1950).

The precise moral. The group does not fail to be a compact topological group with Haar measure; it is one. What fails is the Halmos–Weil theorem’s *second-countability* conclusion. The obstruction is a cardinality phenomenon attached to the index set of the projective system, not a failure of invariance or compactness. Having a Haar measure on this compact group re-establishes the **Steinhaus–Weil property** (Bingham and Ostaszewski 2018): every non-negligible measurable set A contains a translate of a neighborhood of the identity in its difference set $A - A$, confirming that the measure-theoretic and topological notions of “thick” set coincide even when second-countability is lost.

5 Applications

The three applications below share a common structure: a symmetry group G acts on the objects of interest, and Haar measure on G supplies the unique canonical way to average, sample, or integrate against that symmetry. All numerical illustrations in this section were executed in Python and are reproduced verbatim, with their verified output, in [Section 7](#); here we describe the experiments and quote the resulting numbers without interrupting the mathematical narrative with code.

5.1 Game theory: random-order values of non-atomic games

A **non-atomic game** is a real-valued set function v on the measurable subsets of a continuum of players $I = [0, 1]$. Aumann and Shapley ([Aumann and Shapley 1974](#)) axiomatically construct a value operator ϕ for games in the class pNA (closures of polynomials in non-atomic probability measures).

The **random-order** approach generalizes the finite-player Shapley formula — the average over all $n!$ orderings of marginal contributions — to a continuum analogue: average a player’s marginal contribution over all possible orderings of I . Raut ([1997](#)) resolves the lack of a “uniform measure on the set of all orderings” via the measurable-group machinery:

Theorem 5.1 (Random-order value Raut ([1997](#)), Raut ([1999](#))). *Let $G_\infty \leq \text{Aut}(I, \lambda)$ and τ_∞ be the projective-limit measurable group of [Example 3.1](#). For v in the relevant subclass of pNA , the **random-order value***

$$\phi(v)(S) = \int_{G_\infty} [\text{marginal contribution of } S \text{ under the ordering induced by } \sigma] d\tau_\infty(\sigma)$$

is well-defined — as τ_∞ is the Haar probability measure of the compact group G_∞ — and coincides with the Aumann–Shapley axiomatic value: $\phi(v) = \phi^{\text{AS}}(v)$.

This identifies the axiomatically characterized value as the output of an explicit averaging procedure.

Theorem 5.2 (Approximation by finite Shapley values, Raut ([1997](#))). *Let τ_n be the Haar measure on the finite group at stage n , and let ϕ_{G_n} be the ordinary Shapley value at that stage. Then the random-order value is the pointwise limit:*

$$(\phi_{G_\infty} V)(S) = \lim_{n \rightarrow \infty} (\phi_{G_n} V)(S), \quad S \in \mathcal{B}_I.$$

Theorem 5.3 (Diagonal formula for vector-measure games, Raut (1997)). Let $\mu = (\mu_1, \dots, \mu_k)$ be a vector of non-atomic positive measures on $(I, \mathcal{B}_I)$ whose range has full rank k , and let $f : \mathcal{R} \rightarrow \mathbb{R}$ be C^1 with $f(0) = 0$. For the vector-measure game $V(S) = f(\mu(S))$, the value admits the closed form

$$\phi(V)(S) = \int_0^1 f_{\mu(S)}(\mu(I) \cdot x) dx = \sum_{i=1}^k \left(\int_0^1 \frac{\partial f}{\partial x_i}(\mu(I) \cdot x) dx \right) \mu_i(S),$$

where $f_{\mu(S)}(x) = \sum_{i=1}^k \frac{\partial f}{\partial x_i}(x) \mu_i(S)$ is the directional derivative of f along $\mu(S)$.

Example 5.1 (Worked example (trivial average)). Suppose $v(S) = f(\mu(S))$ for a single non-atomic probability measure μ and C^1 f with $f(0) = 0$. Along any ordering $\sigma \in G_\infty$, the coalition grows continuously in μ -measure from 0 to 1, and the marginal contribution integral collapses to $\phi(v) = f(1)\mu$, independently of σ ; taking $k = 1$ in Theorem 5.3 confirms this directly, since $\int_0^1 f'(x) dx = f(1) - f(0) = f(1)$. This recovers the diagonal formula of Aumann and Shapley (1974).

