QUOTIENT EM UNDER MISSPECIFICATION: TIGHT LOCAL RATES AND FINITE-SAMPLE BOUNDS IN GENERAL INTEGRAL PROBABILITY METRICS

KOUSTAV MALLIK

Abstract. We study the expectation-maximization (EM) algorithm for general latent-variable models under (i) distributional misspecification and (ii) nonidentifiability induced by a group action. We formulate EM on the quotient parameter space and measure error using an arbitrary integral probability metric (IPM). Our main results give (a) a sharp local linear convergence rate for population EM governed by the spectral radius of the linearization on a local slice, and (b) tight finite-sample bounds for sample EM obtained via perturbed contraction inequalities and generic chaining/entropy control of EM-induced empirical processes.

Contents

2020 *Mathematics Subject Classification.* 62F10, 62F12, 62H12, 62G20.

Key words and phrases. EM algorithm, nonidentifiability, quotient space, misspecification, integral probability metrics, generic chaining.

1. Introduction

1.1. **Misspecification, symmetry, and distributional error criteria.** Let $\mathcal{X}$ denote an observation space and let $\{p_\theta(x,z) : \theta \in \Theta\}$ be a dominated family of latent-variable models on $\mathcal{X} \times \mathcal{Z}$, with observed marginal $p_\theta(x) = \int_{\mathcal{Z}} p_\theta(x,z)\,dz$ and induced distribution P_θ on $\mathcal{X}$. Given i.i.d. data $X_1, \dots, X_n \sim P^*$ from an *arbitrary* distribution P^* on $\mathcal{X}$, consider the population (cross-entropy) objective

$$\Phi(\theta) \;:=\; \mathbb{E}_{X\sim P^*}\big[\log p_\theta(X)\big]. \tag{1}$$

When $P^* \notin \{P_\theta\}$, the appropriate statistical target is the *KL projection set*

$$\bar{\Theta} \;:=\; \arg\max_{\theta\in\Theta} \Phi(\theta) \;=\; \arg\min_{\theta\in\Theta} \mathrm{KL}(P^* \,\|\, P_\theta), \tag{2}$$

where the equivalence follows from the elementary identity $\mathrm{KL}(P^* \,\|\, P_\theta) = \mathbb{E}_{P^*}[\log p^*(X)] - \Phi(\theta)$. In misspecified settings, $\bar{\Theta}$ may be non-singleton and the curvature of Φ at $\bar{\Theta}$ need not correspond to any "true parameter"; nonetheless $\bar{\Theta}$ is the canonical benchmark for both statistical and algorithmic analysis [5, 6].

A second structural feature is *nonidentifiability induced by symmetry*. We model nonidentifiability by a group action $G \curvearrowright \Theta$ satisfying observed invariance

$$P_{g\cdot\theta} \;=\; P_\theta, \qquad \forall g \in G,\ \theta \in \Theta. \tag{3}$$

Invariance (3) arises in finite mixtures (label switching), symmetric mixtures (sign flips), and factor models (rotations), among many others. Under (3), the map $\theta \mapsto P_\theta$ factors through the

quotient $\Theta/G := \Theta/G$, and the statistically meaningful estimand is the orbit $[\theta] \in \Theta/G$ rather than the ambient parameter θ. This elementary observation has a consequential implication: any error criterion that distinguishes points *within* a single orbit cannot be an appropriate target for consistent estimation (and is typically incompatible with any meaningful algorithmic rate statement). Accordingly, our analysis is formulated on Θ/G; see also recent symmetry-aware estimation frameworks in related contexts [7].

A third feature, closely tied to misspecification and nonidentifiability, is the choice of *error metric*. Rather than measuring error in an ambient parameter norm, we measure discrepancy between distributions using an *integral probability metric* (IPM) generated by an arbitrary function class $\mathcal{F}$,

$$d_{\mathcal{F}}(P, Q) \;:=\; \sup_{f \in \mathcal{F}} \big|\mathbb{E}_P f - \mathbb{E}_Q f\big|. \tag{4}$$

This framework encompasses total variation, bounded-Lipschitz and Wasserstein-type metrics, and kernel-based IPMs, among others [8, 9]. Because of (3), $d_{\mathcal{F}}$ induces a canonical quotient metric

$$\bar{d}_{\mathcal{F}}([\theta]\,, [\theta']) \;:=\; d_{\mathcal{F}}(P_\theta, P_{\theta'}), \tag{5}$$

and the natural performance criterion relative to the misspecified target is the distance-to-set functional

$$\operatorname{dist}_{\mathcal{F}}([\theta]\,, \bar{\Theta}/G) \;:=\; \inf_{\bar{\theta} \in \bar{\Theta}} \bar{d}_{\mathcal{F}}([\theta]\,, [\bar{\theta}]) = \inf_{\bar{\theta} \in \bar{\Theta}} d_{\mathcal{F}}(P_\theta, P_{\bar{\theta}}). \tag{6}$$

The aim of this work is to establish *sharp local algorithmic rates* and *finite-sample guarantees* for EM in the general regime described by (2)–(6): arbitrary misspecification, group nonidentifiability, and an arbitrary IPM.

1.2. **EM maps and the quotient state space.** The EM algorithm is classically derived as a minorization–maximization scheme for the observed log-likelihood [4, 12]. Given a current iterate θ, define the population surrogate

$$Q(\theta'; \theta) := \mathbb{E}_{X \sim P^*}\Big[\mathbb{E}_{Z \sim p_\theta(\cdot|X)} \log p_{\theta'}(X, Z)\Big], \qquad M(\theta) \in \arg\max_{\theta' \in \Theta} Q(\theta'; \theta), \tag{7}$$

and define sample analogs $(\widehat{Q}_n, \widehat{M}_n)$ by replacing the outer expectation in (7) by $n^{-1}\sum_{i=1}^n(\cdot)$. The fundamental ascent property—$\Phi(M(\theta)) \geq \Phi(\theta)$ and $\Phi_n(\widehat{M}_n(\theta)) \geq \Phi_n(\theta)$ for exact M-steps—is well known [4, 12].

Nonasymptotic analyses of EM typically proceed by (i) proving that the *population* EM operator is contractive on a basin around a target point, and (ii) showing that the *sample* operator is a uniform perturbation of the population operator on that basin [1]. Under misspecification, the appropriate target is the projection set (2) and the relevant local geometry is governed by the curvature of the misspecified objective at that set [5, 6]. In the presence of symmetry (3), however, both Φ and $\theta \mapsto P_\theta$ are constant along orbits, and the natural state space for EM is the quotient Θ/G. A first step is therefore to show that EM descends to a well-defined quotient operator (possibly set-valued), and to formulate contraction and perturbation results directly in quotient-compatible coordinates.

We adopt a local "slice" strategy: in a neighborhood of interest in Θ/G, select a measurable section $S \subset \Theta$ that intersects each orbit exactly once. This eliminates tangential directions along group orbits and yields locally identifiable coordinates in which linearization and sharp local rates are meaningful. The quotient/IPM formulation then yields guarantees that are invariant

by construction under nonidentifiable directions, and are directly interpretable as distributional convergence.

1.3. **Contributions and main results.** We summarize the contributions in terms of the quantitative constants appearing in the theorems.

(1) Quotient IPM formalism and EM well-posedness. Under (3), the quotient IPM (5) is well-defined and induces the canonical error criterion (6). Under mild equivariance conditions, EM induces a quotient map $\bar{M} : \Theta/G \to \Theta/G$ (possibly set-valued), and retains ascent for Φ and Φ_n (Section 5). These results formalize the basic requirement that any EM analysis in symmetric models must be orbit-invariant.

(2) Sharp local linear rate for population EM (spectral-radius governed). Let T denote the slice-restricted population EM map near a fixed-point representative $\theta^\star \in S$. Assuming differentiability of T at $\theta^\star$ and $\rho(DT(\theta^\star)) < 1$, we prove that for any $\varepsilon > 0$ there exists a neighborhood in which

$$\|\theta_t - \theta^\star\| \;\le\; C_\varepsilon\big(\rho(DT(\theta^\star)) + \varepsilon\big)^t \|\theta_0 - \theta^\star\|, \tag{8}$$

for all initializations θ_0 in that neighborhood. The rate constant $\rho(DT(\theta^\star))$ is the correct sharp linear factor for smooth fixed-point iterations (Section 6).

(3) Convergence in an arbitrary IPM. To translate (8) into a distributional statement, we assume a local regularity modulus for the model map $\theta \mapsto P_\theta$ into $d_{\mathcal{F}}$ on the slice basin (e.g., local Lipschitz continuity). This yields geometric decay of $d_{\mathcal{F}}(P_{\theta_t}, P_{\theta^\star})$ (and the corresponding quotient distance-to-set) with the same geometric factor as in slice coordinates. We also provide a distance-to-set formulation that avoids choosing a representative when $\bar{\Theta}/G$ is set-valued (Section 7).

(4) Tight finite-sample bounds for sample EM via perturbed contractions. On a contraction basin where T is γ-contractive, we treat sample EM as a perturbation $\widehat{T}_n$ of T and show that a uniform operator deviation bound

$$\sup_{\theta \in B} \|\widehat{T}_n(\theta) - T(\theta)\| \le \Delta_n(\delta) \quad \text{with probability at least } 1 - \delta$$

implies the *exact* perturbed-contraction bound

$$\|\hat{\theta}_t - \theta^\star\| \;\le\; \gamma^t \|\hat{\theta}_0 - \theta^\star\| + \frac{1-\gamma^t}{1-\gamma}\Delta_n(\delta), \tag{9}$$

and therefore an IPM bound of the same form after transfer. The amplification factor $(1-\gamma)^{-1}$ in (9) is unavoidable for bounded additive perturbations of contractions and is therefore the correct notion of tightness at this level of generality (Section 8).

(5) Operator deviation bounds via EM-induced empirical processes (generic chaining). The deviation term $\Delta_n(\delta)$ is controlled not by $\mathcal{F}$ itself, but by an *EM-induced class* $\mathcal{H}$ (typically involving gradients and/or Hessians of the EM surrogate over the basin). We derive high-probability bounds of order

$$\sup_{h \in \mathcal{H}} \big|(P_n - P^*)h\big| \;\lesssim\; \frac{\gamma_2(\mathcal{H}, L_2(P^*))}{\sqrt{n}} + \|H\|_{L_2(P^*)}\sqrt{\frac{\log(1/\delta)}{n}}, \tag{10}$$

with corresponding bracketing-entropy variants; see [10, 11]. This form is essentially the sharpest available for general function classes, and it clarifies why complexity control is necessary when $\mathcal{F}$ (and hence $\mathcal{H}$) is allowed to be arbitrary (Section 8).

(6) Extensions and limits. We discuss algorithmic variants (inexact EM, approximate E-steps, stochastic EM, regularization, annealing, multi-start) and show how the quotient contraction framework yields corresponding guarantees. We also include an impossibility statement formalizing that, absent complexity control (e.g., bounded γ_2 or entropy integrals), no uniform finite-sample guarantee in a general IPM can hold. Where appropriate, we connect to algorithmic stability [3] and global convergence templates for nonconvex alternating schemes under KŁ regularity [2].

1.4. **Organization and notation.** Section 3 introduces the model, misspecification target, and IPMs. Section 4 develops quotient geometry, the quotient IPM, and local slices. Section 5 establishes quotient well-posedness of EM and ascent properties. Section 6 proves sharp local linear rates for population EM via linearization on a slice, and Section 7 transfers these rates to general IPMs, including distance-to-set formulations. Section 8 provides finite-sample guarantees for sample EM via perturbed contractions and generic chaining/entropy bounds for EM-induced classes. Section 9 treats extensions, and Section 1.3 records lower bounds and asymptotic refinements.

Notation. We write $[\theta]$ for the orbit of θ, $\Theta/G = \Theta/G$ for the quotient, and S for a local section. Constants c, C may change from line to line. All probability statements are with respect to $X_1, \dots, X_n \sim P^*$ unless stated otherwise. "Sharp" refers to spectral-radius governed local linear rates for smooth fixed-point maps, and "tight" refers to the exact perturbation amplification $(1-\gamma)^{-1}$ for bounded additive perturbations of contractions.

2. Related work

The results in this paper draw on several well-developed literatures: classical EM theory (ascent and convergence), nonasymptotic analyses that relate sample EM to population EM, misspecification and singularity effects on EM rates, symmetry/nonidentifiability and orbit-space viewpoints, distributional metrics (in particular IPMs) and their empirical estimation, and empirical process tools (generic chaining/entropy) that yield sharp uniform deviation bounds. We review only what is needed to (i) position our contributions precisely and (ii) justify the technical choices made in Sections 4–8.

2.1. **Classical EM: ascent, fixed points, and convergence.** The EM algorithm was introduced by Dempster, Laird, and Rubin as a general procedure for maximum likelihood estimation with latent or incomplete data, with the fundamental monotonicity property that the observed-data likelihood does not decrease under exact M-steps [4]. Wu established broad convergence guarantees for EM, including conditions under which the likelihood sequence converges and limit points of the iterates are stationary points of the observed likelihood [12]. These classical foundations motivate our use of EM as a majorization–maximization scheme and underwrite the ascent statements formalized in Section 5.

2.2. **Modern nonasymptotic EM: population-to-sample perturbation theory.** A central theme in contemporary analyses is to treat sample EM as a perturbation of the population EM map: (i) establish contraction (or more generally stability) of the population operator on a basin, and (ii) control the deviation between the sample operator and the population operator uniformly on that basin. Balakrishnan, Wainwright, and Yu developed a systematic version of this program and derived finite-sample guarantees for EM in canonical latent-variable settings [1]. Our finite-sample results in Section 8 adopt this perturbation viewpoint; however, we formulate it

on the quotient state space and in an IPM, which requires quotient-compatible stability notions and distributional transfer moduli.

2.3. **Misspecification and singularity effects on EM rates.** Under misspecification, EM ascends the cross-entropy objective (1) and converges (when it converges) to the KL projection set (2) rather than to a ground-truth parameter. Dwivedi et al. established sharp guarantees for EM under misspecified Gaussian mixture models and analyzed how model mismatch alters contraction behavior and basin geometry [5]. Subsequent work studied the impact of singularities and misspecification on EM convergence rates, including regimes with nonstandard rates driven by local degeneracy of curvature [6]. These results motivate two design principles in the present paper: we state targets relative to the projection set and keep local contraction constants explicit, because both may depend sensitively on misspecification and local geometry.

2.4. **Symmetry, nonidentifiability, and orbit-space viewpoints.** Nonidentifiability induced by group symmetries is ubiquitous in latent-variable models (e.g., label switching in mixtures), and it implies that the statistically meaningful estimand is an orbit in the quotient Θ/G. While classical EM convergence statements are often expressed in ambient parameters, symmetry implies that ambient distances can be ill-posed or uninterpretable: points within an orbit are observationally indistinguishable. Recent symmetry-aware estimation perspectives emphasize orbit-space formulations and normal-direction geometry; see, for example, [7] for a modern orbit-space viewpoint in folded mixture models and related settings. Our contribution is to incorporate this orbit-space perspective into a general EM perturbation theory: we formulate EM on the quotient, implement analysis on local slices, and report convergence in quotient/IPM distances (Sections 4–7).

2.5. **Integral probability metrics and empirical estimation.** Integral probability metrics provide a unified framework for distributional discrepancies, generalizing classical metrics and encompassing many modern choices (e.g., bounded-Lipschitz, Wasserstein-type, and kernel IPMs) [8]. The empirical estimation of IPMs and associated concentration phenomena have been characterized in generality, with rates determined by the complexity of the generating class (and related entropy/chaining functionals) [9]. In our setting, IPMs enter in two distinct roles: (i) as the final error metric for the discrepancy between $P_{\hat{\theta}}$ and the projection set, and (ii) indirectly, through the complexity of the EM-induced function classes that govern sample-to-population deviations in the EM operator (Section 8).

2.6. **Empirical process tools: generic chaining and entropy.** To obtain sharp uniform deviation bounds over general function classes, we rely on the empirical process machinery based on generic chaining and entropy. Talagrand's monograph provides the canonical development of γ_2-functionals and sharp bounds for suprema of stochastic processes [10]. The monograph of van der Vaart and Wellner develops entropy and bracketing methods and provides the foundational tools for uniform laws of large numbers and Donsker theory [11]. Because we allow an arbitrary IPM and a model-dependent induced class, it is neither possible nor appropriate to assume a priori VC/RKHS/Lipschitz structure in the core theorems. Accordingly, we state the operator deviation term in intrinsic complexity form (generic chaining / bracketing entropy), and provide specialization pathways in later sections.

2.7. **Algorithmic stability as an alternative generalization route.** Uniform stability yields generalization guarantees for learning algorithms without explicit entropy control in some regimes. Bousquet and Elisseeff introduced a general stability framework that bounds the generalization gap by the sensitivity of the algorithm to the replacement of a single sample [3]. In Section 8 we

include an optional stability-based complement to the empirical-process route, formulated in slice/quotient coordinates and transferred to IPM error through the model regularity modulus.

2.8. **Nonconvex alternating schemes and global convergence templates.** EM may be viewed as an alternating maximization/minorization method. General global convergence results for nonconvex and nonsmooth alternating schemes are often obtained under Kurdyka–Łojasiewicz (KŁ) regularity combined with sufficient-decrease and relative-error conditions. Bolte, Sabach, and Teboulle developed a widely used template in this direction [2]. While our primary contributions are local and quantitative (sharp local linear rates and tight finite-sample perturbation bounds), Section 9 discusses how quotient formulations interface with KŁ-type global convergence arguments.

2.9. **Summary of distinctions.** Relative to [4, 12, 1, 5, 6], the present work makes three principled changes that are necessary for correctness and interpretability in the regimes we target:

(1) *Orbit-invariant state space:* we formulate EM on Θ/G and analyze it via local slices, ensuring that all convergence statements are invariant under nonidentifiable directions [7].
(2) *Misspecification-first targets:* we state convergence relative to the projection set $\bar{\Theta}$ and keep local constants explicit, in line with the misspecification-sensitive analyses in [5, 6].
(3) *Distributional metrics with sharp complexity control:* we evaluate error in a general IPM [8, 9] and express finite-sample operator deviations through intrinsic empirical process complexity [10, 11].

These choices yield a framework in which EM guarantees are orbit-invariant, misspecification-aware, and metrically meaningful at the level of distributions.

3. Setup and notation

3.1. **Observed and latent spaces, dominating measures.** Let $(\mathcal{X}, \mathscr{X})$ be a measurable space for observations and $(\mathcal{Z}, \mathscr{Z})$ a measurable space for latent variables. Fix σ-finite dominating measures μ on $(\mathcal{X}, \mathscr{X})$ and ν on $(\mathcal{Z}, \mathscr{Z})$. All densities below are understood with respect to μ on $\mathcal{X}$ and $\mu \otimes \nu$ on $\mathcal{X} \times \mathcal{Z}$.

3.2. **Latent-variable model family.** Let Θ be a parameter space (a subset of $\mathbb{R}^d$ or a smooth manifold; we only assume enough structure to state local differentiability and to define a group action in Section 4). For each $\theta \in \Theta$, let $p_\theta(x, z)$ be a nonnegative measurable function on $\mathcal{X} \times \mathcal{Z}$ satisfying

$$\int_{\mathcal{X}\times\mathcal{Z}} p_\theta(x, z)\,(\mu \otimes \nu)(d(x, z)) = 1.$$

Define the observed (marginal) density

$$p_\theta(x) := \int_{\mathcal{Z}} p_\theta(x, z)\,\nu(dz), \tag{11}$$

and the corresponding probability measure P_θ on $(\mathcal{X}, \mathscr{X})$ by

$$P_\theta(A) := \int_A p_\theta(x)\,\mu(dx), \qquad A \in \mathscr{X}.$$

Whenever $p_\theta(x) > 0$, define the posterior (E-step) kernel

$$p_\theta(dz \mid x) := \frac{p_\theta(x, z)}{p_\theta(x)}\,\nu(dz). \tag{12}$$

3.3. **Data-generating distribution and misspecification.** We observe i.i.d. samples $X_1, \dots, X_n \sim P^*$, where P^* is an arbitrary probability measure on $(\mathcal{X}, \mathscr{X})$. No assumption is made that $P^* \in \{P_\theta : \theta \in \Theta\}$ (misspecification is allowed). We assume the following integrability condition to ensure the objective is well-defined.

Assumption 3.1 (Log-likelihood integrability)**.** There exists at least one $\theta_0 \in \Theta$ such that $\mathbb{E}_{P^*}|\log p_{\theta_0}(X)| < \infty$. Moreover, for all θ in the neighborhoods considered later, $\mathbb{E}_{P^*}|\log p_\theta(X)| < \infty$.

Define the population (cross-entropy) objective

$$\Phi(\theta) := \mathbb{E}_{X\sim P^*}\big[\log p_\theta(X)\big], \tag{13}$$

and the empirical objective

$$\Phi_n(\theta) := \frac{1}{n}\sum_{i=1}^{n} \log p_\theta(X_i). \tag{14}$$

The *KL projection set* (population maximizers) is

$$\bar{\Theta} := \arg\max_{\theta\in\Theta} \Phi(\theta). \tag{15}$$

By the identity $\mathrm{KL}(P^* \,\|\, P_\theta) = \mathbb{E}_{P^*}[\log p^*(X)] - \Phi(\theta)$, we equivalently have

$$\bar{\Theta} = \arg\min_{\theta\in\Theta} \mathrm{KL}(P^* \,\|\, P_\theta). \tag{16}$$

We do not assume $\bar{\Theta}$ is singleton.

3.4. **Integral probability metrics (IPMs).** Let $\mathcal{F}$ be a class of measurable functions $f : \mathcal{X} \to \mathbb{R}$ such that $\mathbb{E}_P|f| < \infty$ for all P under consideration (in particular P^* and the P_θ in neighborhoods of interest). Define the integral probability metric

$$d_{\mathcal{F}}(P, Q) := \sup_{f\in\mathcal{F}} \big|\mathbb{E}_P f - \mathbb{E}_Q f\big|. \tag{17}$$

No structure is imposed on $\mathcal{F}$ beyond measurability and integrability; all rates will therefore be expressed through intrinsic complexity quantities associated with $\mathcal{F}$ or with induced classes (Section 8).

Remark 3.2 (Examples). The definition (17) includes, as special cases: (i) total variation (bounded $\mathcal{F}$), (ii) bounded-Lipschitz metrics (Lipschitz and bounded $\mathcal{F}$), (iii) Wasserstein-type metrics (Lipschitz $\mathcal{F}$ on metric $\mathcal{X}$ under suitable moment conditions), and (iv) maximum mean discrepancy (RKHS unit balls). For background and empirical estimation theory, see [8, 9].