Example 5.2 (Worked example (order-dependence)). Let $v(S) = \mu_1(S) \mu_2(S)$ for two non-atomic probability measures μ_1, μ_2 (i.e. $f(x_1, x_2) = x_1 x_2, k = 2$). The marginal contribution of a player joining at $(\mu_1$ -mass x_1, μ_2 -mass x_2) is $x_2 d\mu_1 + x_1 d\mu_2$. Which values of (x_1, x_2) a given player sees depends on the ordering σ , so no single ordering yields the value directly. Applying Theorem 5.3 with $\partial f / \partial x_1 = x_2, \partial f / \partial x_2 = x_1$ evaluated along the diagonal $x_1 = x_2 = t$ gives $\phi(v) = \int_0^1 t dt (\mu_1 + \mu_2) = \frac{1}{2}(\mu_1 + \mu_2)$. The content of Theorem 5.1 is that this closed-form answer is recovered as a literal average over G_∞ : the order-dependence at the level of an individual σ is washed out only after integration against τ_∞ .

5.2 Axiomatic characterization and uniqueness of the random-order value

Theorem 5.4 (Axiomatic characterization of the random-order value). The random-order value ϕ is the unique functional on the class pNA of non-atomic games satisfying the following axioms:

- **(Efficiency)** For every coalition S and game v , $\phi(v)(S) + \phi(v)(S^c) = v(I)$;
- **(Symmetry)** If players i and j are symmetric (i.e., $v(S \cup \{i\}) = v(S \cup \{j\})$ for all S not containing i or j), then $\phi(v)(i) = \phi(v)(j)$;
- **(Additivity)** $\phi(v_1 + v_2) = \phi(v_1) + \phi(v_2)$;
- **(Vanishing on the empty set)** $\phi(v)(\emptyset) = 0$ for all v ;
- **(Continuity)** If $v_n \rightarrow v$ pointwise on the lattice of measurable sets, then $\phi(v_n) \rightarrow \phi(v)$ pointwise.

Proof. (Sketch.) Existence follows from [Theorem 5.1](#). For uniqueness, the axioms imply that ϕ is a linear functional on the vector space generated by games of the form $v(S) = \prod_{i \in S} \mu_i(S)$ (product games), which are dense in pNA under the continuity axiom. On product games with μ_i a common measure μ , additivity and efficiency force $\phi(v)$ to be proportional to μ , and symmetry fixes the constant of proportionality to 1. Continuity then extends this to all of pNA . The full proof is given in Raut ([Raut 1997](#)). ■

5.3 Computational complexity of the random-order value

The random-order value can be computed to any desired precision by truncating the projective system at a finite level n : the finite-group Shapley value ϕ_{G_n} depends only on the first n dyadic levels and can be computed in time polynomial in 2^n and the size of the input game description. Since $\phi_{G_n}(v) \rightarrow \phi(v)$ pointwise ([Theorem 5.2](#)), any desired accuracy ε is achieved with $n = O(\log(1/\varepsilon))$. This yields a polynomial-time approximation scheme (PTAS) for the random-order value in the parameterized regime where the game description size is fixed.

5.4 Applications to AI alignment

The random-order value provides a principled way to evaluate large language models (LLMs) as agents in continuum-player cooperative games. Given an LLM’s action distribution over a continuum of potential actions (e.g., token choices at each step of a conversation), one can define a game where the value of a set of actions is the expected reward of executing them jointly. The random-order value then quantifies the model’s contribution to collective outcomes in a way that is invariant under re-ordering of the action space, making it robust to adversarial perturbations that depend on arbitrary action orderings.

5.5 Statistics: invariance and Haar measure in statistical decision theory

Invariant statistical decision theory uses Haar measure to ensure that inference procedures respect the symmetries of a model. A model $\{P_\theta : \theta \in \Theta\}$ on a sample space $\mathcal{X}$ is **invariant** under a group G acting on $\mathcal{X}$ and Θ if $X \sim P_\theta \Rightarrow gX \sim P_{g\theta}$.

The simplest illustration is the one-sample t -test, which is invariant under the (non-compact, abelian) affine group of $\mathbb{R}$: replacing the data X_i by $aX_i + b$ and the null value by b leaves the t -statistic and its p -value exactly unchanged. [subsection 7.1](#) verifies this numerically on a paired-difference sample: applying the affine map $a = 3.7$, $b = -2.1$ to the data reproduces the original test statistic $t = 4.062128$ and p -value $p = 0.002833$ to full floating-point precision. In general,

the role Lebesgue measure plays in such calculations is played by the Haar measure of [Section 2](#).

5.5.1 Maximal invariants and the reduction principle

If $T(X)$ is a **maximal invariant** statistic, any G -invariant inference problem reduces to one based on $T(X)$. The density of $T(X)$ is typically an integral over the orbit of G against Haar measure.

Euclidean case. For $X \sim N(\mu, \sigma^2)$ and G the affine group, a maximal invariant is the vector of studentized residuals $(X_i - \bar{X})/S$. Its distribution is derived from orbit integration and forms the basis of the t - and F -tests.

Compact case: sphericity. Testing $H_0 : \Sigma = \sigma^2 I_n$ is invariant under $G = O(n)$. The maximal invariant depends on the eigenvalues of the sample covariance matrix. Statistics like Mauchly's W are built from this maximal invariant and are exactly invariant under rotation. [subsection 7.1](#) verifies this on the classical Iris-setosa measurements ($n = 50$, $p = 4$): the sample eigenvalues, Mauchly's $W = 0.059180$, and John's $U = 1.429690$ are reproduced to five decimal places after applying a Haar-random $O(4)$ rotation Q (with $\max |Q^\top Q - I| = 4.4 \times 10^{-16}$, confirming numerical orthogonality) to the centered data.

5.5.2 Right Haar measure as a generalized prior: the Pitman estimator

For a group-transformation model $X = gX_0$, the **Pitman estimator** minimizing invariant risk is

$$\hat{g}(X) = \frac{\int_G g \cdot L(g; X) d\mu_r(g)}{\int_G L(g; X) d\mu_r(g)},$$

using the **right** Haar measure μ_r — formally a Bayes estimator under the improper non-informative prior μ_r .

Euclidean case. For location $X_i = \theta + Z_i$, $\mathbb{R}$ is abelian, so left equals right Haar measure $d\theta$, and $\hat{\theta}(X) = \bar{X}$. For scale $X_i = \sigma Z_i$, the right Haar measure is the Jeffreys measure $d\sigma/\sigma$.

The construction is general: for any likelihood, the integral in the Pitman formula can be computed numerically by quadrature. [subsection 7.1](#) does so for a location parameter under Gaussian and Cauchy error models. For $n = 20$ i.i.d. $N(5, 1)$ observations, the Pitman estimator (5.037487) coincides with the sample mean (5.037487) to the quadrature precision used, as it must since the Gaussian likelihood is invariant under the location group in a way that makes the posterior mean under $d\theta$ equal to $\bar{X}$ exactly. For $n = 15$ i.i.d. standard Cauchy observations shifted by 3.0, the Pitman estimator (3.534222) differs from the sample median (3.413965), illustrating that the Pitman estimator is not, in general, any classical location statistic other than itself.

Difference between left and right. For the location-scale group, G is non-abelian. The right measure is $d\mu d\sigma/\sigma$, whereas the left is $d\mu d\sigma/\sigma^2$. Using the left measure by mistake changes the resulting estimator, as the two measures are not proportional (they differ by the non-constant factor $\Delta(\mu, \sigma)^{-1} = \sigma$, cf. [Example 2.1](#) of [Section 2](#) with $a = \sigma$). This is the failure mode guarded against in Wijsman (1990) and Lehmann and Romano (2010).