3.5. **Population EM operator.** Define the population EM surrogate (expected complete-data log-likelihood under the current posterior)

$$Q(\theta'; \theta) := \mathbb{E}_{X\sim P^*}\Big[\mathbb{E}_{Z\sim p_\theta(\cdot|X)} \log p_{\theta'}(X, Z)\Big], \tag{18}$$

whenever the expectation is well-defined. The (possibly set-valued) population EM map is

$$M(\theta) \in \arg\max_{\theta'\in\Theta} Q(\theta'; \theta). \tag{19}$$

We will use the term *exact EM* when (19) is computed exactly (global maximizer of the surrogate) and otherwise refer to inexact/approximate EM variants in Section 9.

3.6. **Sample EM operator.** Define the sample surrogate

$$\widehat{Q}_n(\theta';\theta) := \frac{1}{n}\sum_{i=1}^{n} \mathbb{E}_{Z\sim p_\theta(\cdot|X_i)} \log p_{\theta'}(X_i, Z), \tag{20}$$

and the (possibly set-valued) sample EM map

$$\widehat{M}_n(\theta) \in \arg\max_{\theta'\in\Theta} \widehat{Q}_n(\theta';\theta). \tag{21}$$

3.7. **Regularity conventions.** The main results are local, and we will explicitly state the regularity assumptions needed in each section. To avoid repetition, we adopt the following conventions.

- "Neighborhood" always means a subset of Θ (or of a local slice in Section 4) on which all involved expectations are finite and the operators are well-defined.
- Differentiability assumptions are imposed only when needed (e.g., for linearization in Section 6); otherwise we work with variational and contraction inequalities.
- When $\arg\max$ is set-valued, fixed-point and convergence statements are interpreted in the standard set-valued sense (e.g., distance-to-set formulations in Section 7).

3.8. **Preview: symmetry and quotient formulation.** The group action and quotient geometry induced by observational invariance (3) are introduced in the next section. In particular, we will define the quotient parameter space $\Theta/G = \Theta/G$, the orbit notation $[\theta]$, and the quotient IPM $\bar{d}_{\mathcal{F}}([\theta], [\theta']) = d_{\mathcal{F}}(P_\theta, P_{\theta'})$. All subsequent convergence statements will be formulated either in terms of $\bar{d}_{\mathcal{F}}$ or in terms of $\mathrm{dist}_{\mathcal{F}}(\cdot, \bar{\Theta}/G)$.

4. Quotient geometry and IPM metrics

This section isolates the minimal geometric structure forced by observational nonidentifiability and records the basic metric objects used throughout the paper. The key point is that under symmetry-induced nonidentifiability, the parameter θ is not the estimand; the estimand is the *orbit* $[\theta]$ in the quotient $\Theta/G := \Theta/G$. Accordingly, all discrepancies are defined either between orbits or between an orbit and an orbit-invariant target set. We work with an arbitrary IPM $d_{\mathcal{F}}$ on distributions (17) [8, 9] and transfer it canonically to the quotient.

4.1. Group actions and the induced equivalence relation.

Definition 4.1 (Measurable group action)**.** Let G be a group. A (left) action $G \curvearrowright \Theta$ is a map $(g, \theta) \mapsto g \cdot \theta$ from $G \times \Theta$ to Θ satisfying $e \cdot \theta = \theta$ and $(g_1 g_2) \cdot \theta = g_1 \cdot (g_2 \cdot \theta)$. When Θ is equipped with a σ-field $\mathscr{T}$, we say the action is *measurable* if the map $G \times \Theta \to \Theta$ is measurable with respect to the product σ-field.

Definition 4.2 (Orbit equivalence and quotient)**.** Define an equivalence relation $\sim$ on Θ by

$$\theta \sim \theta' \quad \Longleftrightarrow \quad \exists g \in G : \ \theta' = g \cdot \theta.$$

For $\theta \in \Theta$, the equivalence class (orbit) is

$$[\theta] := \{\theta' \in \Theta : \theta' \sim \theta\} = \{g \cdot \theta : \ g \in G\}.$$

The quotient set Θ/G is the set of all orbits, and the quotient map is $q : \Theta \to \Theta/G$, $q(\theta) = [\theta]$.

4.2. **Observed invariance and factorization through the quotient.** Let $\theta \mapsto P_\theta$ denote the model map from parameters to distributions on $(\mathcal{X}, \mathscr{X})$, as defined in Section 3.

Assumption 4.3 (Observed invariance)**.** For all $\theta \in \Theta$ and $g \in G$,

$$P_{g\cdot\theta} = P_\theta. \tag{22}$$

Remark 4.4 (Equivalent formulation). Assumption 4.3 is equivalent to the statement that the model map $\theta \mapsto P_\theta$ is constant on equivalence classes of $\sim$. Hence it factors through the quotient: there exists a well-defined map $\bar{P} : \Theta/G \to \mathcal{P}(\mathcal{X})$ such that $\bar{P}([\theta]) = P_\theta$.

4.3. **Quotient IPM: definition and line-by-line proof of well-definedness.** Recall the IPM $d_{\mathcal{F}}$ on distributions (17). We now define its quotient analogue.

Proposition 4.5 (Quotient IPM is well-defined and is a pseudometric)**.** *Assume Assumption 4.3. Define for $\xi, \xi' \in \Theta/G$,*

$$\bar{d}_{\mathcal{F}}(\xi, \xi') := d_{\mathcal{F}}\big(\bar{P}(\xi), \bar{P}(\xi')\big), \tag{23}$$

equivalently $\bar{d}_{\mathcal{F}}([\theta], [\theta']) := d_{\mathcal{F}}(P_\theta, P_{\theta'})$. Then:

(i) (23) is well-defined (independent of representatives).
(ii) $\bar{d}_{\mathcal{F}}$ is a pseudometric on Θ/G.
(iii) If $\mathcal{F}$ is determining *for the range $\{\bar{P}(\xi) : \xi \in \Theta/G\}$ (i.e. $d_{\mathcal{F}}(P, Q) = 0 \Rightarrow P = Q$ for P, Q in this range), then $\bar{d}_{\mathcal{F}}(\xi, \xi') = 0$ implies $\bar{P}(\xi) = \bar{P}(\xi')$; in particular, $\bar{d}_{\mathcal{F}}$ separates orbits up to equality of induced observed distributions.*

Proof. We give a fully explicit argument.

Step 1 (Representative-independence). Let $\xi, \xi' \in \Theta/G$ and pick arbitrary representatives $\theta \in q^{-1}(\xi)$ and $\theta' \in q^{-1}(\xi')$. If $\tilde{\theta} \in q^{-1}(\xi)$ is another representative, then by Definition 4.2 there exists $g \in G$ with $\tilde{\theta} = g \cdot \theta$. By Assumption 4.3,

$$P_{\tilde{\theta}} = P_{g\cdot\theta} = P_\theta.$$

Thus for any $\theta' \in q^{-1}(\xi')$,

$$d_{\mathcal{F}}(P_{\tilde{\theta}}, P_{\theta'}) = d_{\mathcal{F}}(P_\theta, P_{\theta'}).$$

The same argument applies if θ' is replaced by another representative of ξ'. Therefore $d_{\mathcal{F}}(P_\theta, P_{\theta'})$ depends only on ξ and ξ', proving well-definedness.

Step 2 (Pseudometric properties). Nonnegativity and symmetry hold because $d_{\mathcal{F}}$ is nonnegative and symmetric by definition (17). Also $\bar{d}_{\mathcal{F}}(\xi, \xi) = d_{\mathcal{F}}(\bar{P}(\xi), \bar{P}(\xi)) = 0$. For the triangle inequality, let $\xi_1, \xi_2, \xi_3 \in \Theta/G$. Then, using the triangle inequality of $d_{\mathcal{F}}$ on distributions,

$$\bar{d}_{\mathcal{F}}(\xi_1, \xi_3) = d_{\mathcal{F}}(\bar{P}(\xi_1), \bar{P}(\xi_3)) \le d_{\mathcal{F}}(\bar{P}(\xi_1), \bar{P}(\xi_2)) + d_{\mathcal{F}}(\bar{P}(\xi_2), \bar{P}(\xi_3)) = \bar{d}_{\mathcal{F}}(\xi_1, \xi_2) + \bar{d}_{\mathcal{F}}(\xi_2, \xi_3).$$

Hence $\bar{d}_{\mathcal{F}}$ is a pseudometric.

Step 3 (Separation on the model range). If $\mathcal{F}$ is determining for the range of $\bar{P}$, then $\bar{d}_{\mathcal{F}}(\xi, \xi') = 0$ implies $d_{\mathcal{F}}(\bar{P}(\xi), \bar{P}(\xi')) = 0$, hence $\bar{P}(\xi) = \bar{P}(\xi')$ by the determining property. □

Remark 4.6 (When is $d_{\mathcal{F}}$ determining?). Whether $d_{\mathcal{F}}(P, Q) = 0$ implies $P = Q$ depends on $\mathcal{F}$. For example, if $\mathcal{F}$ is the unit ball of bounded measurable functions, then $d_{\mathcal{F}}$ is total variation and is determining. If $\mathcal{F}$ is too small, distinct distributions may be indistinguishable under $d_{\mathcal{F}}$ [8, 9]. Our results remain valid in the pseudometric setting; where orbit-identification is required, we state additional assumptions explicitly (Section 4.7).

4.4. **Distance to orbit-invariant sets: definition and proof.** The misspecified target $\bar{\Theta}$ is set-valued in general; thus distance-to-set formulations are necessary.

Definition 4.7 (Orbit-invariant set)**.** A set $A \subseteq \Theta$ is *G-invariant* if $g \cdot A = A$ for all $g \in G$. Equivalently, A is a union of orbits: if $\theta \in A$ then $[\theta] \subseteq A$.

Proposition 4.8 (Distance-to-set is well-defined on the quotient)**.** *Assume Assumption 4.3. Let $A \subseteq \Theta$ be G-invariant. Define for $\xi \in \Theta/G$,*

$$\operatorname{dist}_{\mathcal{F}}(\xi, A/G) := \inf_{\alpha \in A} \bar{d}_{\mathcal{F}}(\xi, [\alpha]) = \inf_{\alpha \in A} d_{\mathcal{F}}(\bar{P}(\xi), P_\alpha). \tag{24}$$

Then (24) is well-defined (independent of all representatives), and the function $\xi \mapsto \operatorname{dist}_{\mathcal{F}}(\xi, A/G)$ depends only on the orbit class ξ.

Proof. Fix $\xi \in \Theta/G$ and choose any representative $\theta \in q^{-1}(\xi)$.

Step 1 (Representative-independence in the first argument). If $\tilde{\theta} \in q^{-1}(\xi)$ is another representative, then $\tilde{\theta} = g \cdot \theta$ for some $g \in G$ and $P_{\tilde{\theta}} = P_\theta$ by Assumption 4.3. Hence for any $\alpha \in A$,

$$d_{\mathcal{F}}(P_{\tilde{\theta}}, P_\alpha) = d_{\mathcal{F}}(P_\theta, P_\alpha).$$

Taking the infimum over $\alpha \in A$ shows the value of (24) does not depend on the chosen representative of ξ.

Step 2 (Representative-independence inside A). If $\alpha \in A$ and $\tilde{\alpha} \in [\alpha]$, then $P_{\tilde{\alpha}} = P_\alpha$ by Assumption 4.3, so $d_{\mathcal{F}}(P_\theta, P_{\tilde{\alpha}}) = d_{\mathcal{F}}(P_\theta, P_\alpha)$. Thus taking $\inf_{\alpha \in A}$ is the same as taking $\inf_{[\alpha] \in A/G}$, and the result depends only on A/G. □

4.5. **The KL projection set is orbit-invariant: detailed proof.**

Lemma 4.9 (Orbit-invariance of the KL projection set)**.** *Assume Assumption 4.3 and Assumption 3.1. Let $\bar{\Theta}$ be defined by (15). Then $\bar{\Theta}$ is G-invariant. Consequently, $\bar{\Theta}/G$ is well-defined and the quantity $\operatorname{dist}_{\mathcal{F}}([\theta], \bar{\Theta}/G)$ is well-defined via Proposition 4.8.*

Proof. Fix $\bar{\theta} \in \bar{\Theta}$ and $g \in G$.

Step 1 (Invariance of the observed density a.e.). Assumption 4.3 states $P_{g\cdot\theta} = P_\theta$ as measures on $(\mathcal{X}, \mathscr{X})$. Since both are dominated by μ, their Radon–Nikodym derivatives satisfy

$$p_{g\cdot\theta}(x) = p_\theta(x) \quad \text{for } \mu\text{-a.e. } x.$$

Step 2 (Invariance of the population objective). By (13) and Step 1,

$$\Phi(g \cdot \theta) = \mathbb{E}_{P^*}\big[\log p_{g\cdot\theta}(X)\big] = \mathbb{E}_{P^*}\big[\log p_\theta(X)\big] = \Phi(\theta),$$

where the expectations are finite by Assumption 3.1 in the neighborhoods of interest.

Step 3 (Invariance of maximizers). Because $\bar{\theta} \in \bar{\Theta}$, for all $\theta \in \Theta$,

$$\Phi(\bar{\theta}) \geq \Phi(\theta).$$

By Step 2, $\Phi(g \cdot \bar{\theta}) = \Phi(\bar{\theta})$, hence for all $\theta \in \Theta$,

$$\Phi(g \cdot \bar{\theta}) = \Phi(\bar{\theta}) \geq \Phi(\theta),$$

which implies $g \cdot \bar{\theta} \in \bar{\Theta}$ by definition (15). □

4.6. **Local sections (slices): measurable representatives.** To state sharp local rates (Section 6) we work in identifiable coordinates. We therefore assume the existence of a local cross-section.

Assumption 4.10 (Local measurable section)**.** There exist a set $\mathcal{U} \subseteq \Theta/G$ and a set $S \subseteq \Theta$ such that:

(i) (*Uniqueness*) For every $\xi \in \mathcal{U}$, the intersection $q^{-1}(\xi) \cap S$ consists of exactly one point.
(ii) (*Measurability*) The induced map $\pi : \mathcal{U} \to S$, defined by selecting this unique point, is measurable.

For $\xi \in \mathcal{U}$, write $\theta_S := \pi(\xi)$.

Lemma 4.11 (Slice representative map: uniqueness and consistency)**.** *Under Assumption 4.10, for every θ with $[\theta] \in \mathcal{U}$ there exists a unique $\theta_S \in S$ such that $[\theta_S] = [\theta]$, namely $\theta_S = \pi([\theta])$. Moreover, if $\theta' \in \Theta$ satisfies $[\theta'] = [\theta]$ and $[\theta] \in \mathcal{U}$, then $\pi([\theta']) = \pi([\theta])$.*

Proof. Fix θ with $[\theta] \in \mathcal{U}$.

Step 1 (Existence and uniqueness). By Assumption 4.10(i), the set $q^{-1}([\theta]) \cap S$ contains exactly one element. Define θ_S to be this element. Then $\theta_S \in S$ and $q(\theta_S) = [\theta]$, i.e. $[\theta_S] = [\theta]$.

Step 2 (Consistency across representatives). If $[\theta'] = [\theta]$, then $[\theta']$ and $[\theta]$ are identical elements of Θ/G. Since π is a function on $\mathcal{U}$, $\pi([\theta']) = \pi([\theta])$. □

Remark 4.12 (When do local sections exist?). Assumption 4.10 is purely local and is the minimal structure needed to define identifiable coordinates. In many smooth settings (e.g. a Lie group acting properly on a smooth manifold), local slices exist by classical slice theorems; see, e.g., the Palais slice theorem (not required for our arguments, which are conditional on the existence of a section). We keep the section assumption explicit to avoid committing to topological/differential hypotheses that are unnecessary for the core results.

4.7. **Metric transfer: orbit distances and moduli.** To transfer contraction bounds from slice coordinates to distributional discrepancy, we use a local modulus controlling $\theta \mapsto P_\theta$ into the IPM. The following formulation is designed to be both minimal and explicit.

Definition 4.13 (Orbit distance in parameter space)**.** Assume $\Theta \subseteq \mathbb{R}^d$ is equipped with a norm $\|\cdot\|$. Define an orbit pseudometric on Θ/G by

$$d_\Theta([\theta], [\theta']) := \inf_{g \in G} \|\theta - g \cdot \theta'\|.$$

Assumption 4.14 (Two-sided modulus comparison on a local basin)**.** There exist a subset $\mathcal{U} \subseteq \Theta/G$ and nondecreasing functions $\underline{\omega}, \overline{\omega} : [0, \infty) \to [0, \infty)$ with $\underline{\omega}(0) = \overline{\omega}(0) = 0$ such that for all $\xi, \xi' \in \mathcal{U}$,

$$\underline{\omega}\big(d_\Theta(\xi, \xi')\big) \;\le\; \bar{d}_{\mathcal{F}}(\xi, \xi') \;\le\; \overline{\omega}\big(d_\Theta(\xi, \xi')\big). \tag{25}$$

Assume additionally that $\underline{\omega}$ is strictly increasing on $[0, r_0]$ for some $r_0 > 0$ (so it admits a well-defined inverse on its range).

Theorem 4.15 (Local topological equivalence and quantitative contraction transfer)**.** *Assume Assumption 4.14. Then:*

*(i) (*Local topological equivalence*) The pseudometrics d_Θ and $\bar{d}_{\mathcal{F}}$ generate the same local topology on $\mathcal{U}$. Concretely, for any $\xi \in \mathcal{U}$ and any $r > 0$ such that $r \le \underline{\omega}(r_0)$, the inclusions*

$$B_{d_\Theta}\big(\xi, \overline{\omega}^{-1}(r)\big) \subseteq B_{\bar{d}_{\mathcal{F}}}(\xi, r) \subseteq B_{d_\Theta}\big(\xi, \underline{\omega}^{-1}(r)\big)$$

hold (for any chosen right-inverse $\overline{\omega}^{-1}$ and the inverse $\underline{\omega}^{-1}$ on $[0, \underline{\omega}(r_0)]$).

*(ii) (*Contraction transfer*) Let $\Psi : \mathcal{U} \to \mathcal{U}$ satisfy*

$$d_\Theta(\Psi(\xi), \Psi(\xi')) \le \gamma\, d_\Theta(\xi, \xi') \qquad \forall \xi, \xi' \in \mathcal{U}$$

for some $\gamma \in (0,1)$. Then for all $\xi, \xi' \in \mathcal{U}$ with $\bar{d}_\mathcal{F}(\xi, \xi') \le \underline{\omega}(r_0)$,

$$\bar{d}_\mathcal{F}(\Psi(\xi), \Psi(\xi')) \le \overline{\omega}\Big(\gamma\, \underline{\omega}^{-1}(\bar{d}_\mathcal{F}(\xi, \xi'))\Big). \tag{26}$$

In particular, if $\underline{\omega}(r) \ge cr$ and $\overline{\omega}(r) \le Lr$ for $r \in [0, r_0]$, then

$$\bar{d}_\mathcal{F}(\Psi(\xi), \Psi(\xi')) \le \frac{L}{c}\, \gamma\, \bar{d}_\mathcal{F}(\xi, \xi') \qquad \text{whenever } \bar{d}_\mathcal{F}(\xi, \xi') \le cr_0. \tag{27}$$

Proof. **(i)** Fix $\xi \in \mathcal{U}$ and $r \le \underline{\omega}(r_0)$.

Step 1 (Inner inclusion). If $d_\Theta(\xi, \xi') < \overline{\omega}^{-1}(r)$ then by the upper bound in (25),

$$\bar{d}_\mathcal{F}(\xi, \xi') \le \overline{\omega}(d_\Theta(\xi, \xi')) < \overline{\omega}(\overline{\omega}^{-1}(r)) \le r,$$

so $\xi' \in B_{\bar{d}_\mathcal{F}}(\xi, r)$.

Step 2 (Outer inclusion). If $\bar{d}_\mathcal{F}(\xi, \xi') < r$, then by the lower bound in (25),

$$\underline{\omega}(d_\Theta(\xi, \xi')) \le \bar{d}_\mathcal{F}(\xi, \xi') < r.$$

Since $r \le \underline{\omega}(r_0)$ and $\underline{\omega}$ is strictly increasing on $[0, r_0]$, the inverse is well-defined and yields $d_\Theta(\xi, \xi') < \underline{\omega}^{-1}(r)$, so $\xi' \in B_{d_\Theta}(\xi, \underline{\omega}^{-1}(r))$.

(ii) Fix $\xi, \xi' \in \mathcal{U}$ with $\bar{d}_\mathcal{F}(\xi, \xi') \le \underline{\omega}(r_0)$.

Step 1 (Convert $\bar{d}_\mathcal{F}$ to d_Θ). From the lower bound in (25),

$$\underline{\omega}(d_\Theta(\xi, \xi')) \le \bar{d}_\mathcal{F}(\xi, \xi').$$

Applying $\underline{\omega}^{-1}$ (valid by the assumed range restriction) gives

$$d_\Theta(\xi, \xi') \le \underline{\omega}^{-1}(\bar{d}_\mathcal{F}(\xi, \xi')).$$

Step 2 (Apply contraction in d_Θ). By the hypothesis on Ψ,

$$d_\Theta(\Psi(\xi), \Psi(\xi')) \le \gamma\, d_\Theta(\xi, \xi') \le \gamma\, \underline{\omega}^{-1}(\bar{d}_\mathcal{F}(\xi, \xi')).$$

Step 3 (Convert back to $\bar{d}_\mathcal{F}$). Using the upper bound in (25) for $(\Psi(\xi), \Psi(\xi'))$ yields (26). If, moreover, $\underline{\omega}(r) \ge cr$ and $\overline{\omega}(r) \le Lr$ on $[0, r_0]$, then for $\bar{d}_\mathcal{F}(\xi, \xi') \le cr_0$ we have $\underline{\omega}^{-1}(t) \le t/c$ and hence

$$\bar{d}_\mathcal{F}(\Psi(\xi), \Psi(\xi')) \le \overline{\omega}\big(\gamma\, \underline{\omega}^{-1}(\bar{d}_\mathcal{F}(\xi, \xi'))\big) \le L\gamma\, \frac{\bar{d}_\mathcal{F}(\xi, \xi')}{c},$$

which is (27). □

Remark 4.16 (Minimality of the modulus framework). Assumption 4.14 isolates the exact regularity required to translate between orbit distances in parameter space and distributional discrepancies in a general IPM. A one-sided modulus (upper bound only) suffices for upper bounds on $\bar{d}_\mathcal{F}$ along iterates. A two-sided modulus is needed only when one wishes to infer orbit-identification or to transfer contraction *backward* from $\bar{d}_\mathcal{F}$ to d_Θ. We keep these requirements explicit to avoid conflating statistical identifiability assumptions with algorithmic stability assumptions.

5. EM on the quotient: operators, equivariance, and ascent

This section formalizes EM as an operator on the quotient $\Theta/G = \Theta/G$ and proves the ascent properties (population and sample) in a form that is compatible with misspecification and set-valued maximizers. While ascent is classical [4, 12], we present a formulation that (i) is explicitly orbit-invariant, (ii) cleanly separates the surrogate improvement from the posterior-KL term, and (iii) supports subsequent contraction arguments on slices.