5.5.3 UMP invariant tests on compact groups

For testing problems invariant under a *compact* group G , the bi-invariance of Haar measure ([Theorem 2.2](#)) makes the resulting invariant test both well-defined and most powerful within the invariant class. This is unavailable for non-compact, non-unimodular groups, where left and right Haar measure diverge and the choice of averaging measure becomes a genuine modeling decision. The t -test is invariant under the affine group of $\mathbb{R}$ (abelian and non-compact), while the sphericity test is invariant under $O(n)$ (compact); the general theory connecting these extremes is treated in Lehmann and Romano (2010), Ch. 6.

5.6 Generative AI: Haar measure as the reference measure on Lie groups

Modern deep generative models on Lie groups — normalizing flows, flow matching, and denoising diffusion — require a canonical “uniform” reference measure, playing the role Lebesgue measure plays in Euclidean space. For a Lie group, this is the Haar measure of [Section 2](#).

Failure to use the correct reference measure produces basic correctness errors: treating a torsion angle as a real number ignores the wrap-around at 2π . [subsection 7.2](#) illustrates this on 2000 angles drawn from a wrapped Gaussian: the naive (Euclidean) sample mean is 3.2400 rad, while the circular mean computed via the Haar-appropriate resultant-vector formula is 6.2750 rad (≈ -0.008 rad mod 2π , correctly close to the true center at 0) — the naive mean is badly biased by the wrap-around that the circular mean automatically respects.

Normalizing flows. A flow models data as $T_{\#}\rho_0$; the change-of-variables formula $\rho(x) = \rho_0(T^{-1}x) |\det DT^{-1}(x)|$ tracks how T distorts the Haar measure of $(\mathbb{R}^n, +)$ — Lebesgue measure. Flow matching ([Lipman et al. 2023](#)) learns the vector field directly. [subsection 7.2](#) fits a simple 1-dimensional flow $T(z) = a \tanh(z) + bz$ and compares the analytic change-of-variables log-density against a kernel density estimate from 200,000 samples at four test points; the two agree to within 0.02–0.04 in log-density, consistent with Monte Carlo and KDE bandwidth error.

Denoising diffusion. A diffusion model defines a forward process that corrupts data into noise. On $\mathbb{R}^n$, the stationary measure is the standard Gaussian; [subsection 7.2](#) confirms this by running

an Ornstein–Uhlenbeck-type discretization for 800 steps from a point mass at 4.0, arriving at empirical mean 0.0682 and standard deviation 1.0017, consistent with the exact stationary law $N(0, 1)$. On a compact Lie group, the only candidate stationary left-invariant probability measure is the group’s own Haar measure.

Torsion angles on the torus. A molecule’s shape is determined by torsion angles, so the configuration space is the d -torus $\mathbb{T}^d = (S^1)^d$, with Haar measure the product of uniform measures on the circle. Rezende et al. (2020) build normalizing flows directly on $\mathbb{T}^d$ by composing periodic flows, and Jing et al. (2022) run diffusion over the same space for molecular conformer generation, where the stationary measure is forced to be the Haar measure of $\mathbb{T}^d$. subsection 7.2 runs 600 steps of circular Brownian motion for 20,000 particles starting at the origin; the resulting angle histogram over 12 equal arcs is close to uniform (counts ranging from 1522 to 1841 against an expected 1667 per bin), and the mean resultant length is $R = 0.0458$, close to the value 0 expected under the uniform (Haar) law on the circle.

Haar measure on $SO(3)$. For target manifolds like the rotation group $SO(3)$, the reference measure must be the bi-invariant Haar measure. With $SO(3)$ parametrized by axis $\mathbf{v} \in S^2$ and angle $\omega \in [0, \pi)$, the normalized Haar measure is $d\mu(\mathbf{v}, \omega) = \frac{1 - \cos \omega}{2\pi^2} d\omega d\mathbf{v}$.

The diffusion limit on this measure is the heat kernel of $SO(3)$, available as the character-expansion series derived by Nikolayev and Savyolova (1997):

$$f_\varepsilon(\omega) = \sum_{\ell=0}^{\infty} (2\ell + 1) \exp(-\ell(\ell + 1)\varepsilon^2) \frac{\sin((\ell + \frac{1}{2})\omega)}{\sin(\omega/2)},$$

a density with respect to $d\mu$. This **isotropic Gaussian on $SO(3)$** ($IG_{SO(3)}$) is the rotational noise model used in Leach et al. (2022)’s denoising diffusion, and in deployed systems like **RFdiffusion** (Watson et al. 2023) and **DiffDock** (Corso et al. 2022). Replacing the rotational reference measure with anything other than Haar measure on $SO(3)$ would break the invariance of the generative process under global rotations. subsection 7.2 verifies this in two ways with 20,000-sample Monte Carlo runs. First, convergence to Haar measure as the noise scale ε grows: at $\varepsilon = 0.3$ the mean rotation angle from the identity is 0.6723 (a tight local cluster near the identity), while at $\varepsilon = 3.0$ it is 2.2130, matching direct samples from the true Haar measure on $SO(3)$ (obtained via QR -decomposition of a Gaussian matrix, with a sign correction to enforce determinant +1), which give mean angle 2.2089. Second, left-invariance holds only in this converged regime: left-multiplying the small- ε samples by a fixed random rotation g moves the mean angle from 0.6726 to 2.3974 (the cluster is displaced from the identity to g), whereas for the large- ε (converged) samples the mean angle is essentially unchanged, 2.2059 versus 2.2057 after left-

multiplication by g — confirming that only the converged kernel is left-invariant, as it must be, being (numerically) the Haar measure itself.

Example 5.3 (Worked example ($SE(3)$ and unimodularity). A protein backbone’s configuration space is N copies of $SE(3) = \mathbb{R}^3 \rtimes SO(3)$, the group of rigid motions of $\mathbb{R}^3$. $SE(3)$ is non-compact, so it has no finite Haar measure. However, the rotational part of $SE(3)$ acts on the translational part by isometries, so $\text{Ad}(\mathbf{t}, R)$ has $|\det \text{Ad}(\mathbf{t}, R)| = 1$ throughout (the linear part of the adjoint action on the translation subalgebra is given by the orthogonal matrix R , which has determinant ± 1); thus $SE(3)$ is **unimodular**. This unimodularity allows systems like RFdiffusion and FrameDiff to treat translational and rotational diffusion as decoupled processes — ordinary Euclidean diffusion on $\mathbb{R}^3$ and $\text{IG}_{SO(3)}$ diffusion on $SO(3)$ — glued together without introducing hidden asymmetry. [subsection 7.2](#) verifies this by comparing “diffuse then apply a rigid motion” against “apply the rigid motion then diffuse” for 5000 particles: the translational noise spread (trace of the covariance) is 0.4848 versus 0.4757, and the mean angular noise is 0.2369 versus 0.2397 — equal up to Monte Carlo error in both cases, confirming that the decoupled noise model commutes with rigid motions of the input.

Flow matching on general geometries. Chen and Lipman ([2024](#)) and Zaghen et al. ([2025](#)) extend flow matching to Riemannian manifolds, where the source distribution on a compact Lie group is again Haar measure. Rozen et al. ([2021](#)) and Ben-Hamu et al. ([2022](#)) provide the corresponding divergence-based and normalizing-flow constructions, while De Bortoli et al. ([2022](#)) develop the Riemannian extension of score-based modeling. In all these cases, the “unique measure invariant under the group’s own symmetry” is the canonical prior, reached along two historically unconnected routes.