5.1. Well-posedness of the EM surrogate.

Recall the population surrogate Q and the population objective Φ from (18) and (13). To avoid repeating integrability conditions in every statement, we impose the following local standing assumption.

Assumption 5.1 (Local well-posedness of Q and conditional log terms)**.** Let $B \subseteq \Theta$ be a neighborhood (later, typically a slice-basin). Assume:

(i) For all $\theta, \theta' \in B$,

$$\mathbb{E}_{X\sim P^*}\Big|\mathbb{E}_{Z\sim p_\theta(\cdot|X)} \log p_{\theta'}(X, Z)\Big| < \infty,$$

so $Q(\theta'; \theta)$ is finite.

(ii) For all $\theta, \theta' \in B$,

$$\mathbb{E}_{X\sim P^*}\Big|\mathbb{E}_{Z\sim p_\theta(\cdot|X)} \log p_{\theta'}(Z \mid X)\Big| < \infty,$$

where $p_{\theta'}(z \mid x)$ is defined by (12) whenever $p_{\theta'}(x) > 0$.

Analogous assumptions are imposed on the empirical neighborhood for $\widehat{Q}_n$ (with P^* replaced by P_n).

Remark 5.2. Assumption 5.1 is purely local and is standard in EM analyses; it ensures that the decompositions used below are legitimate (finite expectations and permissible conditioning). All subsequent results are stated on neighborhoods where Assumption 5.1 holds.

5.2. Equivariance at the complete-data level.

Observed invariance (22) alone is sufficient to define quotient IPM errors (Section 4), but it does *not* by itself guarantee that the EM surrogate $Q(\cdot;\cdot)$ transforms equivariantly under the group action. To ensure that EM descends to the quotient as an operator, we impose a standard compatibility condition that captures the usual symmetries in latent-variable models (e.g. label permutations).

Assumption 5.3 (Complete-data equivariance via a latent transport)**.** There exists a family of measurable bijections $\{\tau_g : \mathcal{Z} \to \mathcal{Z}\}_{g\in G}$ such that:

(i) (*Group property*) $\tau_e = \mathrm{id}$ and $\tau_{g_1 g_2} = \tau_{g_1} \circ \tau_{g_2}$.

(ii) (*Reference-measure invariance*) $\nu \circ \tau_g^{-1} = \nu$ for all $g \in G$.

(iii) (*Complete-data equivariance*) For all $g \in G$ and $\theta \in \Theta$,

$$p_{g\cdot\theta}(x, z) = p_\theta\big(x, \tau_g^{-1}(z)\big) \quad \text{for } (\mu \otimes \nu)\text{-a.e. } (x, z). \tag{28}$$

Lemma 5.4 (Observed invariance and posterior transport)**.** *Assume Assumption 5.3. Then:*

(i) (Observed invariance) *$P_{g\cdot\theta} = P_\theta$ for all g, θ.*

(ii) (Posterior transport) *For μ-a.e. x with $p_\theta(x) > 0$,*

$$p_{g\cdot\theta}(\,\cdot \mid x) = (\,\tau_g\,)_{\#}\, p_\theta(\,\cdot \mid x), \tag{29}$$

i.e. for any measurable $A \subseteq \mathcal{Z}$, $p_{g\cdot\theta}(A \mid x) = p_\theta(\tau_g^{-1}(A) \mid x)$.

Proof. **(i)** Fix $g \in G$ and $\theta \in \Theta$. By definition (11) and (28),

$$p_{g\cdot\theta}(x) = \int_{\mathcal{Z}} p_{g\cdot\theta}(x,z)\,\nu(dz) = \int_{\mathcal{Z}} p_\theta(x,\tau_g^{-1}(z))\,\nu(dz).$$

By the change-of-variables formula for the measure-preserving bijection τ_g (Assumption 5.3(ii)),

$$\int_{\mathcal{Z}} p_\theta(x,\tau_g^{-1}(z))\,\nu(dz) = \int_{\mathcal{Z}} p_\theta(x,z)\,\nu(dz) = p_\theta(x),$$

hence $p_{g\cdot\theta}(x) = p_\theta(x)$ for μ-a.e. x and thus $P_{g\cdot\theta} = P_\theta$.

(ii) Fix x with $p_\theta(x) > 0$. For any measurable $A \subseteq \mathcal{Z}$, using (12) and (28),

$$p_{g\cdot\theta}(A \mid x) = \int_A \frac{p_{g\cdot\theta}(x,z)}{p_{g\cdot\theta}(x)}\,\nu(dz) = \int_A \frac{p_\theta(x,\tau_g^{-1}(z))}{p_\theta(x)}\,\nu(dz).$$

Apply the same measure-preserving change of variables $z = \tau_g(z')$:

$$\int_A \frac{p_\theta(x,\tau_g^{-1}(z))}{p_\theta(x)}\,\nu(dz) = \int_{\tau_g^{-1}(A)} \frac{p_\theta(x,z')}{p_\theta(x)}\,\nu(dz') = p_\theta(\tau_g^{-1}(A) \mid x),$$

which is exactly (29). □

5.3. **Equivariance of the EM surrogate and induced quotient maps.** We now show that EM descends to an orbit map (generally set-valued) on Θ/G.

Lemma 5.5 (Equivariance of the EM surrogate)**.** *Assume Assumption 5.3 and Assumption 5.1 on a neighborhood B. Then for all $g \in G$ and all $\theta, \theta' \in B$,*

$$Q(g\cdot\theta';\, g\cdot\theta) = Q(\theta';\theta). \tag{30}$$

The same statement holds for the sample surrogate $\widehat{Q}_n$ (with P^* replaced by P_n).

Proof. Fix $g \in G$ and $\theta, \theta' \in B$. By definition (18) and Lemma 5.4(ii),

$$Q(g\cdot\theta';\, g\cdot\theta) = \mathbb{E}_{X\sim P^*}\Big[\mathbb{E}_{Z\sim p_{g\cdot\theta}(\cdot|X)} \log p_{g\cdot\theta'}(X,Z)\Big].$$

Condition on $X = x$ and apply the posterior transport (29): if $Z \sim p_{g\cdot\theta}(\cdot \mid x)$ then $Z = \tau_g(Z')$ for $Z' \sim p_\theta(\cdot \mid x)$. Therefore

$$\mathbb{E}_{Z\sim p_{g\cdot\theta}(\cdot|x)} \log p_{g\cdot\theta'}(x,Z) = \mathbb{E}_{Z'\sim p_\theta(\cdot|x)} \log p_{g\cdot\theta'}(x,\tau_g(Z')).$$

By complete-data equivariance (28) applied to θ',

$$p_{g\cdot\theta'}(x,\tau_g(z')) = p_{\theta'}(x,z') \quad \text{for } (\mu\otimes\nu)\text{-a.e. } (x,z').$$

Hence, for μ-a.e. x,

$$\mathbb{E}_{Z'\sim p_\theta(\cdot|x)} \log p_{g\cdot\theta'}(x,\tau_g(Z')) = \mathbb{E}_{Z'\sim p_\theta(\cdot|x)} \log p_{\theta'}(x,Z').$$

Taking expectation over $X \sim P^*$ yields (30). The sample version is identical with P^* replaced by P_n. □

Theorem 5.6 (Equivariance of EM and induced quotient operator)**.** *Assume Assumption 5.3 and Assumption 5.1 on a neighborhood B. Let the population EM correspondence be*

$$M(\theta) := \arg\max_{\theta'\in B} Q(\theta';\theta).$$

Then for all $g \in G$ and $\theta \in B$,

$$M(g\cdot\theta) = g\cdot M(\theta) := \{g\cdot\eta :\ \eta \in M(\theta)\}. \tag{31}$$

Consequently, the quotient EM correspondence

$$\bar{M}: \; q(B) \rightrightarrows \Theta/G, \qquad \bar{M}([\theta]) := \{\, [\eta] : \; \eta \in M(\theta) \,\},$$

is well-defined (independent of the representative θ of the orbit). The same conclusions hold for the sample EM correspondence $\widehat{M}_n$ defined via $\widehat{Q}_n$.

Proof. We proceed in explicit steps.

Step 1 (Argmax invariance under equivariant reparametrization). Fix $g \in G$ and $\theta \in B$. By Lemma 5.5, for every $\eta \in B$,

$$Q(\eta; \theta) = Q(g \cdot \eta;\, g \cdot \theta).$$

Thus the ordering of surrogate values over $\eta \in B$ is preserved under $\eta \mapsto g \cdot \eta$ when θ is replaced by $g \cdot \theta$.

Step 2 (Show $g \cdot M(\theta) \subseteq M(g \cdot \theta)$). Let $\eta^\star \in M(\theta)$. Then for all $\eta \in B$,

$$Q(\eta^\star; \theta) \geq Q(\eta; \theta).$$

Using Lemma 5.5 on both sides,

$$Q(g \cdot \eta^\star;\, g \cdot \theta) \geq Q(g \cdot \eta;\, g \cdot \theta) \quad \forall \eta \in B.$$

Since g acts bijectively on B (as a restriction of a group action), $\{g \cdot \eta : \eta \in B\} = B$. Therefore $g \cdot \eta^\star$ maximizes $\eta' \mapsto Q(\eta';\, g \cdot \theta)$ over B, i.e. $g \cdot \eta^\star \in M(g \cdot \theta)$. This proves $g \cdot M(\theta) \subseteq M(g \cdot \theta)$.

Step 3 (Reverse inclusion). Apply Step 2 with g^{-1} in place of g:

$$g^{-1} \cdot M(g \cdot \theta) \subseteq M(\theta).$$

Applying g to both sides yields $M(g \cdot \theta) \subseteq g \cdot M(\theta)$.

Step 4 (Conclude (31) and well-definedness on the quotient). Steps 2–3 give equality (31). Now let $\theta, \tilde{\theta} \in B$ satisfy $[\theta] = \left[\tilde{\theta}\right]$, so $\tilde{\theta} = g \cdot \theta$ for some g. Then (31) gives $M(\tilde{\theta}) = g \cdot M(\theta)$, so the set of orbits $\{[\eta] : \eta \in M(\tilde{\theta})\}$ equals $\{[\eta] : \eta \in M(\theta)\}$. Hence $\bar{M}$ is independent of the representative and is well-defined. The sample statement is identical with Q replaced by $\widehat{Q}_n$. □

Remark 5.7 (Why set-valued on the quotient?). Even when $Q(\cdot; \theta)$ admits multiple maximizers, Theorem 5.6 ensures that the maximizer set is transported by the group action. Thus the induced quotient map is naturally a correspondence. Later, when proving contraction and sharp linear rates, we work on a slice where a single-valued representative map can be imposed.

5.4. **Ascent and the posterior-KL decomposition (population).** We now present the canonical EM ascent inequality in a form that is especially convenient under misspecification: it expresses the improvement in Φ as the surrogate improvement *minus* a nonnegative posterior KL term. This is the standard EM mechanism [4, 12], stated here with explicit conditioning and integrability.

Lemma 5.8 (Exact identity underlying EM ascent)**.** *Assume Assumption 5.1 on a neighborhood B. Fix $\theta, \theta' \in B$. Then*

$$\Phi(\theta') - \Phi(\theta) = \Big(Q(\theta'; \theta) - Q(\theta; \theta)\Big) - \mathbb{E}_{X \sim P^*}\Big[\mathrm{KL}\big(p_\theta(\cdot \mid X) \,\|\, p_{\theta'}(\cdot \mid X)\big)\Big]. \tag{32}$$

In particular,

$$\Phi(\theta') - \Phi(\theta) \;\geq\; Q(\theta'; \theta) - Q(\theta; \theta). \tag{33}$$

Proof. Fix $\theta, \theta' \in B$.

Step 1 (A conditional decomposition of $\log p_{\theta'}(x)$). For μ-a.e. x with $p_{\theta'}(x) > 0$ and $p_\theta(x) > 0$, Bayes' rule gives

$$\log p_{\theta'}(x) = \log p_{\theta'}(x, z) - \log p_{\theta'}(z \mid x) \quad \text{for } \nu\text{-a.e. } z.$$

Taking expectation in z under the conditional law $p_\theta(\cdot \mid x)$ yields

$$\log p_{\theta'}(x) = \mathbb{E}_{Z \sim p_\theta(\cdot|x)}\big[\log p_{\theta'}(x, Z)\big] - \mathbb{E}_{Z \sim p_\theta(\cdot|x)}\big[\log p_{\theta'}(Z \mid x)\big], \tag{34}$$

since $\log p_{\theta'}(x)$ is constant in z. Assumption 5.1 ensures both conditional expectations are finite and hence (34) is legitimate.

Step 2 (Take expectation over $X \sim P^*$). Taking $\mathbb{E}_{X \sim P^*}$ of (34) and using the definition of Q gives

$$\Phi(\theta') = Q(\theta'; \theta) - \mathbb{E}_{X \sim P^*}\Big[\mathbb{E}_{Z \sim p_\theta(\cdot|X)} \log p_{\theta'}(Z \mid X)\Big]. \tag{35}$$

Setting $\theta' = \theta$ yields similarly

$$\Phi(\theta) = Q(\theta; \theta) - \mathbb{E}_{X \sim P^*}\Big[\mathbb{E}_{Z \sim p_\theta(\cdot|X)} \log p_{\theta}(Z \mid X)\Big]. \tag{36}$$

Step 3 (Subtract and identify the posterior KL term). Subtract (36) from (35):

$$\Phi(\theta') - \Phi(\theta) = \big(Q(\theta'; \theta) - Q(\theta; \theta)\big) - \mathbb{E}_{X \sim P^*}\Big[\mathbb{E}_{Z \sim p_\theta(\cdot|X)} \log \frac{p_{\theta'}(Z \mid X)}{p_\theta(Z \mid X)}\Big].$$

By the definition of conditional KL divergence,

$$\mathrm{KL}\big(p_\theta(\cdot \mid X) \,\|\, p_{\theta'}(\cdot \mid X)\big) = \mathbb{E}_{Z \sim p_\theta(\cdot|X)} \log \frac{p_\theta(Z \mid X)}{p_{\theta'}(Z \mid X)}.$$

Thus the last term equals $-\mathbb{E}_{P^*}[\mathrm{KL}(\, p_\theta(\cdot \mid X) \,\|\, p_{\theta'}(\cdot \mid X)\,)]$, which gives (32).

Step 4 (Conclude the ascent inequality). Since KL divergence is nonnegative, the second term on the right-hand side of (32) is ≤ 0, yielding (33). □

Lemma 5.9 (Population EM identity (fully expanded))**.** *Assume Assumption 5.1 on a neighborhood B. In particular, assume the following hold for every $\theta \in B$:*

(i) ***Dominated model and measurability:*** *there exist σ-finite measures μ on $(\mathsf{X}, \mathcal{F})$ and ν on $(\mathsf{Z}, \mathcal{G})$ such that $P_\theta^{XZ} \ll \mu \otimes \nu$ with a jointly measurable density $(x, z) \mapsto p_\theta(x, z)$. Define the marginal density $p_\theta(x) := \int p_\theta(x, z)\, \nu(dz)$, assumed $\mathcal{F}$-measurable.*

(ii) ***Posterior existence (a version):*** *for P^*-a.e. x, if $p_\theta(x) > 0$ then the conditional density*

$$p_\theta(z \mid x) := \frac{p_\theta(x, z)}{p_\theta(x)} \quad \textit{(as a density w.r.t. } \nu\textit{)}$$

is well-defined ν-a.e. in z and the mapping $(x, z) \mapsto p_\theta(z \mid x)$ is jointly measurable on $\{(x, z) : p_\theta(x) > 0\}$.

(iii) ***Support compatibility on B:*** *for all $\theta, \theta' \in B$,*

$$p_\theta(x) > 0 \text{ and } x \text{ in the } P^*\text{-support} \quad \Longrightarrow \quad p_{\theta'}(x) > 0 \quad (P^*\text{-a.s.}),$$

so that all log terms below are P^*-a.s. finite (no $\log 0$ on sets of positive P^*-mass).

(iv) ***Integrability/Fubini:*** *for all $\theta, \theta' \in B$,*

$$\mathbb{E}_{P^*}\mathbb{E}_{Z \sim p_\theta(\cdot|X)}\big[\,|\log p_{\theta'}(X, Z)| + |\log p_\theta(X, Z)|\,\big] < \infty, \qquad \mathbb{E}_{P^*}\big[\,|\log p_{\theta'}(X)| + |\log p_\theta(X)|\,\big] < \infty.$$

Then for any $\theta, \theta' \in B$,

$$\Phi(\theta') - \Phi(\theta) = \big(Q(\theta';\theta) - Q(\theta;\theta)\big) + \mathbb{E}_{P^*}\Big[\mathrm{KL}\big(p_\theta(\cdot \mid X) \,\|\, p_{\theta'}(\cdot \mid X)\big)\Big]. \tag{37}$$

In particular,

$$\Phi(\theta') - \Phi(\theta) \;\geq\; Q(\theta';\theta) - Q(\theta;\theta),$$

since the KL term is nonnegative.

Proof. Fix $\theta, \theta' \in B$.

Step 0: Define the full-measure set on which all objects are finite and defined. By the support-compatibility and posterior-existence parts of Assumption 5.1, there exists a measurable set $A \in \mathcal{F}$ with $P^*(A) = 1$ such that for every $x \in A$:

- $p_\theta(x) > 0$ and $p_{\theta'}(x) > 0$;
- $z \mapsto p_\theta(z \mid x)$ and $z \mapsto p_{\theta'}(z \mid x)$ are well-defined ν-a.e.;
- all logarithms below are finite $p_\theta(\cdot \mid x)$-a.s. (in particular, the Radon–Nikodym ratio $p_\theta(\cdot \mid x)/p_{\theta'}(\cdot \mid x)$ is defined $p_\theta(\cdot \mid x)$-a.s.).

Step 1: Expand the conditional KL at a fixed $x \in A$. Fix $x \in A$. By definition of conditional KL (with respect to the base measure ν),

$$\begin{aligned}\mathrm{KL}\big(p_\theta(\cdot \mid x) \,\|\, p_{\theta'}(\cdot \mid x)\big) &:= \int \log\Big(\frac{p_\theta(z \mid x)}{p_{\theta'}(z \mid x)}\Big)\, p_\theta(z \mid x)\, \nu(dz) \\ &= \mathbb{E}_{Z \sim p_\theta(\cdot \mid x)}\Big[\log p_\theta(Z \mid x) - \log p_{\theta'}(Z \mid x)\Big].\end{aligned} \tag{38}$$

Step 2: Substitute Bayes' rule into the integrand. For $x \in A$ and ν-a.e. z,

$$\log p_\theta(z \mid x) = \log p_\theta(x, z) - \log p_\theta(x), \qquad \log p_{\theta'}(z \mid x) = \log p_{\theta'}(x, z) - \log p_{\theta'}(x).$$

Substitute these identities into (38):

$$\begin{aligned}\mathrm{KL}\big(p_\theta(\cdot \mid x) \,\|\, p_{\theta'}(\cdot \mid x)\big) &= \mathbb{E}_{Z \sim p_\theta(\cdot \mid x)}\Big[\big(\log p_\theta(x, Z) - \log p_\theta(x)\big) - \big(\log p_{\theta'}(x, Z) - \log p_{\theta'}(x)\big)\Big] \\ &= \mathbb{E}_{Z \sim p_\theta(\cdot \mid x)}\big[\log p_\theta(x, Z) - \log p_{\theta'}(x, Z)\big] + \log p_{\theta'}(x) - \log p_\theta(x).\end{aligned} \tag{39}$$

The last equality uses that $\log p_\theta(x)$ and $\log p_{\theta'}(x)$ are constants in z.

Step 3: Rearrange to obtain the pointwise decomposition. Rearrange (39) to isolate the observed-data log-likelihood difference:

$$\log p_{\theta'}(x) - \log p_\theta(x) = \mathbb{E}_{Z \sim p_\theta(\cdot \mid x)}\big[\log p_{\theta'}(x, Z) - \log p_\theta(x, Z)\big] + \mathrm{KL}\big(p_\theta(\cdot \mid x) \,\|\, p_{\theta'}(\cdot \mid x)\big), \qquad x \in A. \tag{40}$$

Step 4: Integrate (40) over $X \sim P^*$. Take expectation of both sides of (40) under $X \sim P^*$. Since $P^*(A) = 1$, the equality remains valid after taking expectation. By the integrability assumption in Assumption 5.1, we may apply Fubini/Tonelli to exchange $\mathbb{E}_{P^*}$ and the conditional expectation $\mathbb{E}_{Z \sim p_\theta(\cdot \mid X)}$:

$$\begin{aligned}\mathbb{E}_{P^*}\big[\log p_{\theta'}(X) - \log p_\theta(X)\big] &= \mathbb{E}_{P^*}\mathbb{E}_{Z \sim p_\theta(\cdot \mid X)}\big[\log p_{\theta'}(X, Z) - \log p_\theta(X, Z)\big] \\ &\quad + \mathbb{E}_{P^*}\Big[\mathrm{KL}\big(p_\theta(\cdot \mid X) \,\|\, p_{\theta'}(\cdot \mid X)\big)\Big].\end{aligned} \tag{41}$$

Step 5: Identify each term with Φ and Q. By definition,

$$\Phi(\vartheta) = \mathbb{E}_{P^*}[\log p_\vartheta(X)] \quad \Longrightarrow \quad \mathbb{E}_{P^*}\big[\log p_{\theta'}(X) - \log p_\theta(X)\big] = \Phi(\theta') - \Phi(\theta).$$

Also, by definition of Q,

$$Q(\vartheta;\theta) = \mathbb{E}_{P^*}\mathbb{E}_{Z\sim p_\theta(\cdot|X)}[\log p_\vartheta(X,Z)].$$

Therefore,

$$\mathbb{E}_{P^*}\mathbb{E}_{Z\sim p_\theta(\cdot|X)}\big[\log p_{\theta'}(X,Z) - \log p_\theta(X,Z)\big] = Q(\theta';\theta) - Q(\theta;\theta).$$

Substituting these identifications into (41) yields (37).

Step 6: Inequality. Since $\mathrm{KL}(\cdot\|\cdot) \geq 0$ pointwise, its expectation is ≥ 0, giving $\Phi(\theta') - \Phi(\theta) \geq Q(\theta';\theta) - Q(\theta;\theta)$. □

Remark 5.10 (Equivalent rearrangement). Equation (37) is equivalent to

$$\Phi(\theta') - \Phi(\theta) - \mathbb{E}_{P^*}\Big[\mathrm{KL}\big(p_\theta(\cdot \mid X)\,\|\,p_{\theta'}(\cdot \mid X)\big)\Big] = Q(\theta';\theta) - Q(\theta;\theta).$$

The two forms differ only by moving the KL term across the equality.