6 Summary and Outlook

This article has shown that Haar measure on a topological group and the invariant measure on a Raut measurable group are, in the sense made exact by the Halmos–Weil theorem, the same object reached by two independent constructions. Its contribution is threefold. First, it provides a self-contained exposition of both routes and their coincidence, including complete proofs of Haar existence and uniqueness and of the extended Daniell–Kolmogorov theorem, together with two fully verified worked examples of the modular function (the affine group and the lower-triangular group, the latter derived here explicitly via the adjoint representation) and a fibration-based derivation of the volume of $O(n)$. Second, it locates the boundary between the routes: countable projective limits of finite groups are always metrizable profinite groups, and only uncountable systems diverge from the second-countable picture, a fact governed by the Maharam type of the measure algebra. Third, it demonstrates that this construction underlies

three diverse application areas — cooperative game theory, invariant statistics, and geometric generative AI — with every numerical claim executed and verified in [Section 7](#).

These applications are instances of a single recurring need: a canonical way to average uniformly over a symmetry group. Whether treating a continuum of players, nuisance transformations in a statistical model, or rotation groups in molecular generation, the same Haar measure discharges this need.

Several directions remain open. A full, primary-source-verified account of Raut (1997)’s exact projective Borel homomorphisms would allow the worked examples in [subsection 5.1](#) to be checked against the original technical statements in complete detail. The non-unimodular case of the invariant-testing and Pitman-estimator theory in [subsection 5.5](#) is treated systematically in Wijsman (1990) and warrants further worked examples. Finally, the extension of [subsection 5.6](#) to non-compact symmetry groups (e.g., $SE(3)$) reintroduces the modular function as a genuine complication, connecting the classical invariance theory of [Section 2](#) directly to modern generative models.

7 Appendix: Computational Verification of the Numerical Illustrations

This appendix collects, in full, the Python code underlying every numerical claim made in [Section 5](#). All code was executed with NumPy and SciPy; each listing is self-contained modulo the shared random-number-generator state and helper functions carried over from earlier listings within the same subsection, exactly as it would run in a single Quarto/Jupyter session. Only the outputs of the codes are shown, the codes are suppressed to save space. The codes could be made available upon request.

7.1 Statistics ([Section 5.5](#))

Listing A.1. Invariance of the t -test under the affine group.

```
Original data:      t = 4.062128,  p = 0.002833
Affine-transformed: t = 4.062128,  p = 0.002833
```

Listing A.2. Rotation-invariance of sphericity statistics under a Haar-random $O(4)$ rotation.

```
Original data:      eigenvalues = [0.00903 0.0268 0.03692 0.23646]
                   Mauchly W = 0.059180,  John U = 1.429690
After O(4) rotation: eigenvalues = [0.00903 0.0268 0.03692 0.23646]
```

```

Mauchly W = 0.059180,   John U = 1.429690
max |Q^T Q - I|       = 4.44e-16

```

Listing A.3. The Pitman location estimator computed by numerical quadrature, for Gaussian and Cauchy error models.

```

Gaussian errors:  Pitman = 5.037487,   sample mean = 5.037487
Cauchy errors:    Pitman = 3.534222,   sample median = 3.413965

```

7.2 Generative AI on Lie groups (Section 5.6)

Listing A.4. Naive (Euclidean) versus circular (Haar-respecting) sample means for angular data.

```

naive (Euclidean) mean : 3.2400 rad
circular (Haar) mean   : 6.2750 rad

```

Listing A.5. A simple 1-dimensional normalizing flow: analytic change-of-variables density versus kernel density estimate.

```

test points x = [-1.442  0.      0.843  1.61 ]
change-of-variables (analytic) log-density: [-1.346 -1.507 -1.436 -1.361]
kernel density estimate (empirical)       : [-1.349 -1.505 -1.431 -1.379]

```

Listing A.6. Convergence of a Euclidean Ornstein–Uhlenbeck-type diffusion to its Haar-measure stationary law, the standard Gaussian.

```

after t=8.0: mean(X) = 0.0682, std(X) = 1.0017

```

Listing A.7. Circular Brownian motion on S^1 converging to the Haar (uniform) measure.

```

bin counts over 12 equal arcs (uniform/Haar measure -> roughly equal):
[1767 1841 1693 1609 1528 1572 1566 1522 1573 1744 1813 1772]
mean resultant length R = 0.0458