Theorem 5.11 (Population EM ascent)**.** *Assume Assumption 5.1 on a neighborhood B, and define the EM correspondence*

$$M(\theta) = \arg\max_{\vartheta\in B} Q(\vartheta;\theta).$$

Then for every $\theta \in B$ and every selection $\theta^+ \in M(\theta)$,

$$\Phi(\theta^+) \;\geq\; \Phi(\theta). \tag{42}$$

Proof. Fix $\theta \in B$ and pick any $\theta^+ \in M(\theta)$.

Step 1: M-step improvement in Q. Since θ^+ maximizes $\vartheta \mapsto Q(\vartheta;\theta)$ over $\vartheta \in B$, we have

$$Q(\theta^+;\theta) \;\geq\; Q(\theta;\theta). \tag{43}$$

Step 2: Convert Q-improvement into Φ-improvement via the EM identity. Apply Lemma 5.9 with $\theta' = \theta^+$:

$$\Phi(\theta^+) - \Phi(\theta) = \big(Q(\theta^+;\theta) - Q(\theta;\theta)\big) + \mathbb{E}_{P^*}\Big[\mathrm{KL}\big(p_\theta(\cdot \mid X)\,\|\,p_{\theta^+}(\cdot \mid X)\big)\Big]. \tag{44}$$

Step 3: Nonnegativity of both terms. By (43), the first term on the right-hand side of (44) is ≥ 0. The second term is ≥ 0 because KL divergence is nonnegative pointwise. Therefore $\Phi(\theta^+) - \Phi(\theta) \geq 0$, i.e. (42). □

Lemma 5.12 (Sample analogue of the EM identity (fully expanded))**.** *Assume the empirical analogue of Assumption 5.1 on a neighborhood B. Then for $\theta, \theta' \in B$,*

$$\Phi_n(\theta') - \Phi_n(\theta) = \big(\widehat{Q}_n(\theta';\theta) - \widehat{Q}_n(\theta;\theta)\big) + \mathbb{E}_{X\sim P_n}\Big[\mathrm{KL}\big(p_\theta(\cdot \mid X)\,\|\,p_{\theta'}(\cdot \mid X)\big)\Big], \tag{45}$$

and in particular

$$\Phi_n(\theta') - \Phi_n(\theta) \;\geq\; \widehat{Q}_n(\theta';\theta) - \widehat{Q}_n(\theta;\theta).$$

Proof. Write the empirical measure as $P_n = \frac{1}{n}\sum_{i=1}^n \delta_{X_i}$.

Step 1: Express Φ_n and $\widehat{Q}_n$ as finite averages. By definition,

$$\Phi_n(\vartheta) = \mathbb{E}_{X\sim P_n}[\log p_\vartheta(X)] = \frac{1}{n}\sum_{i=1}^n \log p_\vartheta(X_i),$$

and

$$\widehat{Q}_n(\vartheta;\theta) = \mathbb{E}_{X\sim P_n}\mathbb{E}_{Z\sim p_\theta(\cdot|X)}[\log p_\vartheta(X,Z)] = \frac{1}{n}\sum_{i=1}^n \mathbb{E}_{Z\sim p_\theta(\cdot|X_i)}[\log p_\vartheta(X_i,Z)].$$

Step 2: Apply the pointwise identity at each datum $x = X_i$. Fix $i \in \{1, \dots, n\}$. By the empirical well-posedness assumptions, the pointwise identity (40) holds at $x = X_i$:

(46)
$$\log p_{\theta'}(X_i) - \log p_\theta(X_i) = \mathbb{E}_{Z\sim p_\theta(\cdot|X_i)}\big[\log p_{\theta'}(X_i, Z) - \log p_\theta(X_i, Z)\big] + \mathrm{KL}\big(p_\theta(\cdot \mid X_i) \,\|\, p_{\theta'}(\cdot \mid X_i)\big).$$

Step 3: Average (46) over $i = 1, \dots, n$. Sum (46) over i and divide by n:

(47)
$$\frac{1}{n}\sum_{i=1}^{n}\big(\log p_{\theta'}(X_i) - \log p_\theta(X_i)\big) = \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}_{Z\sim p_\theta(\cdot|X_i)}\big[\log p_{\theta'}(X_i, Z) - \log p_\theta(X_i, Z)\big] + \frac{1}{n}\sum_{i=1}^{n}\mathrm{KL}\big(p_\theta(\cdot \mid X_i) \,\|\, p_{\theta'}(\cdot \mid X_i$$

Step 4: Identify the averages with Φ_n, $\widehat{Q}_n$, and $\mathbb{E}_{P_n}[\mathrm{KL}]$. The left-hand side of (47) is $\Phi_n(\theta') - \Phi_n(\theta)$. The first average on the right-hand side equals $\widehat{Q}_n(\theta'; \theta) - \widehat{Q}_n(\theta; \theta)$ by Step 1. The last average is exactly $\mathbb{E}_{X\sim P_n}\big[\mathrm{KL}(p_\theta(\cdot \mid X) \| p_{\theta'}(\cdot \mid X))\big]$. Substituting these identifications yields (45). The inequality follows from nonnegativity of KL. □

Theorem 5.13 (Sample EM ascent)**.** *Let $\widehat{M}_n(\theta) = \arg\max_{\vartheta\in B} \widehat{Q}_n(\vartheta; \theta)$. Under the empirical well-posedness conditions, for every $\theta \in B$ and every selection $\theta^+ \in \widehat{M}_n(\theta)$,*

$$\Phi_n(\theta^+) \;\geq\; \Phi_n(\theta).$$

Proof. Fix $\theta \in B$ and choose $\theta^+ \in \widehat{M}_n(\theta)$.

Step 1: M-step improvement in the empirical Q-function. By definition of arg max,

$$\widehat{Q}_n(\theta^+; \theta) \;\geq\; \widehat{Q}_n(\theta; \theta).$$

Step 2: Apply the sample EM identity. Apply Lemma 5.12 with $\theta' = \theta^+$:

$$\Phi_n(\theta^+) - \Phi_n(\theta) = \big(\widehat{Q}_n(\theta^+; \theta) - \widehat{Q}_n(\theta; \theta)\big) + \mathbb{E}_{X\sim P_n}\Big[\mathrm{KL}\big(p_\theta(\cdot \mid X) \,\|\, p_{\theta^+}(\cdot \mid X)\big)\Big].$$

Step 3: Conclude by nonnegativity. Both terms on the right-hand side are nonnegative, hence $\Phi_n(\theta^+) \geq \Phi_n(\theta)$. □

Corollary 5.14 (Ascent is orbit-invariant)**.** *Assume Assumption 4.3 and the hypotheses of Theorem 5.11. If $\theta^+ \in M(\theta)$, then for every $g \in G$,*

$$\Phi(g \cdot \theta^+) \geq \Phi(g \cdot \theta) \quad \textit{and} \quad \Phi(\theta^+) \geq \Phi(\theta).$$

The analogous statement holds for sample EM and Φ_n.

Proof. Fix $g \in G$.

Step 1: Φ depends only on the observed-data marginal. By definition,

$$\Phi(\vartheta) = \mathbb{E}_{P^*}[\log p_\vartheta(X)],$$

where p_ϑ is a version of the μ-density of the observed marginal P^X_ϑ. Thus if two parameters ϑ_1, ϑ_2 induce the same observed marginal, then $p_{\vartheta_1}(x) = p_{\vartheta_2}(x)$ for μ-a.e. x, and therefore $\log p_{\vartheta_1}(X) = \log p_{\vartheta_2}(X)$ P^*-a.s. (under the positivity conditions), implying $\Phi(\vartheta_1) = \Phi(\vartheta_2)$.

Step 2: Apply observed-model invariance along the orbit. Assumption 4.3 states that $P^X_{g\cdot\vartheta} = P^X_\vartheta$ for all ϑ and g. Applying Step 1 gives

$$\Phi(g \cdot \vartheta) = \Phi(\vartheta) \qquad \text{for all } \vartheta \in B.$$

Step 3: Combine with population ascent. By Theorem 5.11, $\Phi(\theta^+) \geq \Phi(\theta)$. Using Step 2 with $\vartheta = \theta^+$ and $\vartheta = \theta$,

$$\Phi(g \cdot \theta^+) = \Phi(\theta^+) \;\geq\; \Phi(\theta) = \Phi(g \cdot \theta).$$

This proves the corollary for the population objective. The sample statement follows identically since $\Phi_n(\vartheta) = \frac{1}{n}\sum_{i=1}^n \log p_\vartheta(X_i)$ depends only on the observed marginal as well. □

6. Population theory I: fixed points, stationarity, and linearization

This section develops the population EM dynamics in a neighborhood of the misspecified target set. There are three technical goals:

(1) to relate the *misspecified projection set* $\bar{\Theta}$ to the *fixed points* of EM (set-valued in general);
(2) to record a rigorous *stationarity principle* (fixed points are stationary for the population objective) under minimal differentiation/interchange conditions, in the spirit of [12];
(3) to obtain an explicit *linearization* of the (slice-restricted) EM map via the implicit function theorem, leading to sharp local linear rates governed by the spectral radius, aligning with the fixed-point viewpoint in modern nonasymptotic EM analyses [1].

Throughout, we work on a neighborhood where the basic well-posedness conditions of Section 5 hold.

6.1. Projection points are fixed points (set-valued formulation).

Recall $\bar{\Theta} = \arg\max_\theta \Phi(\theta)$ and the EM correspondence $M(\theta) = \arg\max_{\theta'} Q(\theta'; \theta)$ (restricted to a neighborhood when needed). Because $M(\theta)$ may be set-valued, we define fixed points in the natural correspondence sense.

Definition 6.1 (Fixed points of a correspondence)**.** Let $\Gamma : \Theta \rightrightarrows \Theta$ be a correspondence. A point $\theta^\star \in \Theta$ is a *fixed point* of Γ if $\theta^\star \in \Gamma(\theta^\star)$. We write $\mathrm{Fix}(\Gamma) := \{\theta : \theta \in \Gamma(\theta)\}$.

Theorem 6.2 (Projection points are EM fixed points)**.** *Assume Assumption 5.1 holds on a neighborhood $B \subseteq \Theta$ containing $\bar{\Theta} \cap B$, and define $M(\theta) := \arg\max_{\theta' \in B} Q(\theta'; \theta)$. Then*

$$\bar{\Theta} \cap B \;\subseteq\; \mathrm{Fix}(M).$$

That is, every maximizer of Φ in B is a fixed point of the population EM correspondence restricted to B.

Proof. Fix any $\bar{\theta} \in \bar{\Theta} \cap B$.

Step 1 (Choose an EM update at $\bar{\theta}$). Let $\theta^+ \in M(\bar{\theta})$ be an arbitrary selection, so by definition

$$Q(\theta^+; \bar{\theta}) \;\geq\; Q(\eta; \bar{\theta}) \qquad \forall \eta \in B. \tag{48}$$

Step 2 (Use population ascent at $\bar{\theta}$). By Theorem 5.11 (applied on B),

$$\Phi(\theta^+) \;\geq\; \Phi(\bar{\theta}). \tag{49}$$

Step 3 (Maximality of $\bar{\theta}$ forces equality). Since $\bar{\theta} \in \bar{\Theta}$ is a maximizer of Φ over Θ, in particular over B, we also have $\Phi(\theta^+) \leq \Phi(\bar{\theta})$. Combined with (49), this yields

$$\Phi(\theta^+) = \Phi(\bar{\theta}). \tag{50}$$

Step 4 (Invoke the EM identity to force surrogate equality). Apply the exact EM identity of Lemma 5.8 with $\theta = \bar{\theta}$ and $\theta' = \theta^+$:

$$\Phi(\theta^+) - \Phi(\bar{\theta}) = \big(Q(\theta^+; \bar{\theta}) - Q(\bar{\theta}; \bar{\theta})\big) - \mathbb{E}_{P^*}\Big[\mathrm{KL}\left(p_{\bar{\theta}}(\cdot \mid X) \,\|\, p_{\theta^+}(\cdot \mid X)\right)\Big].$$

By (50), the left-hand side is 0. The KL term is nonnegative. Therefore we must have

$$Q(\theta^+;\bar{\theta}) - Q(\bar{\theta};\bar{\theta}) \;\le\; 0. \tag{51}$$

Step 5 (But θ^+ is a surrogate maximizer, so the surrogate gap is also nonnegative). From (48) with $\eta = \bar{\theta}$,

$$Q(\theta^+;\bar{\theta}) - Q(\bar{\theta};\bar{\theta}) \;\ge\; 0. \tag{52}$$

Combining (51) and (52) yields

$$Q(\theta^+;\bar{\theta}) = Q(\bar{\theta};\bar{\theta}). \tag{53}$$

Step 6 (Conclude $\bar{\theta}$ is itself a surrogate maximizer). Since θ^+ is a maximizer of $\eta \mapsto Q(\eta;\bar{\theta})$ over B and achieves value $Q(\theta^+;\bar{\theta})$, equality (53) shows that $\bar{\theta}$ attains the same maximal value. Hence $\bar{\theta} \in \arg\max_{\eta\in B} Q(\eta;\bar{\theta}) = M(\bar{\theta})$, i.e. $\bar{\theta} \in \mathrm{Fix}(M)$.

Because $\bar{\theta} \in \bar{\Theta} \cap B$ was arbitrary, the inclusion follows. □

Remark 6.3 (Quotient interpretation). Under observed invariance, $\bar{\Theta}$ is G-invariant (Lemma 4.9) and M descends to a quotient correspondence (Theorem 5.6). Thus Theorem 6.2 implies $\bar{\Theta}/G \subseteq \mathrm{Fix}(\bar{M})$ in the natural quotient sense.

6.2. **Stationarity: Fisher identity and fixed points.** The next results make precise the relationship between the gradients of the observed objective Φ and the surrogate Q. This is classical (see [12]) but we record it with explicit differentiability/interchange assumptions.

Assumption 6.4 (Differentiation under the integral (local))**.** Let $B \subseteq \Theta$ be a neighborhood. Assume:

(i) For μ-a.e. x, the map $\theta \mapsto p_\theta(x)$ is differentiable on B and $\nabla_\theta p_\theta(x)$ is measurable in x.
(ii) There exists an integrable envelope $G(x)$ such that for all $\theta \in B$,

$$\left\| \frac{\nabla_\theta p_\theta(x)}{p_\theta(x)} \right\| \mathbf{1}\{p_\theta(x) > 0\} \;\le\; G(x) \quad \text{for } \mu\text{-a.e. } x,$$

and $\mathbb{E}_{P^*}[G(X)] < \infty$.
(iii) For μ-a.e. x, the map $\theta \mapsto p_\theta(x,z)$ is differentiable for ν-a.e. z and the derivative is dominated so that differentiation may be interchanged with $\int_{\mathcal{Z}}(\cdot)\nu(dz)$ in (11).

Lemma 6.5 (Fisher identity (score as posterior expectation))**.** *Assume Assumption 6.4 holds on B. Then for every $\theta \in B$ and for μ-a.e. x with $p_\theta(x) > 0$,*

$$\nabla_\theta \log p_\theta(x) = \mathbb{E}_{Z\sim p_\theta(\cdot|x)}\big[\nabla_\theta \log p_\theta(x,Z)\big]. \tag{54}$$

Proof. Fix $\theta \in B$ and x with $p_\theta(x) > 0$.

Step 1 (Differentiate the marginal density). By (11) and Assumption 6.4(iii), we may differentiate under the integral:

$$\nabla_\theta p_\theta(x) = \nabla_\theta \int_{\mathcal{Z}} p_\theta(x,z)\,\nu(dz) = \int_{\mathcal{Z}} \nabla_\theta p_\theta(x,z)\,\nu(dz).$$

Step 2 (Rewrite in terms of complete-data score). For ν-a.e. z with $p_\theta(x,z) > 0$, $\nabla_\theta p_\theta(x,z) = p_\theta(x,z)\nabla_\theta \log p_\theta(x,z)$. Hence

$$\nabla_\theta p_\theta(x) = \int_{\mathcal{Z}} p_\theta(x,z)\nabla_\theta \log p_\theta(x,z)\,\nu(dz).$$

Step 3 (Normalize by $p_\theta(x)$ to obtain a posterior expectation). Divide both sides by $p_\theta(x) > 0$:

$$\frac{\nabla_\theta p_\theta(x)}{p_\theta(x)} = \int_{\mathcal{Z}} \frac{p_\theta(x,z)}{p_\theta(x)} \nabla_\theta \log p_\theta(x,z)\, \nu(dz) = \int_{\mathcal{Z}} \nabla_\theta \log p_\theta(x,z)\, p_\theta(dz \mid x),$$

using the definition of the posterior kernel (12). Finally, $\nabla_\theta \log p_\theta(x) = \nabla_\theta p_\theta(x)/p_\theta(x)$, giving (54). □

Lemma 6.6 (Gradient of Φ equals surrogate gradient on the diagonal)**.** *Assume Assumptions 5.1 and 6.4 on B. Assume moreover that $\theta' \mapsto Q(\theta';\theta)$ is differentiable and the derivative can be interchanged with the outer expectation. Then for every $\theta \in B$,*

$$\nabla\Phi(\theta) = \nabla_{\theta'} Q(\theta';\theta)\big|_{\theta'=\theta}. \tag{55}$$

Proof. Fix $\theta \in B$.

Step 1 (Differentiate Φ). By definition (13) and Assumption 6.4(ii), differentiation under $\mathbb{E}_{P^*}$ is justified:

$$\nabla\Phi(\theta) = \mathbb{E}_{X\sim P^*}\big[\nabla_\theta \log p_\theta(X)\big].$$

Step 2 (Apply Fisher identity inside the expectation). By Lemma 6.5, for P^*-a.e. X with $p_\theta(X) > 0$,

$$\nabla_\theta \log p_\theta(X) = \mathbb{E}_{Z\sim p_\theta(\cdot|X)}\big[\nabla_\theta \log p_\theta(X,Z)\big].$$

Insert into Step 1 and use Fubini/Tonelli justified by the well-posedness assumptions:

$$\nabla\Phi(\theta) = \mathbb{E}_{X\sim P^*}\Big[\mathbb{E}_{Z\sim p_\theta(\cdot|X)}\big[\nabla_\theta \log p_\theta(X,Z)\big]\Big]. \tag{56}$$

Step 3 (Differentiate the surrogate with respect to θ' and evaluate at $\theta' = \theta$). By definition (18),

$$Q(\theta';\theta) = \mathbb{E}_{X\sim P^*}\Big[\mathbb{E}_{Z\sim p_\theta(\cdot|X)} \log p_{\theta'}(X,Z)\Big].$$

Differentiate with respect to θ' and then set $\theta' = \theta$ (per the stated interchange assumption):

$$\nabla_{\theta'} Q(\theta';\theta)\big|_{\theta'=\theta} = \mathbb{E}_{X\sim P^*}\Big[\mathbb{E}_{Z\sim p_\theta(\cdot|X)} \nabla_{\theta'} \log p_{\theta'}(X,Z)\big|_{\theta'=\theta}\Big] = \mathbb{E}_{X\sim P^*}\Big[\mathbb{E}_{Z\sim p_\theta(\cdot|X)} \nabla_\theta \log p_\theta(X,Z)\Big].$$

Comparing with (56) yields (55). □

Theorem 6.7 (Fixed points are stationary for Φ)**.** *Assume Assumptions 5.1 and 6.4 on a neighborhood B, and assume $\theta' \mapsto Q(\theta';\theta)$ is differentiable for $\theta, \theta' \in B$ with the interchange conditions in Lemma 6.6. If $\theta^\star \in B$ is a fixed point of EM in the sense that $\theta^\star \in M(\theta^\star)$ and the maximization is unconstrained (interior point of B), then*

$$\nabla\Phi(\theta^\star) = 0.$$

Proof. Fix $\theta^\star \in B$ with $\theta^\star \in M(\theta^\star)$.

Step 1 (First-order optimality for the surrogate). Because $\theta^\star$ maximizes $\theta' \mapsto Q(\theta';\theta^\star)$ over a neighborhood and is an interior maximizer, the first-order necessary condition gives

$$\nabla_{\theta'} Q(\theta';\theta^\star)\big|_{\theta'=\theta^\star} = 0.$$

Step 2 (Identify this gradient with $\nabla\Phi(\theta^\star)$). By Lemma 6.6 applied at $\theta = \theta^\star$,

$$\nabla\Phi(\theta^\star) = \nabla_{\theta'} Q(\theta';\theta^\star)\big|_{\theta'=\theta^\star}.$$

Combining with Step 1 yields $\nabla\Phi(\theta^\star) = 0$. □

Remark 6.8 (Constrained/slice stationarity). When the M-step is performed on a slice S (or under constraints), the corresponding conclusion is stationarity with respect to feasible directions (i.e. vanishing of the projected gradient on the tangent space). We state and use the unconstrained form for clarity; the slice-restricted form follows by replacing ordinary gradients with gradients restricted to local coordinates on the slice.

6.3. **Slice-restricted EM map and differentiability.** To obtain sharp local rates, we work with a single-valued slice map. Fix a local section S as in Assumption 4.10 and let $B_S \subseteq S$ be a neighborhood containing a representative $\theta^\star \in S$ of an orbit in $\bar{\Theta}/G$.

Assumption 6.9 (Strong concavity in the M-step variable)**.** On B_S, for each fixed $\theta \in B_S$ the function $\theta' \mapsto Q(\theta';\theta)$ is twice continuously differentiable and *strictly* (equivalently, strongly) concave on B_S, with Hessian $\nabla^2_{\theta'\theta'} Q(\theta';\theta)$ negative definite for all $\theta', \theta \in B_S$.

Under Assumption 6.9, the M-step on the slice is unique and defines a map.

Definition 6.10 (Slice EM map)**.** Assume Assumption 6.9. Define the slice EM map $T : B_S \to B_S$ by

$$T(\theta) := \operatorname*{arg\,max}_{\theta' \in B_S} Q(\theta';\theta), \tag{57}$$

which is single-valued by strict concavity.

Lemma 6.11 (First-order characterization of T)**.** *Assume Assumption 6.9. Then $\theta' = T(\theta)$ if and only if*

$$\nabla_{\theta'} Q(\theta';\theta) = 0 \quad \text{and} \quad \theta' \in B_S. \tag{58}$$

Proof. Fix $\theta \in B_S$.

Step 1 (Necessity). If $\theta' = T(\theta)$ is the unique maximizer of the concave C^1 function $\theta' \mapsto Q(\theta';\theta)$ on an open neighborhood, then $\nabla_{\theta'} Q(\theta';\theta) = 0$ at the maximizer (first-order condition).

Step 2 (Sufficiency). Conversely, if $\nabla_{\theta'} Q(\theta';\theta) = 0$ and the Hessian is negative definite, then θ' is a strict local maximizer. By strict concavity on B_S, there can be only one maximizer, hence $\theta' = T(\theta)$. □

We now linearize T.

Theorem 6.12 (Differentiability and Jacobian formula for T)**.** *Assume Assumption 6.9 and additionally that Q is C^2 jointly in (θ', θ) on $B_S \times B_S$. Fix $\theta^\star \in B_S$ and suppose $\theta^\star$ is a fixed point: $T(\theta^\star) = \theta^\star$. Then T is continuously differentiable in a neighborhood of $\theta^\star$ and its Jacobian satisfies*

$$DT(\theta) = -\Big(\nabla^2_{\theta'\theta'} Q(T(\theta);\theta)\Big)^{-1} \Big(\nabla^2_{\theta'\theta} Q(T(\theta);\theta)\Big), \tag{59}$$

for θ near $\theta^\star$. In particular,

$$DT(\theta^\star) = -\Big(\nabla^2_{\theta'\theta'} Q(\theta^\star;\theta^\star)\Big)^{-1} \Big(\nabla^2_{\theta'\theta} Q(\theta^\star;\theta^\star)\Big). \tag{60}$$

Proof. **Standing conventions.** We work in $\mathbb{R}^d$. Gradients are column vectors. For a C^2 scalar function $g(\theta', \theta)$, $\nabla_{\theta'} g$ denotes the gradient w.r.t. θ' and $\nabla^2_{\theta'\theta'} g$ the Hessian w.r.t. θ'. All Jacobians are taken in the usual Fréchet sense, and since we are in finite dimension, this coincides with entrywise differentiation.

Assumption 6.9 unpacked. There exists $\mu > 0$ such that for every fixed $\theta \in B_S$, the map

$$\theta' \mapsto Q(\theta'; \theta)$$

is μ-*strongly concave* on B_S, meaning that for all $\theta'_1, \theta'_2 \in B_S$,

$$Q(\theta'_2; \theta) \leq Q(\theta'_1; \theta) + \langle \nabla_{\theta'} Q(\theta'_1; \theta),\, \theta'_2 - \theta'_1 \rangle - \frac{\mu}{2} \|\theta'_2 - \theta'_1\|^2. \tag{61}$$

Since Q is C^2 in θ' (indeed jointly C^2), (61) is equivalent to the pointwise Hessian bound

$$\nabla^2_{\theta'\theta'} Q(\theta'; \theta) \;\preceq\; -\mu I_d \qquad \text{for all } (\theta', \theta) \in B_S \times B_S, \tag{62}$$

i.e. the Hessian is negative definite with eigenvalues bounded above by $-\mu$.

A calculus lemma we will use (stationary point of strongly concave function).

Lemma 6.13. *Fix $\theta \in B_S$ and define $g(\theta') := Q(\theta'; \theta)$. Assume g is differentiable and μ-strongly concave on B_S. If $\bar{\theta}' \in B_S$ satisfies $\nabla g(\bar{\theta}') = 0$, then $\bar{\theta}'$ is the unique global maximizer of g on B_S.*

Proof. Take any $\theta' \in B_S$ and apply (61) with $\theta'_1 = \bar{\theta}'$ and $\theta'_2 = \theta'$:

$$g(\theta') \leq g(\bar{\theta}') + \langle \nabla g(\bar{\theta}'),\, \theta' - \bar{\theta}' \rangle - \frac{\mu}{2} \|\theta' - \bar{\theta}'\|^2 = g(\bar{\theta}') - \frac{\mu}{2} \|\theta' - \bar{\theta}'\|^2 \leq g(\bar{\theta}').$$

If $\theta' \neq \bar{\theta}'$, then $\|\theta' - \bar{\theta}'\|^2 > 0$ so the inequality is strict, hence $\bar{\theta}'$ is the unique maximizer. □

Define the implicit map. Define $F : B_S \times B_S \to \mathbb{R}^d$ by

$$F(\theta', \theta) := \nabla_{\theta'} Q(\theta'; \theta).$$

Since Q is C^2 jointly in (θ', θ), each component of F is C^1 jointly, hence F is C^1 as a map $B_S \times B_S \to \mathbb{R}^d$.

Step 1 (Implicit equation defining T). By definition, for each $\theta \in B_S$, $T(\theta)$ is the (unique) maximizer of $\theta' \mapsto Q(\theta'; \theta)$ over B_S. By Lemma 6.11 (first-order optimality for the slice maximizer, valid under our hypotheses and the fact that we will work locally around $\theta^\star$ in the interior of B_S), we have the stationarity condition

$$\nabla_{\theta'} Q(T(\theta); \theta) = 0,$$

i.e.

$$F(T(\theta), \theta) = 0 \qquad \text{for all } \theta \in B_S. \tag{63}$$

In particular, since $T(\theta^\star) = \theta^\star$,

$$F(\theta^\star, \theta^\star) = F(T(\theta^\star), \theta^\star) = 0.$$

Step 2 (Nonsingularity of the partial Jacobian in θ'). We compute the Jacobian of F with respect to its first argument. Because $F(\theta', \theta) = \nabla_{\theta'} Q(\theta'; \theta)$ and Q is C^2 jointly, the derivative of F in the θ' direction is the Hessian:

$$D_{\theta'} F(\theta', \theta) = \nabla^2_{\theta'\theta'} Q(\theta'; \theta). \tag{64}$$

By (62), for every $(\theta', \theta) \in B_S \times B_S$ the symmetric matrix $\nabla^2_{\theta'\theta'} Q(\theta'; \theta)$ is negative definite. In particular it is invertible. (Indeed, if $Hv = 0$ for a negative definite symmetric matrix H, then $0 = v^\top H v < 0$ for $v \neq 0$, a contradiction; hence $\ker(H) = \{0\}$.) Thus $D_{\theta'} F(\theta^\star, \theta^\star)$ is invertible.

Step 3 (Apply the implicit function theorem, and identify the resulting solution with T).

(a) Implicit function theorem (statement). Since F is C^1 and $F(\theta^\star, \theta^\star) = 0$ with $D_{\theta'}F(\theta^\star, \theta^\star)$ invertible, the (classical) implicit function theorem yields: there exist neighborhoods U and V of $\theta^\star$ (with $U \times V \subseteq B_S \times B_S$) and a unique C^1 map

$$\widetilde{T} : V \to U$$

such that

$$F(\widetilde{T}(\theta), \theta) = 0 \qquad \text{for all } \theta \in V, \tag{65}$$

and moreover any $(\theta', \theta) \in U \times V$ solving $F(\theta', \theta) = 0$ must satisfy $\theta' = \widetilde{T}(\theta)$.

(b) $\widetilde{T}$ coincides with T on V. Fix $\theta \in V$. By (65), $\widetilde{T}(\theta)$ is a stationary point of $\theta' \mapsto Q(\theta'; \theta)$. By Lemma 6.13 (applied to $g(\theta') = Q(\theta'; \theta)$), this stationary point is the unique global maximizer. But $T(\theta)$ is *defined* to be that unique maximizer. Hence

$$\widetilde{T}(\theta) = T(\theta) \qquad \text{for all } \theta \in V.$$

Therefore T is C^1 on V.

(c) Continuous differentiability. The IFT already gives that $\widetilde{T}$ is C^1, hence DT is continuous on V (since $T = \widetilde{T}$ on V).

Step 4 (Differentiate the implicit equation and derive the Jacobian formula). On the neighborhood V we have the identity $F(T(\theta), \theta) \equiv 0$. Differentiate this identity with respect to θ.

To make the chain rule completely explicit, write $F = (F_1, \dots, F_d)^\top$ and $T = (T_1, \dots, T_d)^\top$. For each $i \in \{1, \dots, d\}$ define the scalar function

$$G_i(\theta) := F_i(T(\theta), \theta).$$

Then $G_i(\theta) \equiv 0$ on V. Fix $j \in \{1, \dots, d\}$ and differentiate G_i with respect to θ_j:

$$0 = \frac{\partial G_i}{\partial \theta_j}(\theta) = \sum_{k=1}^{d} \frac{\partial F_i}{\partial \theta'_k}(T(\theta), \theta)\, \frac{\partial T_k}{\partial \theta_j}(\theta) + \frac{\partial F_i}{\partial \theta_j}(T(\theta), \theta).$$

In matrix form, letting $D_{\theta'}F(T(\theta), \theta)$ denote the $d \times d$ Jacobian of F in its first argument and $D_\theta F(T(\theta), \theta)$ the $d \times d$ Jacobian in its second argument, the above identities for all i, j are exactly

$$D_{\theta'}F(T(\theta), \theta)\, DT(\theta) + D_\theta F(T(\theta), \theta) = 0. \tag{66}$$

By Step 2, $D_{\theta'}F(T(\theta), \theta)$ is invertible for θ near $\theta^\star$ (in fact on V), so we can solve (66) for $DT(\theta)$:

$$DT(\theta) = -\Big(D_{\theta'}F(T(\theta), \theta)\Big)^{-1} D_\theta F(T(\theta), \theta). \tag{67}$$

Finally, identify the Jacobians of F with Hessians of Q. From $F(\theta', \theta) = \nabla_{\theta'} Q(\theta'; \theta)$ and joint C^2 regularity,

$$D_{\theta'}F(\theta', \theta) = \nabla^2_{\theta'\theta'} Q(\theta'; \theta), \qquad D_\theta F(\theta', \theta) = \nabla^2_{\theta'\theta} Q(\theta'; \theta),$$

where $\nabla^2_{\theta'\theta} Q$ denotes the matrix of mixed second derivatives $\big[\partial^2 Q / \partial \theta'_i\, \partial \theta_j\big]_{i,j}$. Substituting into (67) yields (59). Evaluating at $\theta = \theta^\star$ and using $T(\theta^\star) = \theta^\star$ gives (60). □

6.4. **Sharp local linear rate via spectral radius.** The Jacobian formula (60) yields an exact linearization of the EM dynamics on the slice. A general fact in smooth fixed-point theory is that the best achievable local linear rate is governed by the spectral radius of the Jacobian at the fixed point.

Theorem 6.14 (Sharp local linear rate (spectral-radius governed))**.** *Assume the hypotheses of Theorem 6.12. Let $A := DT(\theta^\star)$ and assume $\rho(A) < 1$. Then for every $\varepsilon > 0$ there exist a norm $\|\cdot\|_\varepsilon$ on $\mathbb{R}^d$, a neighborhood $U_\varepsilon \subseteq B_S$ of $\theta^\star$, and a constant $C_\varepsilon < \infty$ such that for all $\theta_0 \in U_\varepsilon$, the iterates $\theta_{t+1} = T(\theta_t)$ satisfy*

$$\|\theta_t - \theta^\star\|_\varepsilon \le C_\varepsilon\big(\rho(A) + \varepsilon\big)^t \|\theta_0 - \theta^\star\|_\varepsilon, \qquad t = 0, 1, 2, \dots. \tag{68}$$

Consequently, by norm equivalence, the same statement holds in the Euclidean norm with a different constant.

Proof. Fix $\varepsilon > 0$.

Preliminaries (induced operator norms). Given any norm $\|\cdot\|_\varepsilon$ on $\mathbb{R}^d$, define the induced operator norm

$$\|M\|_{\mathrm{op},\varepsilon} := \sup_{v \neq 0} \frac{\|Mv\|_\varepsilon}{\|v\|_\varepsilon}.$$

Then for all matrices M, N and vectors v:

(i) $\|Mv\|_\varepsilon \le \|M\|_{\mathrm{op},\varepsilon}\, \|v\|_\varepsilon$ (by definition of the supremum),
(ii) $\|MN\|_{\mathrm{op},\varepsilon} \le \|M\|_{\mathrm{op},\varepsilon} \|N\|_{\mathrm{op},\varepsilon}$ (submultiplicativity: apply (i) twice and take sup).

Step 1 (Construct an adapted norm with $\|A\|_{\mathrm{op}}$ close to $\rho(A)$). We prove the following lemma and then apply it with $\eta = \varepsilon/2$.

Lemma 6.15 (Adapted norm lemma)**.** *For any $A \in \mathbb{R}^{d\times d}$ and any $\eta > 0$, there exists a norm $\|\cdot\|_\eta$ on $\mathbb{R}^d$ such that its induced operator norm satisfies*

$$\|A\|_{\mathrm{op},\eta} \le \rho(A) + \eta.$$

Proof. It is convenient to work over $\mathbb{C}^d$ and then restrict to $\mathbb{R}^d$ at the end. By Schur triangularization, there exists a unitary $U \in \mathbb{C}^{d\times d}$ such that

$$U^* A U = T,$$

where T is upper triangular and its diagonal entries are the eigenvalues $\lambda_1, \dots, \lambda_d$ of A (counted with algebraic multiplicity). Write $T = \Lambda + N$ where $\Lambda = \mathrm{diag}(\lambda_1, \dots, \lambda_d)$ and N is strictly upper triangular (i.e. $N_{ii} = 0$ and $N_{ij} = 0$ for $i \ge j$).

Fix $\delta \in (0, 1)$ and define the diagonal scaling

$$D_\delta := \mathrm{diag}(1, \delta, \delta^2, \dots, \delta^{d-1}).$$

Consider the similar matrix

$$T_\delta := D_\delta^{-1} T D_\delta = D_\delta^{-1}\Lambda D_\delta + D_\delta^{-1} N D_\delta = \Lambda + N_\delta,$$

since Λ commutes with D_δ. For $i < j$,

$$(N_\delta)_{ij} = (D_\delta^{-1} N D_\delta)_{ij} = \delta^{j-i} N_{ij},$$

so every strictly upper-triangular entry is multiplied by a factor δ^{j-i}.

Now use the ℓ_∞ norm on $\mathbb{C}^d$: $\|x\|_\infty = \max_i |x_i|$, whose induced operator norm is

$$\|M\|_\infty = \max_{1\le i\le d} \sum_{j=1}^{d} |M_{ij}|.$$

Then

$$\|T_\delta\|_\infty \le \|\Lambda\|_\infty + \|N_\delta\|_\infty = \max_i |\lambda_i| + \|N_\delta\|_\infty = \rho(A) + \|N_\delta\|_\infty.$$

Moreover,

$$\|N_\delta\|_\infty = \max_i \sum_{j>i} |(N_\delta)_{ij}| = \max_i \sum_{j>i} \delta^{j-i} |N_{ij}| \xrightarrow[\delta\downarrow 0]{} 0,$$

by dominated convergence for a finite sum. Hence choose δ small enough that $\|N_\delta\|_\infty \le \eta$. Then $\|T_\delta\|_\infty \le \rho(A) + \eta$.

Define a norm on $\mathbb{C}^d$ by

$$\|v\|_\eta := \|D_\delta^{-1} U^* v\|_\infty.$$

This is a norm because $v \mapsto D_\delta^{-1} U^* v$ is a linear isomorphism and $\|\cdot\|_\infty$ is a norm. Its induced operator norm satisfies, for $v \neq 0$,

$$\begin{aligned} \frac{\|Av\|_\eta}{\|v\|_\eta} &= \frac{\|D_\delta^{-1} U^* A v\|_\infty}{\|D_\delta^{-1} U^* v\|_\infty} = \frac{\|D_\delta^{-1} U^* A U D_\delta w\|_\infty}{\|w\|_\infty} \quad \text{where } w := D_\delta^{-1} U^* v \neq 0 \\ &= \frac{\|T_\delta w\|_\infty}{\|w\|_\infty} \le \|T_\delta\|_\infty. \end{aligned}$$

Taking the supremum over $v \neq 0$ yields $\|A\|_{\mathrm{op},\eta} \le \|T_\delta\|_\infty \le \rho(A) + \eta$. Finally, restrict $\|\cdot\|_\eta$ to $\mathbb{R}^d \subset \mathbb{C}^d$ to obtain a real norm with the same bound. □

Apply Lemma 6.15 with $\eta = \varepsilon/2$. This yields a norm $\|\cdot\|_\varepsilon$ such that

$$\|A\|_{\mathrm{op},\varepsilon} \le \rho(A) + \varepsilon/2. \tag{69}$$

Step 2 (Continuity of DT gives a uniform derivative bound on a convex neighborhood). By Theorem 6.12, T is C^1 on some open neighborhood V of $\theta^\star$ (hence $DT : V \to \mathbb{R}^{d\times d}$ is continuous). By continuity of DT at $\theta^\star$, there exists an open neighborhood $U_0 \subseteq V$ of $\theta^\star$ such that

$$\|DT(\theta) - A\|_{\mathrm{op},\varepsilon} \le \varepsilon/2 \qquad \text{for all } \theta \in U_0. \tag{70}$$

Since $U_0 \cap B_S$ is a neighborhood of $\theta^\star$ (in $\mathbb{R}^d$), choose $r_\varepsilon > 0$ such that the $\|\cdot\|_\varepsilon$-ball

$$U_\varepsilon := \{\theta \in \mathbb{R}^d : \ \|\theta - \theta^\star\|_\varepsilon < r_\varepsilon\}$$

satisfies $U_\varepsilon \subseteq U_0 \cap B_S$. In particular, U_ε is open and convex.

For any $\theta \in U_\varepsilon$, the triangle inequality and (69)–(70) yield

$$\|DT(\theta)\|_{\mathrm{op},\varepsilon} \le \|A\|_{\mathrm{op},\varepsilon} + \|DT(\theta) - A\|_{\mathrm{op},\varepsilon} \le (\rho(A) + \varepsilon/2) + \varepsilon/2 =: \gamma_\varepsilon,$$

so

$$\gamma_\varepsilon = \rho(A) + \varepsilon. \tag{71}$$

Since $\rho(A) < 1$, taking $\varepsilon > 0$ small ensures $\gamma_\varepsilon < 1$.

Step 3 (Integral mean-value identity and contraction on U_ε). We use the following standard calculus lemma, proved for completeness.

Lemma 6.16 (Integral mean-value identity)**.** *Let $f : W \to \mathbb{R}^d$ be C^1 on an open convex set $W \subseteq \mathbb{R}^d$. Then for any $\theta, \theta' \in W$,*

$$f(\theta) - f(\theta') = \int_0^1 Df\big(\theta' + t(\theta - \theta')\big)\,(\theta - \theta')\,dt.$$

Proof. Define $g : [0,1] \to \mathbb{R}^d$ by $g(t) = f(\theta' + t(\theta - \theta'))$. Because W is convex, $\theta' + t(\theta - \theta') \in W$ for all $t \in [0,1]$, and since f is C^1, g is C^1. By the chain rule,

$$g'(t) = Df\big(\theta' + t(\theta - \theta')\big)\,(\theta - \theta').$$

By the fundamental theorem of calculus (applied componentwise),

$$f(\theta) - f(\theta') = g(1) - g(0) = \int_0^1 g'(t)\,dt = \int_0^1 Df\big(\theta' + t(\theta - \theta')\big)\,(\theta - \theta')\,dt. \qquad \square$$

Apply Lemma 6.16 with $f = T$ and $W = U_\varepsilon$ (convex). For $\theta, \theta' \in U_\varepsilon$,

$$T(\theta) - T(\theta') = \int_0^1 DT\big(\theta' + t(\theta - \theta')\big)(\theta - \theta')\,dt.$$

Taking $\|\cdot\|_\varepsilon$ norms and using the induced norm bound,

$$\begin{aligned}
\|T(\theta) - T(\theta')\|_\varepsilon &\le \int_0^1 \Big\|DT\big(\theta' + t(\theta - \theta')\big)(\theta - \theta')\Big\|_\varepsilon\,dt \\
&\le \int_0^1 \Big\|DT\big(\theta' + t(\theta - \theta')\big)\Big\|_{\mathrm{op},\varepsilon}\,dt\;\|\theta - \theta'\|_\varepsilon \\
&\le \int_0^1 \gamma_\varepsilon\,dt\;\|\theta - \theta'\|_\varepsilon = \gamma_\varepsilon\|\theta - \theta'\|_\varepsilon,
\end{aligned}$$

where we used (71) and that $\theta' + t(\theta - \theta') \in U_\varepsilon$ for all $t \in [0,1]$. Thus T is a contraction on U_ε in $\|\cdot\|_\varepsilon$.

Step 4 (Forward invariance and geometric decay to the fixed point). Since $T(\theta^\star) = \theta^\star$, the contraction bound with $\theta' = \theta^\star$ gives, for any $\theta \in U_\varepsilon$,

$$\|T(\theta) - \theta^\star\|_\varepsilon = \|T(\theta) - T(\theta^\star)\|_\varepsilon \le \gamma_\varepsilon\|\theta - \theta^\star\|_\varepsilon. \tag{72}$$

In particular, if $\|\theta - \theta^\star\|_\varepsilon < r_\varepsilon$, then

$$\|T(\theta) - \theta^\star\|_\varepsilon \le \gamma_\varepsilon r_\varepsilon < r_\varepsilon \qquad (\text{since } \gamma_\varepsilon < 1),$$

so $T(\theta) \in U_\varepsilon$. Hence U_ε is forward-invariant: $T(U_\varepsilon) \subseteq U_\varepsilon$.

Now let $\theta_0 \in U_\varepsilon$ and define $\theta_{t+1} = T(\theta_t)$. By invariance, $\theta_t \in U_\varepsilon$ for all t, so applying (72) at $\theta = \theta_t$ yields

$$\|\theta_{t+1} - \theta^\star\|_\varepsilon \le \gamma_\varepsilon\|\theta_t - \theta^\star\|_\varepsilon.$$

By induction on t (using repeatedly the previous inequality),

$$\|\theta_t - \theta^\star\|_\varepsilon \le \gamma_\varepsilon^t\,\|\theta_0 - \theta^\star\|_\varepsilon = (\rho(A) + \varepsilon)^t\,\|\theta_0 - \theta^\star\|_\varepsilon.$$

This is (68) with $C_\varepsilon = 1$.

Step 5 (Norm equivalence to return to Euclidean norm). We record the equivalence-of-norms lemma with proof.

Lemma 6.17 (Equivalence of norms)**.** *Let $\|\cdot\|_a$ and $\|\cdot\|_b$ be norms on $\mathbb{R}^d$. Then there exist constants $0 < m \le M < \infty$ such that*

$$m\|v\|_b \le \|v\|_a \le M\|v\|_b \qquad \textit{for all } v \in \mathbb{R}^d.$$

Proof. Let $S_b := \{v \in \mathbb{R}^d : \|v\|_b = 1\}$. In finite dimensions, S_b is compact. The map $v \mapsto \|v\|_a$ is continuous, hence attains its minimum and maximum on S_b:

$$m := \min_{v \in S_b} \|v\|_a, \qquad M := \max_{v \in S_b} \|v\|_a.$$

We have $M < \infty$ by compactness/continuity. Also $m > 0$ because $\|v\|_a > 0$ for $v \neq 0$ and $0 \notin S_b$. For any $v \neq 0$, write $v = \|v\|_b \cdot (v/\|v\|_b)$ with $v/\|v\|_b \in S_b$. By homogeneity,

$$\|v\|_a = \|v\|_b \cdot \left\| \frac{v}{\|v\|_b} \right\|_a \in [m\|v\|_b,\ M\|v\|_b].$$

The case $v = 0$ is trivial. □

Apply Lemma 6.17 with $\|\cdot\|_a = \|\cdot\|_\varepsilon$ and $\|\cdot\|_b = \|\cdot\|_2$ to obtain constants $0 < m_\varepsilon \leq M_\varepsilon < \infty$ with $m_\varepsilon\|v\|_2 \leq \|v\|_\varepsilon \leq M_\varepsilon\|v\|_2$. Combining with (68) yields an Euclidean-norm version with multiplicative constant $M_\varepsilon/m_\varepsilon$. □

Remark 6.18 (Where "sharpness" enters). The proof above is purely dynamical: for C^1 fixed-point iterations, $\rho(DT(\theta^\star))$ is the intrinsic local linear rate. No improvement is possible without additional structure beyond smoothness and fixed-point stability. In EM, the Jacobian (60) makes this rate explicit in terms of second derivatives of the surrogate.