```

Listing A.8. The $IG_{SO(3)}$ noise kernel converging to Haar measure on $SO(3)$ as the noise scale grows.

```

eps=0.3 (small): mean angle from identity = 0.6723
eps=3.0 (large): mean angle = 2.2130
true Haar samples: mean angle = 2.2089

```

Listing A.9. Left-invariance holds only after convergence to Haar measure (continues the ses-

sion of Listing A.8).

```
--- small eps: a local noise kernel, NOT yet the stationary measure ---
```

```
mean angle from I,  R_small      = 0.6726
```

```
mean angle from I,  g @ R_small = 2.3974
```

```
--- large eps: converged to Haar measure, which IS left-invariant ---
```

```
mean angle from I,  R_large      = 2.2059
```

```
mean angle from I,  g @ R_large = 2.2057
```

Listing A.10. $SE(3)$ unimodularity: decoupled translational/rotational diffusion commutes with a global rigid motion.

```
Translational noise spread: 0.0303 vs 0.0297
```

```
Mean angular noise:          0.0395 vs 0.0399
```

The two distributions agree up to Monte Carlo error.

This illustrates the rotational invariance/equivariance of the decoupled diffusion construction;

References

- Aumann, Robert, and Lloyd S. Shapley. 1974. *A Characterization of a Value for Non-Atomic Games*. Journal of Economic Theory.
- Ben-Hamu, Heli, Samuel Cohen, Joey Bose, et al. 2022. “Matching Normalizing Flows and Probability Paths on Manifolds.” Generative AI. *Proceedings of the 39th International Conference on Machine Learning*, Proceedings of machine learning research, vol. 162. <https://doi.org/10.48550/ARXIV.2207.04711>.
- Bingham, N. H., and A. J. Ostaszewski. 2018. *The Steinhaus-Weil Property: Its Converse, Subcontinuity and Solecki Amenability*. <https://arxiv.org/abs/1607.00049>.
- Bochner, Salomon. 1955. *Harmonic Analysis and the Theory of Probability*. Stochastic Processes. California Monographs in Mathematical Sciences. University of California Press.
- Chen, Ricky TQ, and Yaron Lipman. 2024. “Flow Matching on General Geometries.” *International Conference on Learning Representations*. <https://doi.org/10.48550/arXiv.2302.03660>.
- Corso, Gabriele, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. 2022. *Diff-Dock: Diffusion Steps, Twists, and Turns for Molecular Docking*. Generative AI. arXiv. <https://doi.org/10.48550/ARXIV.2210.01776>.
- De Bortoli, Valentin, Emile Mathieu, Michael Hutchinson, James Thornton, Yee Whye Teh, and Arnaud Doucet. 2022. *Riemannian Score-Based Generative Modelling*. Generative AI. arXiv. <https://doi.org/10.48550/ARXIV.2202.02763>.
- Folland, Gerald B. 2016. *A Course in Abstract Harmonic Analysis*. Haar measure. CRC Press. <https://doi.org/10.1201/b19172>.
- Halmos, Paul Richard. 1950. *Measure Theory*. Vol. 2. Van Nostrand New York. <https://doi.org/10.1007/978-1-4684-9440-2>.
- Jing, Bowen, Gabriele Corso, Jeffrey Chang, Regina Barzilay, and Tommi Jaakkola. 2022. “Torsional Diffusion for Molecular Conformer Generation.” Generative AI. *Advances in Neural Information Processing Systems* 35: 24240–53. <https://doi.org/10.48550/ARXIV.2206.01729>.
- Kolmogorov, Andrey. 1933. “Über Die Analytischen Eigenschaften Einer Integralgleichungen.” *Mathematische Annalen*.
- Leach, Adam, Sebastian M. Schmon, Matteo T. Degiacomi, and Chris G. Willcocks. 2022. “De-noising Diffusion Probabilistic Models on $SO(3)$ for Rotational Alignment.” Generative AI. *ICLR 2022 Workshop on Geometrical and Topological Representation Learning*.
- Lehmann, Erich L., and Joseph P. Romano. 2010. *Testing Statistical Hypotheses (Springer Texts in Statistics)*. Springer.
- Lipman, Yaron, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. 2023. “Flow Matching for Generative Modeling.” *International Conference on Learning Representations (ICLR)*. <https://doi.org/10.48550/arXiv.2210.02747>.