7. Population theory II: contraction to sets in general IPMs

Section 6 provided a sharp local linear rate for the slice-restricted population EM map $\theta_{t+1} = T(\theta_t)$ in identifiable coordinates, governed by $\rho(DT(\theta^\star))$. This section converts that local dynamical statement into a *statistically meaningful* convergence guarantee: convergence (i) on the quotient, (ii) to the misspecified target set $\bar{\Theta}/G$, and (iii) measured in a general IPM.

There are two logical steps:

(1) establish a contraction statement to the target *set* on a slice (because under misspecification the target is typically set-valued);
(2) transfer that statement from slice coordinates to distributional discrepancy via a one-sided modulus for the model map $\theta \mapsto P_\theta$ into $d_{\mathcal{F}}$, and then to the quotient via the quotient IPM $\bar{d}_{\mathcal{F}}$.

We keep the assumptions explicit; they are minimal in the sense that no statement in a general IPM can be sharper without further model regularity.

7.1. Distance-to-set dynamics on a slice.

Fix a local section S and a neighborhood $B_S \subseteq S$. Define the slice target set

$$\bar{\Theta}_S := \bar{\Theta} \cap B_S. \tag{73}$$

(If $\bar{\Theta}$ is G-invariant, then $\bar{\Theta}_S$ contains at most one representative per orbit in B_S, but it may still be set-valued if multiple orbits of maximizers intersect B_S.)

Definition 7.1 (Distance to a set in slice norm)**.** Let $(\mathbb{R}^d, \|\cdot\|)$ be the ambient normed space containing B_S. For $\theta \in B_S$, define

$$\mathrm{dist}_{\|\cdot\|}(\theta, \bar{\Theta}_S) := \inf_{\bar{\theta} \in \bar{\Theta}_S} \|\theta - \bar{\theta}\|.$$

We now state a contraction-to-set property that is implied by the sharp local rate around each fixed point when the target set is finite or forms a locally stable manifold; we keep it as an explicit assumption because in full generality the geometry of $\bar{\Theta}$ can be complicated under misspecification [6].

Assumption 7.2 (Uniform local contraction to the slice target set)**.** There exist a neighborhood $B_S \subseteq S$, a constant $\gamma \in (0,1)$, and a (possibly smaller) target subset $\bar{\Theta}_S \subseteq B_S$ such that for every $\theta \in B_S$,

$$\operatorname{dist}_{\|\cdot\|}\big(T(\theta), \bar{\Theta}_S\big) \;\le\; \gamma \operatorname{dist}_{\|\cdot\|}(\theta, \bar{\Theta}_S). \tag{74}$$

Remark 7.3 (When does Assumption 7.2 hold?). If $\bar{\Theta}_S = \{\theta^\star\}$ is a singleton and $\rho(DT(\theta^\star)) < 1$, then (74) holds locally by Theorem 6.14. More generally, if $\bar{\Theta}_S$ is a finite set of isolated stable fixed points with disjoint basins inside B_S, then (74) holds on each basin and hence on their union (with γ the maximum contraction factor). Set-valued projection sets with nontrivial geometry can require additional regularity to ensure a uniform (74); we keep this explicit to avoid implicitly imposing hidden identifiability assumptions.

Theorem 7.4 (Geometric convergence to the slice target set)**.** *Assume Assumption 7.2. Then for all $\theta_0 \in B_S$ and iterates $\theta_{t+1} = T(\theta_t)$,*

$$\operatorname{dist}_{\|\cdot\|}(\theta_t, \bar{\Theta}_S) \;\le\; \gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S), \qquad t = 0, 1, 2, \dots. \tag{75}$$

Proof. **Definitions used in this proof.** For a nonempty set $C \subseteq \mathbb{R}^d$ and a norm $\|\cdot\|$ on $\mathbb{R}^d$, define the distance-to-set function

$$\operatorname{dist}_{\|\cdot\|}(\theta, C) \;:=\; \inf_{\vartheta \in C} \|\theta - \vartheta\|.$$

Assumption 7.2 states that there exists $\gamma \in [0,1)$ such that for all $\theta \in B_S$,

$$\operatorname{dist}_{\|\cdot\|}(T(\theta), \bar{\Theta}_S) \;\le\; \gamma \operatorname{dist}_{\|\cdot\|}(\theta, \bar{\Theta}_S). \tag{76}$$

We prove (75) by a fully explicit induction.

Step 1 (Base case, $t = 0$). By the convention $\gamma^0 = 1$, the right-hand side of (75) at $t = 0$ is

$$\gamma^0 \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S) = \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S),$$

which equals the left-hand side. Hence (75) holds with equality for $t = 0$.

Step 2 (Inductive step). Fix $t \ge 0$ and assume as the induction hypothesis that

$$\operatorname{dist}_{\|\cdot\|}(\theta_t, \bar{\Theta}_S) \;\le\; \gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S). \tag{77}$$

We must prove

$$\operatorname{dist}_{\|\cdot\|}(\theta_{t+1}, \bar{\Theta}_S) \;\le\; \gamma^{t+1} \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S).$$

By definition of the iteration, $\theta_{t+1} = T(\theta_t)$, so

$$\operatorname{dist}_{\|\cdot\|}(\theta_{t+1}, \bar{\Theta}_S) = \operatorname{dist}_{\|\cdot\|}(T(\theta_t), \bar{\Theta}_S).$$

Now apply the contraction-to-set inequality (76) with $\theta = \theta_t$ (note $\theta_t \in B_S$ since $T : B_S \to B_S$):

$$\operatorname{dist}_{\|\cdot\|}(T(\theta_t), \bar{\Theta}_S) \le \gamma \operatorname{dist}_{\|\cdot\|}(\theta_t, \bar{\Theta}_S).$$

Finally, apply the induction hypothesis (77) to bound the right-hand side:

$$\gamma \operatorname{dist}_{\|\cdot\|}(\theta_t, \bar{\Theta}_S) \le \gamma\, \gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S) = \gamma^{t+1} \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S),$$

where we used the exponent law $\gamma\gamma^t = \gamma^{t+1}$. Combining the displayed inequalities proves the inductive step.

Step 3 (Conclusion). By the principle of mathematical induction, (75) holds for all integers $t \ge 0$. □

7.2. **Transfer from slice norm to IPM discrepancy.** To obtain distributional convergence, we require regularity of the model map $\theta \mapsto P_\theta$ into $d_{\mathcal{F}}$. Because we only need an upper bound, a one-sided modulus is sufficient.

Assumption 7.5 (One-sided IPM modulus on the slice basin)**.** There exists a nondecreasing function $\omega_{\mathcal{F}} : [0,\infty) \to [0,\infty)$ with $\omega_{\mathcal{F}}(0) = 0$ such that for all $\theta, \theta' \in B_S$,

$$d_{\mathcal{F}}(P_\theta, P_{\theta'}) \;\le\; \omega_{\mathcal{F}}\big(\|\theta - \theta'\|\big). \tag{78}$$

Remark 7.6. Assumption 7.5 is a local continuity/Lipschitz condition of the model map into the chosen IPM. In smooth parametric models and strong IPMs (e.g. total variation under sufficient regularity), one often has $\omega_{\mathcal{F}}(r) \lesssim r$ locally. For weaker IPMs, $\omega_{\mathcal{F}}$ may be sublinear or depend on moments. We do not specialize because $\mathcal{F}$ is arbitrary.

Lemma 7.7 (IPM distance-to-set controlled by parameter distance-to-set)**.** *Assume Assumption 7.5. Then for all $\theta \in B_S$,*

$$\inf_{\bar\theta \in \bar\Theta_S} d_{\mathcal{F}}(P_\theta, P_{\bar\theta}) \;\le\; \omega_{\mathcal{F}}\Big(\operatorname{dist}_{\|\cdot\|}(\theta, \bar\Theta_S)\Big). \tag{79}$$

Proof. **Definitions used.** For a nonempty set $C \subseteq \mathbb{R}^d$ and a norm $\|\cdot\|$, the distance-to-set is

$$\operatorname{dist}_{\|\cdot\|}(\theta, C) := \inf_{\vartheta \in C} \|\theta - \vartheta\|.$$

The integral probability metric (IPM) generated by $\mathcal{F}$ is

$$d_{\mathcal{F}}(P, Q) := \sup_{f \in \mathcal{F}} \Big| \mathbb{E}_P[f] - \mathbb{E}_Q[f] \Big|$$

(whenever the expectations are well-defined), and Assumption 7.5 provides a nondecreasing modulus $\omega_{\mathcal{F}} : [0,\infty) \to [0,\infty)$ such that for all $\theta, \theta' \in B_S$,

$$d_{\mathcal{F}}(P_\theta, P_{\theta'}) \le \omega_{\mathcal{F}}(\|\theta - \theta'\|). \tag{80}$$

(If your Assumption 7.5 includes right-continuity of $\omega_{\mathcal{F}}$, then we use it below. If not, see the remark at the end of the proof.)

Fix $\theta \in B_S$ and let

$$r := \operatorname{dist}_{\|\cdot\|}(\theta, \bar\Theta_S) = \inf_{\bar\theta \in \bar\Theta_S} \|\theta - \bar\theta\|.$$

Step 1 (Existence of ϵ-minimizers for an infimum). We claim that for every $\epsilon > 0$ there exists $\bar\theta_\epsilon \in \bar\Theta_S$ such that

$$\|\theta - \bar\theta_\epsilon\| \le r + \epsilon. \tag{81}$$

Indeed, by definition of r as the infimum of the set $\{\|\theta - \bar\theta\| : \bar\theta \in \bar\Theta_S\}$, the statement (81) is exactly the standard property of infima: if no such $\bar\theta_\epsilon$ existed, then $\|\theta - \bar\theta\| > r + \epsilon$ for all $\bar\theta \in \bar\Theta_S$, which would imply $\inf_{\bar\theta \in \bar\Theta_S} \|\theta - \bar\theta\| \ge r + \epsilon$, contradicting the definition of r.

Step 2 (Apply the modulus inequality). For each $\epsilon > 0$, apply (80) with $\theta' = \bar\theta_\epsilon$:

$$d_{\mathcal{F}}(P_\theta, P_{\bar\theta_\epsilon}) \le \omega_{\mathcal{F}}\big(\|\theta - \bar\theta_\epsilon\|\big).$$

Using (81) and monotonicity (nondecreasingness) of $\omega_{\mathcal{F}}$,

$$d_{\mathcal{F}}(P_\theta, P_{\bar\theta_\epsilon}) \le \omega_{\mathcal{F}}(r + \epsilon). \tag{82}$$

Step 3 (Take the infimum over $\bar{\theta}$ and then let $\epsilon \downarrow 0$). Since $\inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_\theta, P_{\bar{\theta}}) \le d_{\mathcal{F}}(P_\theta, P_{\bar{\theta}_\epsilon})$ for every $\epsilon > 0$, combining with (82) yields

$$\inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_\theta, P_{\bar{\theta}}) \le \omega_{\mathcal{F}}(r+\epsilon) \qquad \text{for all } \epsilon > 0.$$

Now take $\epsilon \downarrow 0$. If $\omega_{\mathcal{F}}$ is right-continuous, then $\omega_{\mathcal{F}}(r+\epsilon) \downarrow \omega_{\mathcal{F}}(r)$, and we obtain (79).

Remark (if right-continuity was not assumed). If $\omega_{\mathcal{F}}$ is only assumed nondecreasing, then the above argument yields

$$\inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_\theta, P_{\bar{\theta}}) \le \inf_{\epsilon>0} \omega_{\mathcal{F}}(r+\epsilon) =: \omega_{\mathcal{F}}(r+),$$

where $\omega_{\mathcal{F}}(r+)$ denotes the right-limit envelope. Equivalently, you may replace $\omega_{\mathcal{F}}$ everywhere by its right-continuous modification $\widetilde{\omega}_{\mathcal{F}}(r) := \inf_{\epsilon>0} \omega_{\mathcal{F}}(r+\epsilon)$, which is nondecreasing and right-continuous and satisfies the same modulus inequality (80). □

Theorem 7.8 (Population EM convergence in a general IPM on the slice)**.** *Assume Assumptions 7.2 and 7.5. Then for all $\theta_0 \in B_S$ and $\theta_{t+1} = T(\theta_t)$,*

$$\inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_{\theta_t}, P_{\bar{\theta}}) \;\le\; \omega_{\mathcal{F}}\Big(\gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S)\Big). \tag{83}$$

In particular, if $\omega_{\mathcal{F}}(r) \le L_{\mathcal{F}} r$ on $[0, r_0]$ and

$$\gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S) \le r_0 \qquad \textit{(e.g. it suffices that } \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S) \le r_0\textit{)},$$

then

$$\inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_{\theta_t}, P_{\bar{\theta}}) \;\le\; L_{\mathcal{F}}\, \gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S). \tag{84}$$

Proof. Fix $\theta_0 \in B_S$ and define the iterates $\theta_{t+1} = T(\theta_t)$ for $t \ge 0$.

Step 1 (Geometric contraction in parameter distance-to-set). By Assumption 7.2, T is a contraction *to the set* $\bar{\Theta}_S$: there exists $\gamma \in [0, 1)$ such that for all $\theta \in B_S$,

$$\operatorname{dist}_{\|\cdot\|}(T(\theta), \bar{\Theta}_S) \le \gamma \operatorname{dist}_{\|\cdot\|}(\theta, \bar{\Theta}_S).$$

Applying Theorem 7.4 (proved by induction from the preceding inequality), we obtain for every $t \ge 0$,

$$\operatorname{dist}_{\|\cdot\|}(\theta_t, \bar{\Theta}_S) \le \gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S). \tag{85}$$

Step 2 (Convert parameter distance-to-set into IPM distance-to-set). Apply Lemma 7.7 with $\theta = \theta_t$:

$$\inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_{\theta_t}, P_{\bar{\theta}}) \le \omega_{\mathcal{F}}\Big(\operatorname{dist}_{\|\cdot\|}(\theta_t, \bar{\Theta}_S)\Big).$$

Now substitute the bound (85). Since $\omega_{\mathcal{F}}$ is nondecreasing,

$$\omega_{\mathcal{F}}\Big(\operatorname{dist}_{\|\cdot\|}(\theta_t, \bar{\Theta}_S)\Big) \le \omega_{\mathcal{F}}\Big(\gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S)\Big),$$

which proves (83).

Step 3 (Lipschitz specialization). Assume $\omega_{\mathcal{F}}(r) \le L_{\mathcal{F}} r$ for $r \in [0, r_0]$ and that $\gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S) \le r_0$. Then the argument of $\omega_{\mathcal{F}}$ in (83) lies in $[0, r_0]$, so

$$\omega_{\mathcal{F}}\Big(\gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S)\Big) \le L_{\mathcal{F}}\, \gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S),$$

which yields (84). □

7.3. **Quotient formulation: convergence to $\bar{\Theta}/G$ in the quotient IPM.** We now rewrite the slice bound in an explicitly quotient form. Recall $\bar{d}_{\mathcal{F}}([\theta], [\theta']) = d_{\mathcal{F}}(P_\theta, P_{\theta'})$ and $\mathrm{dist}_{\mathcal{F}}([\theta], \bar{\Theta}/G) = \inf_{\bar{\theta}\in\bar{\Theta}} d_{\mathcal{F}}(P_\theta, P_{\bar{\theta}})$ (Section 4).

Assumption 7.9 (Slice is representative for the quotient basin)**.** Let $\mathcal{U} \subseteq \Theta/G$ be a quotient neighborhood and let S be a section as in Assumption 4.10. Assume $B_S = \pi(\mathcal{U}) \subseteq S$ is the set of slice representatives of $\mathcal{U}$, and that the population quotient EM correspondence admits a single-valued slice representative map $T : B_S \to B_S$ (constructed by selecting the slice representative of an EM update).

Theorem 7.10 (Population EM convergence on the quotient in $\bar{d}_{\mathcal{F}}$)**.** *Assume Assumptions 7.9, 7.2, and 7.5. Let $\xi_0 \in \mathcal{U}$ and define $\theta_0 = \pi(\xi_0) \in B_S$ and $\theta_{t+1} = T(\theta_t)$. Set $\xi_t := [\theta_t] \in \mathcal{U}$. Then*

$$\mathrm{dist}_{\mathcal{F}}(\xi_t, \bar{\Theta}/G) = \inf_{\bar{\theta}\in\bar{\Theta}} \bar{d}_{\mathcal{F}}(\xi_t, [\bar{\theta}]) = \inf_{\bar{\theta}\in\bar{\Theta}} d_{\mathcal{F}}(P_{\theta_t}, P_{\bar{\theta}}) \;\le\; \omega_{\mathcal{F}}\Big(\gamma^t \,\mathrm{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S)\Big). \tag{86}$$

In particular, if $\omega_{\mathcal{F}}(r) \le L_{\mathcal{F}} r$ locally (on a neighborhood containing the arguments above), then the quotient distance-to-set decays geometrically with factor γ.

Proof. **Definitions and basic facts used.**

(i) The quotient point corresponding to θ is its orbit $[\theta] := \{g \cdot \theta : \; g \in G\}$.

(ii) The slice representative map $\pi : \mathcal{U} \to B_S$ satisfies: for every $\xi \in \mathcal{U}$, $\pi(\xi) \in \xi \cap B_S$, and this intersection is a singleton; equivalently,

$$\forall \xi \in \mathcal{U}, \qquad \xi \cap B_S = \{\pi(\xi)\}. \tag{87}$$

(iii) The induced quotient IPM $\bar{d}_{\mathcal{F}}$ is defined on $\mathcal{U} \times \mathcal{U}$ by

$$\bar{d}_{\mathcal{F}}(\xi, \xi') := d_{\mathcal{F}}\big(P_{\pi(\xi)}, P_{\pi(\xi')}\big). \tag{88}$$

(iv) The quotient distance-to-set is

$$\mathrm{dist}_{\mathcal{F}}(\xi, \bar{\Theta}/G) := \inf_{\zeta\in\bar{\Theta}/G} \bar{d}_{\mathcal{F}}(\xi, \zeta),$$

where $\bar{\Theta}/G := \{[\bar{\theta}] : \bar{\theta} \in \bar{\Theta}\}$.

(v) (*Model invariance along orbits*) For any $g \in G$ and any θ, the induced distribution is constant on the orbit:

$$P_{g\cdot\theta} = P_\theta. \tag{89}$$

This is exactly the condition that makes $\bar{d}_{\mathcal{F}}$ well-defined on orbits; it is the structural input behind the quotient construction.

Fix $\xi_0 \in \mathcal{U}$ and define $\theta_0 = \pi(\xi_0) \in B_S$ and $\theta_{t+1} = T(\theta_t)$. By Assumption 7.9, the iterates satisfy $\theta_t \in B_S$ for all t and their orbits satisfy $\xi_t := [\theta_t] \in \mathcal{U}$ for all t (so $\pi(\xi_t)$ is well-defined).

Step 1 (Identify $\pi(\xi_t)$ with θ_t). Since $\xi_t = [\theta_t]$ and $\theta_t \in B_S$, we have $\theta_t \in \xi_t \cap B_S$. By the uniqueness property (87), $\xi_t \cap B_S = \{\pi(\xi_t)\}$, hence

$$\pi(\xi_t) = \theta_t. \tag{90}$$

Step 2 (Expand the quotient distance-to-set as an infimum over $\bar{\theta} \in \bar{\Theta}$). By definition of $\mathrm{dist}_{\mathcal{F}}$ and $\bar{\Theta}/G$,

$$\mathrm{dist}_{\mathcal{F}}(\xi_t, \bar{\Theta}/G) = \inf_{\zeta\in\bar{\Theta}/G} \bar{d}_{\mathcal{F}}(\xi_t, \zeta) = \inf_{\bar{\theta}\in\bar{\Theta}} \bar{d}_{\mathcal{F}}\big(\xi_t, [\bar{\theta}]\,\big). \tag{91}$$

(This is simply rewriting the infimum over the set of orbits $\{[\bar{\theta}] : \bar{\theta} \in \bar{\Theta}\}$ as an infimum over representatives.)

Step 3 (Relate $\bar{d}_{\mathcal{F}}(\xi_t, [\bar{\theta}])$ to $d_{\mathcal{F}}(P_{\theta_t}, P_{\bar{\theta}})$). Fix $\bar{\theta} \in \bar{\Theta}$. By (88) and (90),

$$\bar{d}_{\mathcal{F}}\big(\xi_t, [\bar{\theta}]\big) = d_{\mathcal{F}}\big(P_{\pi(\xi_t)}, P_{\pi([\bar{\theta}])}\big) = d_{\mathcal{F}}\big(P_{\theta_t}, P_{\pi([\bar{\theta}])}\big).$$

Now note that $\pi([\bar{\theta}]) \in [\bar{\theta}]$, so there exists some $g \in G$ with $\pi([\bar{\theta}]) = g \cdot \bar{\theta}$. By orbit-invariance (89), $P_{\pi([\bar{\theta}])} = P_{g\cdot\bar{\theta}} = P_{\bar{\theta}}$. Therefore

$$\bar{d}_{\mathcal{F}}\big(\xi_t, [\bar{\theta}]\big) = d_{\mathcal{F}}\big(P_{\theta_t}, P_{\bar{\theta}}\big). \tag{92}$$

Substituting (92) into (91) yields the claimed identities

$$\operatorname{dist}_{\mathcal{F}}(\xi_t, \bar{\Theta}/G) = \inf_{\bar{\theta}\in\bar{\Theta}} \bar{d}_{\mathcal{F}}(\xi_t, [\bar{\theta}]) = \inf_{\bar{\theta}\in\bar{\Theta}} d_{\mathcal{F}}(P_{\theta_t}, P_{\bar{\theta}}).$$

Step 4 (Reduce the infimum to slice representatives in $\bar{\Theta}_S$). We claim

$$\inf_{\bar{\theta}\in\bar{\Theta}} d_{\mathcal{F}}(P_{\theta_t}, P_{\bar{\theta}}) = \inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_{\theta_t}, P_{\bar{\theta}}), \qquad \bar{\Theta}_S := \bar{\Theta} \cap B_S. \tag{93}$$

The "$\leq$" direction is immediate since $\bar{\Theta}_S \subseteq \bar{\Theta}$. For "$\geq$": fix any $\bar{\theta} \in \bar{\Theta}$ such that its orbit lies in $\mathcal{U}$ (these are exactly the orbits relevant to $\bar{\Theta}/G$ as a subset of $\mathcal{U}$). By Lemma 4.9 (the G-invariance of $\bar{\Theta}$) and the slice property, the slice representative $\bar{\theta}_S := \pi([\bar{\theta}])$ lies in $\bar{\Theta} \cap B_S = \bar{\Theta}_S$ and belongs to the orbit of $\bar{\theta}$. Hence by (89), $P_{\bar{\theta}_S} = P_{\bar{\theta}}$, and therefore

$$d_{\mathcal{F}}(P_{\theta_t}, P_{\bar{\theta}}) = d_{\mathcal{F}}(P_{\theta_t}, P_{\bar{\theta}_S}) \geq \inf_{\tilde{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_{\theta_t}, P_{\tilde{\theta}}).$$

Taking the infimum over $\bar{\theta} \in \bar{\Theta}$ yields the "$\geq$" inequality in (93).

Step 5 (Apply the slice IPM convergence bound). By (93) and Theorem 7.8,

$$\inf_{\bar{\theta}\in\bar{\Theta}} d_{\mathcal{F}}(P_{\theta_t}, P_{\bar{\theta}}) = \inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_{\theta_t}, P_{\bar{\theta}}) \leq \omega_{\mathcal{F}}\Big(\gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S)\Big).$$

Combining this with Steps 2–3 gives (86).

Step 6 (Local Lipschitz specialization). If $\omega_{\mathcal{F}}(r) \leq L_{\mathcal{F}} r$ on an interval containing $r_t := \gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S)$ for the times of interest, then (86) implies

$$\operatorname{dist}_{\mathcal{F}}(\xi_t, \bar{\Theta}/G) \leq L_{\mathcal{F}}\, \gamma^t \operatorname{dist}_{\|\cdot\|}(\theta_0, \bar{\Theta}_S),$$

i.e. geometric decay with factor γ. □

Remark 7.11 (Two-sided moduli and orbit identification). Theorem 7.10 requires only the one-sided modulus (78), because it delivers an upper bound on IPM error. If one wishes to infer a rate in an orbit-distance such as d_Θ (Definition 4.13), then a two-sided modulus as in Assumption 4.14 is needed (Theorem 4.15).

Corollary 7.12 (Objective convergence along population EM)**.** *Under the hypotheses of Theorem 7.10, the sequence $\{\Phi(\theta_t)\}_{t\geq 0}$ is nondecreasing and bounded above by $\sup_{\theta\in\Theta} \Phi(\theta)$. In particular, $\Phi(\theta_t)$ converges as $t \to \infty$.*

Proof. **Step 1 (Monotonicity of $\Phi(\theta_t)$).** By Theorem 5.11 (population EM ascent), for every $\theta \in B_S$ we have

$$\Phi(T(\theta)) \geq \Phi(\theta).$$

Apply this with $\theta = \theta_t$ to obtain

$$\Phi(\theta_{t+1}) = \Phi(T(\theta_t)) \geq \Phi(\theta_t),$$

so $\{\Phi(\theta_t)\}_{t\geq 0}$ is nondecreasing.

Step 2 (Boundedness above). By definition of the supremum, for every $\theta \in \Theta$ we have $\Phi(\theta) \le \sup_{\vartheta\in\Theta} \Phi(\vartheta)$. In particular, for each t,

$$\Phi(\theta_t) \le \sup_{\vartheta\in\Theta} \Phi(\vartheta).$$

(If you prefer to avoid taking $\sup_\Theta$ globally, it suffices to assume Φ is bounded above on the forward-invariant region containing $\{\theta_t\}$; for instance, if the iterates are confined to a compact subset of Θ and Φ is continuous.)

Step 3 (A bounded monotone sequence converges). Let $a_t := \Phi(\theta_t)$. Since (a_t) is nondecreasing and bounded above, the limit exists and equals $\sup_{t\ge0} a_t$. For completeness: set $L := \sup_{t\ge0} a_t < \infty$. For any $\epsilon > 0$, by definition of supremum there exists T such that $L - \epsilon < a_T \le L$. Since (a_t) is nondecreasing, for all $t \ge T$ we have

$$L - \epsilon < a_T \le a_t \le L,$$

so $|a_t - L| < \epsilon$ for all $t \ge T$, i.e. $a_t \to L$.

Therefore $\Phi(\theta_t)$ converges as $t \to \infty$. □

8. Finite-sample theory: sample EM as a perturbed quotient contraction

This section develops nonasymptotic guarantees for *sample* EM by comparing it to *population* EM on a local basin. The logical structure mirrors the now-standard population-to-sample program for EM [1] but is formulated in a way that is compatible with (i) orbit geometry (quotient state space), (ii) misspecified, potentially set-valued targets, and (iii) a general IPM error criterion.

There are three components:

(1) a *deterministic* perturbation lemma for contractions (tight recursion with the unavoidable factor $(1-\gamma)^{-1}$);
(2) a *metric transfer* step converting slice-coordinate bounds into quotient/IPM distance-to-set bounds;
(3) an *operator deviation* analysis bounding $\sup_{\theta\in B_S} \|\hat{T}(\theta) - T(\theta)\|$ via empirical process complexity of an EM-induced class (generic chaining / entropy), following the sharp uniform-deviation theory in [10, 11].

8.1. Sample slice EM map and the operator deviation functional.

Fix a local section S and a slice basin $B_S \subseteq S$ as in Sections 4–7. Assume the population slice EM map $T : B_S \to B_S$ is single-valued (Definition 6.10).

Assumption 8.1 (Uniform strong concavity for population and sample M-steps)**.** There exists $\lambda > 0$ such that:

(i) For every $\theta \in B_S$, the population surrogate $\theta' \mapsto Q(\theta';\theta)$ is twice continuously differentiable on B_S and

$$-\nabla^2_{\theta'\theta'} Q(\theta';\theta) \;\succeq\; \lambda I \qquad \forall \theta,\theta' \in B_S.$$

(ii) For the same λ (or a smaller constant), for every $\theta \in B_S$ the sample surrogate $\theta' \mapsto \hat{Q}(\theta';\theta)$ is twice continuously differentiable on B_S and

$$-\nabla^2_{\theta'\theta'} \hat{Q}(\theta';\theta) \;\succeq\; \lambda I \qquad \forall \theta,\theta' \in B_S,$$

on the event under consideration (typically a high-probability event, quantified later).

Definition 8.2 (Sample slice EM map)**.** Under Assumption 8.1, define the (single-valued) sample slice EM map $\hat{T} : B_S \to B_S$ by

$$\hat{T}(\theta) := \arg\max_{\theta' \in B_S} \hat{Q}(\theta'; \theta).$$

We measure sample-to-population discrepancy through the operator deviation functional

$$\Delta_n := \sup_{\theta \in B_S} \|\hat{T}(\theta) - T(\theta)\|. \tag{94}$$

Theorems 8.4–8.12 are deterministic statements conditional on a bound $\Delta_n \le \varepsilon$; the probabilistic work is to bound Δ_n with high probability in Subsection 8.5.

8.2. **Deterministic perturbation theory for contractions (tight).** We begin with the basic deterministic inequality controlling the iterates of a perturbed contraction. This lemma is elementary but central: it yields the sharp amplification factor $(1-\gamma)^{-1}$.

Assumption 8.3 (Population contraction on the slice basin)**.** There exists $\gamma \in (0,1)$ such that

$$\|T(\theta) - T(\theta')\| \le \gamma \|\theta - \theta'\| \qquad \forall \theta, \theta' \in B_S. \tag{95}$$

Theorem 8.4 (Perturbed contraction recursion (sharp, deterministic))**.** *Assume Assumption 8.3. Let $(\theta_t)_{t\ge 0}$ and $(\hat{\theta}_t)_{t\ge 0}$ be sequences in B_S satisfying*

$$\theta_{t+1} = T(\theta_t), \qquad \hat{\theta}_{t+1} = \hat{T}(\hat{\theta}_t),$$

and suppose Δ_n in (94) is finite. Then for every $t \ge 0$,

$$\|\hat{\theta}_t - \theta_t\| \le \gamma^t \|\hat{\theta}_0 - \theta_0\| + \frac{1-\gamma^t}{1-\gamma}\,\Delta_n. \tag{96}$$

In particular, if $\hat{\theta}_0 = \theta_0$, then

$$\|\hat{\theta}_t - \theta_t\| \le \frac{1-\gamma^t}{1-\gamma}\,\Delta_n \le \frac{\Delta_n}{1-\gamma}. \tag{97}$$

Proof. **Definitions being used.** Assumption 8.3 provides:

(i) (*T is a contraction on B_S.*) There exists $\gamma \in [0,1)$ such that for all $\theta, \theta' \in B_S$,

$$\|T(\theta) - T(\theta')\| \le \gamma \|\theta - \theta'\|. \tag{98}$$

(ii) (*Uniform perturbation size.*) The quantity Δ_n is defined by

$$\Delta_n := \sup_{\theta \in B_S} \|\hat{T}(\theta) - T(\theta)\|, \tag{99}$$

and is finite by hypothesis. In particular, for every $\theta \in B_S$,

$$\|\hat{T}(\theta) - T(\theta)\| \le \Delta_n. \tag{100}$$

(iii) (*Iterates stay in the domain.*) By assumption of the theorem statement, $\theta_t, \hat{\theta}_t \in B_S$ for all t, so (98) and (100) can be applied at each time.

Define the error sequence

$$e_t := \|\hat{\theta}_t - \theta_t\| \qquad (t \ge 0),$$

which is a nonnegative real sequence.

Step 1 (Derive the fundamental one-step inequality). We start from the definitions of the iterates:

$$\hat{\theta}_{t+1} - \theta_{t+1} = \hat{T}(\hat{\theta}_t) - T(\theta_t).$$

Add and subtract $T(\hat{\theta}_t)$:

$$\hat{T}(\hat{\theta}_t) - T(\theta_t) = \big(\hat{T}(\hat{\theta}_t) - T(\hat{\theta}_t)\big) + \big(T(\hat{\theta}_t) - T(\theta_t)\big).$$

Now apply the triangle inequality $\|a+b\| \le \|a\| + \|b\|$:

$$\|\hat{\theta}_{t+1} - \theta_{t+1}\| \le \|\hat{T}(\hat{\theta}_t) - T(\hat{\theta}_t)\| + \|T(\hat{\theta}_t) - T(\theta_t)\|. \tag{101}$$

Bound each term:

- By (100) applied at $\theta = \hat{\theta}_t \in B_S$,
$$\|\hat{T}(\hat{\theta}_t) - T(\hat{\theta}_t)\| \le \Delta_n.$$
- By the contraction property (98) applied to $(\hat{\theta}_t, \theta_t) \in B_S \times B_S$,
$$\|T(\hat{\theta}_t) - T(\theta_t)\| \le \gamma \|\hat{\theta}_t - \theta_t\| = \gamma e_t.$$

Substituting into (101) gives the one-step recursion

$$e_{t+1} \le \Delta_n + \gamma e_t \qquad (t \ge 0). \tag{102}$$

Step 2 (Unroll the recursion by repeated substitution). We now expand (102) explicitly. First apply it once:

$$e_1 \le \Delta_n + \gamma e_0.$$

Apply it again to e_2 and substitute the bound for e_1:

$$e_2 \le \Delta_n + \gamma e_1 \le \Delta_n + \gamma(\Delta_n + \gamma e_0) = \Delta_n(1+\gamma) + \gamma^2 e_0.$$

One more step:

$$e_3 \le \Delta_n + \gamma e_2 \le \Delta_n + \gamma\big(\Delta_n(1+\gamma) + \gamma^2 e_0\big) = \Delta_n(1+\gamma+\gamma^2) + \gamma^3 e_0.$$

The pattern is now clear and can be proved by induction:

Claim 8.5. *For every integer $t \ge 1$,*

$$e_t \le \gamma^t e_0 + \Delta_n \sum_{k=0}^{t-1} \gamma^k. \tag{103}$$

Proof of Claim 8.5. We argue by induction on t. For $t = 1$, (103) reads $e_1 \le \gamma e_0 + \Delta_n$, which is exactly (102) at $t = 0$. Assume (103) holds for some $t \ge 1$. Then by (102),

$$e_{t+1} \le \Delta_n + \gamma e_t \le \Delta_n + \gamma\Big(\gamma^t e_0 + \Delta_n \sum_{k=0}^{t-1} \gamma^k\Big) = \gamma^{t+1} e_0 + \Delta_n \sum_{k=0}^{t} \gamma^k,$$

which is (103) with t replaced by $t+1$. This completes the induction. □

Thus (103) holds for all $t \ge 1$, and it is also true for $t = 0$ (with the empty sum convention $\sum_{k=0}^{-1}(\cdot) = 0$). Therefore, for all $t \ge 0$,

$$e_t \le \gamma^t e_0 + \Delta_n \sum_{k=0}^{t-1} \gamma^k. \tag{104}$$

Step 3 (Evaluate the geometric series and obtain the closed form). Since $\gamma \in [0, 1)$, the finite geometric sum satisfies

$$\sum_{k=0}^{t-1} \gamma^k = \frac{1-\gamma^t}{1-\gamma} \qquad (t \ge 1), \tag{105}$$

and for $t = 0$ the sum is 0, which is consistent with the right-hand side since $(1-\gamma^0)/(1-\gamma) = 0$. Substitute (105) into (104) and recall $e_t = \|\hat{\theta}_t - \theta_t\|$ and $e_0 = \|\hat{\theta}_0 - \theta_0\|$ to get

$$\|\hat{\theta}_t - \theta_t\| \le \gamma^t \|\hat{\theta}_0 - \theta_0\| + \frac{1-\gamma^t}{1-\gamma}\Delta_n,$$

which is exactly (96).

Step 4 (Same-initialization specialization). If $\hat{\theta}_0 = \theta_0$, then the first term vanishes and (96) becomes

$$\|\hat{\theta}_t - \theta_t\| \le \frac{1-\gamma^t}{1-\gamma}\Delta_n.$$

Since $0 \le \gamma^t \le 1$, we have $1 - \gamma^t \le 1$, hence

$$\frac{1-\gamma^t}{1-\gamma}\Delta_n \le \frac{1}{1-\gamma}\Delta_n,$$

yielding (97). □

Remark 8.6 (Why $(1-\gamma)^{-1}$ is unavoidable (sharpness)). The factor $(1-\gamma)^{-1}$ cannot be improved under a *uniform additive perturbation* model. Indeed, even in one dimension, consider the deterministic recursion

$$e_{t+1} = \gamma e_t + \Delta_n, \qquad e_0 = 0,$$

which is the equality case of (102). Solving it exactly by the same unrolling gives

$$e_t = \Delta_n \sum_{k=0}^{t-1} \gamma^k = \Delta_n \frac{1-\gamma^t}{1-\gamma},$$

so the upper bound in (97) is attained. Thus any deterministic bound that depends only on (γ, Δ_n) must, in general, contain a factor of order $(1-\gamma)^{-1}$.

8.3. **Distance-to-set version (misspecified targets).** Under misspecification the target is set-valued, so we need a distance-to-set analogue of the previous theorem. We state it for an arbitrary closed set $A \subseteq B_S$ (later $A = \bar{\Theta}_S$).

Definition 8.7 (One-step set deviation)**.** For a set $A \subseteq B_S$, define

$$\Delta_n(A) := \sup_{\theta \in B_S} \operatorname{dist}\big(\hat{T}(\theta), T(\theta)\big), \qquad \operatorname{dist}(u, v) := \|u - v\|.$$

Note $\Delta_n(A)$ is the same as Δ_n in (94) (the notation emphasizes the role of sets below).

Assumption 8.8 (Contraction to a set for population map)**.** There exist $\gamma \in (0, 1)$ and a nonempty set $A \subseteq B_S$ such that

$$\operatorname{dist}\big(T(\theta), A\big) \;\le\; \gamma \operatorname{dist}(\theta, A) \qquad \forall \theta \in B_S. \tag{106}$$

Theorem 8.9 (Perturbed contraction to a set)**.** *Assume Assumption 8.8. Let $\hat{\theta}_{t+1} = \hat{T}(\hat{\theta}_t)$ with $\hat{\theta}_t \in B_S$ for all t. Then for all $t \ge 0$,*

$$\operatorname{dist}(\hat{\theta}_t, A) \;\le\; \gamma^t \operatorname{dist}(\hat{\theta}_0, A) + \frac{1-\gamma^t}{1-\gamma}\,\Delta_n. \tag{107}$$

Proof. **Definitions used.** For a nonempty set $A \subseteq \mathbb{R}^d$ (in the ambient norm $\|\cdot\|$ fixed throughout this section),

$$\operatorname{dist}(\theta, A) := \inf_{a \in A} \|\theta - a\|.$$

Assumption 8.8 provides the following two ingredients:

(i) (*Contraction to the set A for the population map T.*) There exists $\gamma \in [0,1)$ such that for all $\theta \in B_S$,

$$\operatorname{dist}(T(\theta), A) \le \gamma \operatorname{dist}(\theta, A). \tag{108}$$

(ii) (*Uniform perturbation size.*) With $\Delta_n := \sup_{\theta \in B_S} \|\hat{T}(\theta) - T(\theta)\| < \infty$, we have

$$\|\hat{T}(\theta) - T(\theta)\| \le \Delta_n \qquad \text{for all } \theta \in B_S. \tag{109}$$

We will also use the following basic inequality for distance-to-set.

Lemma (distance-to-set is 1-Lipschitz). For any nonempty set $A \subseteq \mathbb{R}^d$ and any $x, y \in \mathbb{R}^d$,

$$\big|\operatorname{dist}(x, A) - \operatorname{dist}(y, A)\big| \le \|x - y\|. \tag{110}$$

In particular,

$$\operatorname{dist}(x, A) \le \operatorname{dist}(y, A) + \|x - y\|. \tag{111}$$

Proof. Fix $x, y \in \mathbb{R}^d$. For any $a \in A$, the triangle inequality gives

$$\|x - a\| \le \|x - y\| + \|y - a\|.$$

Taking the infimum over $a \in A$ on the right-hand side yields

$$\|x - a\| \le \|x - y\| + \operatorname{dist}(y, A) \qquad \text{for all } a \in A.$$

Now take the infimum over $a \in A$ on the left-hand side to obtain

$$\operatorname{dist}(x, A) \le \|x - y\| + \operatorname{dist}(y, A).$$

Swapping the roles of x and y gives $\operatorname{dist}(y, A) \le \|x - y\| + \operatorname{dist}(x, A)$. Combining the two inequalities yields (110), and (111) is the first inequality written explicitly. □

Step 1 (One-step recursion for the distance-to-set). Set

$$e_t := \operatorname{dist}(\hat{\theta}_t, A) \qquad (t \ge 0),$$

so $e_t \ge 0$ is a real sequence. Using $\hat{\theta}_{t+1} = \hat{T}(\hat{\theta}_t)$, we start with

$$e_{t+1} = \operatorname{dist}(\hat{\theta}_{t+1}, A) = \operatorname{dist}(\hat{T}(\hat{\theta}_t), A).$$

Apply (111) with $x = \hat{T}(\hat{\theta}_t)$ and $y = T(\hat{\theta}_t)$:

$$\operatorname{dist}(\hat{T}(\hat{\theta}_t), A) \le \operatorname{dist}(T(\hat{\theta}_t), A) + \|\hat{T}(\hat{\theta}_t) - T(\hat{\theta}_t)\|. \tag{112}$$

We bound the two terms in (112) using the assumptions:

- By (109) (since $\hat{\theta}_t \in B_S$),
$$\|\hat{T}(\hat{\theta}_t) - T(\hat{\theta}_t)\| \le \Delta_n.$$
- By the contraction-to-set property (76) applied at $\theta = \hat{\theta}_t \in B_S$,
$$\operatorname{dist}(T(\hat{\theta}_t), A) \le \gamma \operatorname{dist}(\hat{\theta}_t, A) = \gamma e_t.$$

Substituting these into (112) yields the one-step recursion

$$e_{t+1} \le \Delta_n + \gamma e_t \qquad (t \ge 0). \tag{113}$$

Step 2 (Unroll the scalar recursion). We now unroll (113) exactly as for the pointwise perturbation theorem. For completeness, we state and prove the unrolling bound.

Lemma 8.10 (Unrolling $e_{t+1} \le \gamma e_t + \Delta$)**.** *Let $(e_t)_{t\ge 0}$ be a nonnegative sequence satisfying $e_{t+1} \le \gamma e_t + \Delta$ for all $t \ge 0$, where $\gamma \in [0,1)$ and $\Delta \ge 0$. Then for all $t \ge 0$,*

$$e_t \le \gamma^t e_0 + \Delta \sum_{k=0}^{t-1} \gamma^k = \gamma^t e_0 + \Delta \frac{1-\gamma^t}{1-\gamma}.$$

Proof. We prove by induction that for all $t \ge 0$,

$$e_t \le \gamma^t e_0 + \Delta \sum_{k=0}^{t-1} \gamma^k,$$

with the convention that $\sum_{k=0}^{-1} \gamma^k := 0$. For $t = 0$ this reads $e_0 \le e_0$, which is true. Assume the bound holds at time t. Then

$$e_{t+1} \le \gamma e_t + \Delta \le \gamma\Big(\gamma^t e_0 + \Delta \sum_{k=0}^{t-1} \gamma^k\Big) + \Delta = \gamma^{t+1} e_0 + \Delta \sum_{k=0}^{t} \gamma^k,$$

which is exactly the inductive claim for $t+1$. This completes the induction. Evaluating the geometric series gives the closed form. □

Apply Lemma 8.10 to (113) with $e_t = \operatorname{dist}(\hat\theta_t, A)$ and $\Delta = \Delta_n$:

$$\operatorname{dist}(\hat\theta_t, A) \le \gamma^t \operatorname{dist}(\hat\theta_0, A) + \Delta_n \frac{1-\gamma^t}{1-\gamma},$$

which is (107). □

8.4. **Transfer to IPM error and to the quotient.** We now convert the slice distance-to-set bound into the distributional criterion $\operatorname{dist}_{\mathcal F}(\left[\hat\theta_t\right], \bar\Theta/G)$. As in Section 7, only a one-sided modulus is needed.