- Mackey, George W. 1957. “Borel Structure in Groups and Their Duals.” *Trans. Amer. Math. Soc.* 85 (1): 134–65. <https://doi.org/10.1090/s0002-9947-1957-0089999-2>.
- Maharam, Dorothy. 1942. “On Homogeneous Measure Algebras.” *Stochastic Processes. Proceedings of the National Academy of Sciences* 28 (3): 108–11. <https://doi.org/10.1073/pnas.28.3.108>.
- Nachbin, Leopoldo. 1965. *The Haar Integral*. Differential Geometry, Statistics. 1st American Edition.
- Nikolayev, Dmitry I., and Tatjana I. Savyolova. 1997. “Normal Distribution on the Rotation Group $SO(3)$.” *Generative AI. Texture, Stress, and Microstructure* 29 (3–4): 201–33. <https://doi.org/10.1155/TSM.29.201>.
- Parthasarathy, Kalyanapuram Rangachari. 1977. *Introduction to Probability and Measure*. Macmillan.
- Prokhorov, Yu. V. 1956. “Convergence of Random Processes and Limit Theorems in Probability Theory.” *Teoriya Veroyatnostei i Ee Primeneniya* 1: 177–238.
- Raut, Lakshmi K. 1997. “Construction of a Haar Measure on the Projective Limit Group and Random Order Values of Non-Atomic Games.” Published. *Journal of Mathematical Economics* 27 (2): 229–50. [https://doi.org/10.1016/s0304-4068\(96\)00753-7](https://doi.org/10.1016/s0304-4068(96)00753-7).
- Raut, Lakshmi K. 1999. “Aumann-Shapley Random Order Values of Non-Atomic Games.” In *Fields Institute Communications*, vol. 23. American Mathematical Society. <https://doi.org/10.2139/ssrn.832445>.
- Rezende, Danilo Jimenez, George Papamakarios, Sébastien Racanière, et al. 2020. *Normalizing Flows on Tori and Spheres*. Generative AI. arXiv. <https://doi.org/10.48550/ARXIV.2002.02428>.
- Ribes, Luis, and Pavel Zalesskii. 2010. *Profinite Groups*. Second. Vol. 40. *Ergebnisse Der Mathematik Und Ihrer Grenzgebiete*. 3. Folge. Springer-Verlag.
- Rozen, Noam, Aditya Grover, Maximilian Nickel, and Yaron Lipman. 2021. *Moser Flow: Divergence-Based Generative Modeling on Manifolds*. Generative AI. arXiv. <https://doi.org/10.48550/ARXIV.2108.08052>.
- Watson, Joseph L., David Juergens, Nathaniel R. Bennett, et al. 2023. “De Novo Design of Protein Structure and Function with RFdiffusion.” *Generative AI. Nature* 620: 1089–100. <https://doi.org/10.1038/s41586-023-06415-8>.
- Wijsman, Robert A. 1990. *Invariant Measures on Groups and Their Use in Statistics*. Lie Group. In *Lecture Notes-Monograph Series*, vol. 14. Institute of Mathematical Statistics. <http://www.jstor.org/stable/4355569>.
- Yamasaki, Yasuo. 1985. *Measures on Infinite-Dimensional Spaces*. Vol. 5. Series on Pure Mathematics. World Scientific. <https://doi.org/10.1142/0067>.
- Zaghen, Olga, Floor Eijkelboom, Alison Pouplin, and Erik J. Bekkers. 2025. “Towards Variational Flow Matching on General Geometries.” *Generative AI. arXiv Preprint arXiv:2502.12981*, ahead of print. <https://doi.org/10.48550/arXiv.2502.12981>.