Assumption 8.11 (IPM modulus on the slice basin)**.** There exists a nondecreasing $\omega_{\mathcal F}$ with $\omega_{\mathcal F}(0) = 0$ such that

$$d_{\mathcal F}(P_\theta, P_{\theta'}) \;\le\; \omega_{\mathcal F}(\|\theta - \theta'\|) \qquad \forall \theta, \theta' \in B_S.$$

Theorem 8.12 (Sample EM: quotient IPM distance-to-set bound)**.** *Let $A = \bar\Theta_S := \bar\Theta \cap B_S$. Assume Assumption 8.8 holds with this A, and assume Assumption 8.11. Let $\hat\theta_{t+1} = \hat T(\hat\theta_t)$ with $\hat\theta_t \in B_S$ and define $\hat\xi_t := \left[\hat\theta_t\right] \in \Theta/G$. Then for all $t \ge 0$,*

$$\operatorname{dist}_{\mathcal F}\left(\hat\xi_t, \bar\Theta/G\right) = \inf_{\bar\theta \in \bar\Theta} d_{\mathcal F}(P_{\hat\theta_t}, P_{\bar\theta}) \;\le\; \omega_{\mathcal F}\left(\gamma^t \operatorname{dist}(\hat\theta_0, \bar\Theta_S) + \frac{1-\gamma^t}{1-\gamma}\,\Delta_n\right). \tag{114}$$

If $\omega_{\mathcal F}(r) \le L_{\mathcal F} r$ on the relevant range, then

$$\operatorname{dist}_{\mathcal F}\left(\hat\xi_t, \bar\Theta/G\right) \;\le\; L_{\mathcal F}\left(\gamma^t \operatorname{dist}(\hat\theta_0, \bar\Theta_S) + \frac{1-\gamma^t}{1-\gamma}\,\Delta_n\right). \tag{115}$$

Proof. **Definitions and assumptions used.**

(i) (*Distance-to-set in parameter space.*) For $C \subseteq \mathbb{R}^d$ nonempty,

$$\operatorname{dist}(\theta, C) := \inf_{c \in C} \|\theta - c\|.$$

(ii) (*Quotient distance-to-set in IPM.*) For $\xi \in \Theta/G$,

$$\operatorname{dist}_{\mathcal{F}}(\xi, \bar{\Theta}/G) := \inf_{\bar{\theta}\in\bar{\Theta}} \bar{d}_{\mathcal{F}}(\xi, [\bar{\theta}]), \qquad \bar{d}_{\mathcal{F}}([\theta], [\theta']) := d_{\mathcal{F}}(P_\theta, P_{\theta'}).$$

In particular, for $\hat{\xi}_t = \left[\hat{\theta}_t\right]$,

$$\operatorname{dist}_{\mathcal{F}}(\hat{\xi}_t, \bar{\Theta}/G) = \inf_{\bar{\theta}\in\bar{\Theta}} d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}}), \tag{116}$$

using the orbit-invariance property $P_{g\cdot\theta} = P_\theta$ that makes $\bar{d}_{\mathcal{F}}$ well-defined.

(iii) (*G-invariance of* $\bar{\Theta}$.) Lemma 4.9 states $g \cdot \bar{\Theta} = \bar{\Theta}$ for all $g \in G$. Combined with orbit-invariance of the model, this implies that replacing $\bar{\theta}$ by any orbit-equivalent representative does not change $P_{\bar{\theta}}$.

(iv) (*Finite-sample IPM modulus.*) Assumption 8.11 provides a nondecreasing $\omega_{\mathcal{F}} : [0,\infty) \to [0,\infty)$ such that for all $\theta, \theta' \in B_S$,

$$d_{\mathcal{F}}(P_\theta, P_{\theta'}) \le \omega_{\mathcal{F}}(\|\theta - \theta'\|). \tag{117}$$

(As in Lemma 7.7, if right-continuity is not assumed, one can replace $\omega_{\mathcal{F}}$ by its right-continuous envelope without changing (117).)

(v) (*Perturbed contraction to the set* $A = \bar{\Theta}_S$.) Assumption 8.8 and the definition $\Delta_n := \sup_{\theta\in B_S} \|\hat{T}(\theta) - T(\theta)\| < \infty$ imply (Theorem 8.9) that for all $t \ge 0$,

$$\operatorname{dist}(\hat{\theta}_t, \bar{\Theta}_S) \le \gamma^t \operatorname{dist}(\hat{\theta}_0, \bar{\Theta}_S) + \frac{1-\gamma^t}{1-\gamma}\,\Delta_n. \tag{118}$$

Fix $t \ge 0$.

Step 1 (Reduce the quotient infimum to slice representatives in $\bar{\Theta}_S$). We claim that

$$\inf_{\bar{\theta}\in\bar{\Theta}} d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}}) = \inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}}), \qquad \bar{\Theta}_S := \bar{\Theta} \cap B_S. \tag{119}$$

The "$\le$" direction holds because $\bar{\Theta}_S \subseteq \bar{\Theta}$. For "$\ge$": fix any $\bar{\theta} \in \bar{\Theta}$ whose orbit is relevant for $\bar{\Theta}/G$ inside the local quotient chart. Let $\bar{\theta}_S := \pi([\bar{\theta}])$. By Lemma 4.9, $\bar{\theta}_S \in \bar{\Theta}$, and by definition of π, $\bar{\theta}_S \in B_S$, hence $\bar{\theta}_S \in \bar{\Theta}_S$. Moreover $\bar{\theta}_S$ lies in the orbit of $\bar{\theta}$, so by orbit-invariance of the model we have $P_{\bar{\theta}_S} = P_{\bar{\theta}}$. Therefore

$$d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}}) = d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}_S}) \ge \inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}}).$$

Taking the infimum over $\bar{\theta} \in \bar{\Theta}$ yields the "$\ge$" inequality in (119), establishing (119). Combining with (116) gives

$$\operatorname{dist}_{\mathcal{F}}(\hat{\xi}_t, \bar{\Theta}/G) = \inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}}). \tag{120}$$

Step 2 (Control IPM distance-to-set by parameter distance-to-set). We now prove the analogue of Lemma 7.7 for the present $\hat{\theta}_t$. Let $r_t := \operatorname{dist}(\hat{\theta}_t, \bar{\Theta}_S)$. By definition of infimum, for every $\epsilon > 0$ there exists $\bar{\theta}_{t,\epsilon} \in \bar{\Theta}_S$ such that

$$\|\hat{\theta}_t - \bar{\theta}_{t,\epsilon}\| \le r_t + \epsilon. \tag{121}$$

Apply the modulus inequality (117) with $\theta = \hat{\theta}_t$ and $\theta' = \bar{\theta}_{t,\epsilon}$:

$$d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}_{t,\epsilon}}) \le \omega_{\mathcal{F}}\big(\|\hat{\theta}_t - \bar{\theta}_{t,\epsilon}\|\big) \le \omega_{\mathcal{F}}(r_t + \epsilon),$$

using (121) and monotonicity of $\omega_{\mathcal{F}}$. Since $\inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}}) \le d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}_{t,\epsilon}})$, we obtain

$$\inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}}) \le \omega_{\mathcal{F}}(r_t + \epsilon) \qquad \text{for all } \epsilon > 0.$$

Letting $\epsilon \downarrow 0$ and using right-continuity (or the right-continuous envelope) yields

$$\inf_{\bar{\theta}\in\bar{\Theta}_S} d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}}) \le \omega_{\mathcal{F}}\big(\operatorname{dist}(\hat{\theta}_t, \bar{\Theta}_S)\big). \tag{122}$$

Combining (120) and (122) gives

$$\operatorname{dist}_{\mathcal{F}}(\hat{\xi}_t, \bar{\Theta}/G) \le \omega_{\mathcal{F}}\big(\operatorname{dist}(\hat{\theta}_t, \bar{\Theta}_S)\big). \tag{123}$$

Step 3 (Insert the perturbed set-contraction bound and conclude). Apply (118) to bound $\operatorname{dist}(\hat{\theta}_t, \bar{\Theta}_S)$, and then apply monotonicity of $\omega_{\mathcal{F}}$:

$$\begin{aligned}\operatorname{dist}_{\mathcal{F}}(\hat{\xi}_t, \bar{\Theta}/G) &\le \omega_{\mathcal{F}}\big(\operatorname{dist}(\hat{\theta}_t, \bar{\Theta}_S)\big) \\ &\le \omega_{\mathcal{F}}\left(\gamma^t \operatorname{dist}(\hat{\theta}_0, \bar{\Theta}_S) + \frac{1-\gamma^t}{1-\gamma}\,\Delta_n\right).\end{aligned}$$

This is exactly (114).

Step 4 (Lipschitz specialization). If $\omega_{\mathcal{F}}(r) \le L_{\mathcal{F}} r$ on an interval containing the argument

$$r_t := \gamma^t \operatorname{dist}(\hat{\theta}_0, \bar{\Theta}_S) + \frac{1-\gamma^t}{1-\gamma}\Delta_n,$$

then applying this bound to (114) yields (115). □

8.5. **Bounding Δ_n via EM-induced empirical processes.** The deterministic bounds above reduce the finite-sample problem to controlling Δ_n. We now relate Δ_n to uniform deviations of (population vs. empirical) gradients of the surrogate.

8.5.1. *Gradient characterization and a stability inequality.* Define the surrogate gradients (in the M-step variable):

$$G(\theta', \theta) := \nabla_{\theta'} Q(\theta'; \theta), \qquad \hat{G}(\theta', \theta) := \nabla_{\theta'} \hat{Q}(\theta'; \theta). \tag{124}$$

Under Assumption 8.1, the first-order conditions characterize the updates:

$$G(T(\theta), \theta) = 0, \qquad \hat{G}(\hat{T}(\theta), \theta) = 0 \qquad (\theta \in B_S). \tag{125}$$

Lemma 8.13 (Strong concavity yields a Lipschitz inverse-gradient bound)**.** *Assume Assumption 8.1(i) for Q. Then for each fixed $\theta \in B_S$, the map $\theta' \mapsto G(\theta', \theta)$ is λ-strongly monotone: for all $\theta'_1, \theta'_2 \in B_S$,*

$$\big\langle G(\theta'_1, \theta) - G(\theta'_2, \theta),\, \theta'_1 - \theta'_2 \big\rangle \;\ge\; \lambda \|\theta'_1 - \theta'_2\|^2. \tag{126}$$

Consequently,

$$\|\theta'_1 - \theta'_2\| \;\le\; \frac{1}{\lambda}\, \|G(\theta'_1, \theta) - G(\theta'_2, \theta)\|. \tag{127}$$

The same statements hold for $\hat{G}$ under Assumption 8.1(ii).

Proof. Fix $\theta \in B_S$.

Step 1 (From strong concavity to strong monotonicity). Assumption 8.1(i) states $-\nabla^2_{\theta'\theta'} Q(\theta';\theta) \succeq \lambda I$. Equivalently, $\nabla^2_{\theta'\theta'} Q(\theta';\theta) \preceq -\lambda I$. For a C^1 map $G(\cdot,\theta) = \nabla_{\theta'} Q(\cdot;\theta)$, this implies strong monotonicity with constant λ: integrate the Hessian along the segment between θ'_1 and θ'_2 to obtain (126) (the standard characterization of strong concavity/monotonicity).

Step 2 (Derive the inverse-Lipschitz bound). Apply Cauchy–Schwarz to (126):

$$\lambda\|\theta'_1 - \theta'_2\|^2 \le \langle G(\theta'_1,\theta) - G(\theta'_2,\theta), \theta'_1 - \theta'_2\rangle \le \|G(\theta'_1,\theta) - G(\theta'_2,\theta)\|\,\|\theta'_1 - \theta'_2\|.$$

If $\theta'_1 \neq \theta'_2$, divide by $\|\theta'_1 - \theta'_2\|$; otherwise the inequality is trivial. This yields (127). □

Theorem 8.14 (Operator deviation controlled by gradient deviations)**.** *Assume Assumption 8.1 and that $T(\theta), \hat{T}(\theta) \in B_S$ for all $\theta \in B_S$. Then*

$$\Delta_n \;\le\; \frac{2}{\lambda} \sup_{\theta\in B_S} \sup_{\theta'\in B_S} \|\hat{G}(\theta',\theta) - G(\theta',\theta)\|. \tag{128}$$

Proof. Fix $\theta \in B_S$.

Step 1 (Add and subtract $G(\hat{T}(\theta),\theta)$). Using (125) and the triangle inequality,

$$\|G(\hat{T}(\theta),\theta) - G(T(\theta),\theta)\| = \|G(\hat{T}(\theta),\theta) - 0\| \le \|G(\hat{T}(\theta),\theta) - \hat{G}(\hat{T}(\theta),\theta)\| + \|\hat{G}(\hat{T}(\theta),\theta)\|.$$

But $\hat{G}(\hat{T}(\theta),\theta) = 0$ by (125), so

$$\|G(\hat{T}(\theta),\theta) - G(T(\theta),\theta)\| \le \|G(\hat{T}(\theta),\theta) - \hat{G}(\hat{T}(\theta),\theta)\|.$$

Similarly, adding and subtracting $\hat{G}(T(\theta),\theta)$ yields

$$\|G(\hat{T}(\theta),\theta) - G(T(\theta),\theta)\| = \|\hat{G}(T(\theta),\theta) - G(T(\theta),\theta)\|.$$

Combining the two displays and taking the maximum,

$$\|G(\hat{T}(\theta),\theta) - G(T(\theta),\theta)\| \le \sup_{\theta'\in B_S} \|\hat{G}(\theta',\theta) - G(\theta',\theta)\|.$$

Step 2 (Invert the strong monotonicity). By Lemma 8.13 applied to $G(\cdot,\theta)$,

$$\|\hat{T}(\theta) - T(\theta)\| \le \frac{1}{\lambda}\|G(\hat{T}(\theta),\theta) - G(T(\theta),\theta)\| \le \frac{1}{\lambda} \sup_{\theta'\in B_S} \|\hat{G}(\theta',\theta) - G(\theta',\theta)\|.$$

Step 3 (Symmetrize to obtain the factor $2/\lambda$). The previous bound is already sufficient. A slightly more conservative but uniform form is obtained by noting that

$$\sup_{\theta'\in B_S} \|\hat{G}(\theta',\theta) - G(\theta',\theta)\| \le \sup_{\theta\in B_S} \sup_{\theta'\in B_S} \|\hat{G}(\theta',\theta) - G(\theta',\theta)\|,$$

and then taking $\sup_{\theta\in B_S}$ of the left-hand side. Writing the right-hand side with a prefactor $2/\lambda$ accommodates the common situation where one bounds separately $\sup\|\hat{G} - G\|$ and an additional event ensuring $\hat{T}(\theta) \in B_S$ uniformly. (Keeping $2/\lambda$ is harmless and matches standard perturbation statements in the EM literature.) Thus (128) holds. □

8.5.2. *EM-induced function class and empirical process bounds.* We now express $\hat{G} - G$ as an empirical process indexed by an EM-induced class.

Under the standing differentiability conditions (as in Assumptions 5.1 and 6.4), the gradient takes the form

$$G(\theta', \theta) = \mathbb{E}_{X \sim P^*}\Big[g_{\theta',\theta}(X)\Big], \qquad \hat{G}(\theta', \theta) = \mathbb{E}_{X \sim P_n}\Big[g_{\theta',\theta}(X)\Big],$$

where the *index function* $g_{\theta',\theta} : \mathcal{X} \to \mathbb{R}^d$ is

$$g_{\theta',\theta}(x) := \mathbb{E}_{Z \sim p_\theta(\cdot|x)}\big[\nabla_{\theta'} \log p_{\theta'}(x, Z)\big]. \tag{129}$$

Define the EM-induced vector-valued class

$$\mathcal{H} := \{\, g_{\theta',\theta} : \ (\theta', \theta) \in B_S \times B_S \,\}. \tag{130}$$

To apply scalar empirical process results, we use the standard reduction: for $u \in \mathbb{S}^{d-1}$, consider the scalar class $\mathcal{H}_u := \{\langle u, g\rangle : g \in \mathcal{H}\}$ and take a union bound over a net on the sphere. We state the cleanest form in terms of a generic complexity functional; see [10, 11].

Theorem 8.15 (Generic chaining bound for the EM-induced class)**.** *Assume the functions in $\mathcal{H}$ are square-integrable under P^* and that $\mathcal{H}$ is P^*-centered (or replace each h by $h - \mathbb{E}_{P^*}h$). Write the $L_2(P^*)$ semimetric as*

$$d(h, h') := \|h - h'\|_{L_2(P^*)}.$$

(Standard additional concentration hypothesis.) *Assume moreover that the class has $L_2(P^*)$-subgaussian increments in the following sense: there exists $K < \infty$ such that for all $h, h' \in \mathcal{H}$, the centered random variable*

$$Z_{h,h'}(X) := (h - h')(X) - \mathbb{E}_{P^*}[(h - h')(X)]$$

satisfies the ψ_2 bound

$$\|Z_{h,h'}(X)\|_{\psi_2} \ \le \ K\, d(h, h'). \tag{131}$$

(For example, (131) holds if every $h(X)$ is subgaussian with $\|h(X) - \mathbb{E}h(X)\|_{\psi_2} \le K\|h\|_{L_2(P^*)}$.)

Then there exists a universal constant $C > 0$ such that for every $\delta \in (0, 1)$, with probability at least $1 - \delta$,

$$\sup_{h \in \mathcal{H}} \big|(P_n - P^*)h\big| \ \le \ CK\left(\frac{\gamma_2(\mathcal{H}, d)}{\sqrt{n}} + \operatorname{diam}(\mathcal{H}, d)\sqrt{\frac{\log(1/\delta)}{n}}\right), \tag{132}$$

where $\gamma_2(\mathcal{H}, d)$ is Talagrand's γ_2 functional of $(\mathcal{H}, d)$ and $\operatorname{diam}(\mathcal{H}, d) := \sup_{h,h' \in \mathcal{H}} d(h, h')$.

Remark (on the assumptions). *Without some concentration control like (131), one generally cannot expect an exponential tail $\sqrt{\log(1/\delta)/n}$ from only L_2 integrability; the high-probability form (132) is the standard subgaussian-increment regime in generic chaining references.*

Proof. **Step 0 (Notation and γ_2 definition).** Let $X_1, \dots, X_n \overset{\text{i.i.d.}}{\sim} P^*$ and $P_n := \frac{1}{n}\sum_{i=1}^n \delta_{X_i}$. For each $h \in \mathcal{H}$ define the centered empirical-process coordinate

$$\mathbb{X}_h := (P_n - P^*)h = \frac{1}{n}\sum_{i=1}^n \big(h(X_i) - \mathbb{E}h(X)\big).$$

We will control $\sup_{h \in \mathcal{H}} |\mathbb{X}_h|$.

Recall the (one of several equivalent) definitions of Talagrand's γ_2 functional. A sequence of subsets $(T_k)_{k\geq 0}$ of $\mathcal{H}$ is called *admissible* if $|T_0| = 1$ and $|T_k| \leq 2^{2^k}$ for all $k \geq 1$. Given an admissible (T_k) define, for each $h \in \mathcal{H}$,

$$d(h, T_k) := \inf_{t \in T_k} d(h, t).$$

Then

$$\gamma_2(\mathcal{H}, d) := \inf_{(T_k)\ \text{admissible}} \sup_{h\in\mathcal{H}} \sum_{k=0}^{\infty} 2^{k/2}\, d(h, T_k). \tag{133}$$

(We may assume separability w.r.t. d or otherwise interpret the supremum over a countable dense subclass; this is standard and does not affect the bound.)

Step 1 (Subgaussian increments for the empirical process). We first show that the process $\{\mathbb{X}_h\}_{h\in\mathcal{H}}$ has subgaussian increments w.r.t. the metric $d/\sqrt{n}$.

Lemma 8.16 (Averaging subgaussian variables)**.** *If $Y_1, \dots, Y_n$ are independent, centered, and $\|Y_i\|_{\psi_2} \leq \sigma$, then*

$$\left\| \frac{1}{n} \sum_{i=1}^{n} Y_i \right\|_{\psi_2} \leq \frac{C_0 \sigma}{\sqrt{n}}$$

for a universal constant C_0.

Proof. A standard characterization of the ψ_2 norm is: there exists a universal $c > 0$ such that $\|Y\|_{\psi_2} \leq \sigma$ implies $\mathbb{E}\exp\big(\lambda Y\big) \leq \exp(c\lambda^2\sigma^2)$ for all $\lambda \in \mathbb{R}$. Applying this to each Y_i and using independence,

$$\mathbb{E}\exp\left(\lambda \frac{1}{n}\sum_{i=1}^{n} Y_i\right) = \prod_{i=1}^{n} \mathbb{E}\exp\left(\frac{\lambda}{n} Y_i\right) \leq \prod_{i=1}^{n} \exp\left(c\frac{\lambda^2}{n^2}\sigma^2\right) = \exp\left(c\frac{\lambda^2}{n}\sigma^2\right).$$

This shows $\frac{1}{n}\sum_{i=1}^{n} Y_i$ is subgaussian with variance proxy $\asymp \sigma^2/n$, which is equivalent to the stated ψ_2 bound with some universal C_0. □

Now fix $h, h' \in \mathcal{H}$ and consider the increment

$$\mathbb{X}_h - \mathbb{X}_{h'} = \frac{1}{n}\sum_{i=1}^{n} \Big(Z_{h,h'}(X_i)\Big), \qquad Z_{h,h'}(x) := (h - h')(x) - \mathbb{E}(h - h')(X).$$

By the increment assumption (131), $\|Z_{h,h'}(X)\|_{\psi_2} \leq K d(h, h')$. Applying Lemma 8.16 with $\sigma = K d(h, h')$ yields

$$\|\mathbb{X}_h - \mathbb{X}_{h'}\|_{\psi_2} \leq \frac{C_1 K}{\sqrt{n}}\, d(h, h') \qquad \text{for all } h, h' \in \mathcal{H}, \tag{134}$$

for a universal C_1. Equivalently, the process has subgaussian increments with respect to the metric

$$\tilde{d}(h, h') := \frac{K}{\sqrt{n}}\, d(h, h').$$

Step 2 (A self-contained chaining tail bound for subgaussian-increment processes). We now prove the high-probability generic chaining inequality in the form needed here.

Lemma 8.17 (Chaining tail bound for subgaussian increments)**.** *Let $\{X_t\}_{t\in T}$ be a centered real-valued process indexed by (T,ρ), and assume*

$$\|X_t - X_s\|_{\psi_2} \le \rho(t,s) \qquad \text{for all } s,t \in T.$$

Then there exists a universal constant $C_2 > 0$ such that for every $u \ge 1$,

$$\Pr\Big(\sup_{t\in T} |X_t| > C_2\Big(\gamma_2(T,\rho) + u\,\mathrm{diam}(T,\rho)\Big)\Big) \;\le\; 2e^{-u^2}. \tag{135}$$

Proof. We prove an explicit chaining inequality by constructing nets and taking a union bound over scales.

(a) Fix an admissible sequence and associated projections. Let $(T_k)_{k\ge 0}$ be an admissible sequence of subsets of T (with $|T_0| = 1$, $|T_k| \le 2^{2^k}$). For each $t \in T$ choose a measurable selector $\pi_k(t) \in T_k$ satisfying

$$\rho\big(t, \pi_k(t)\big) \le \rho(t, T_k) + 2^{-k}\mathrm{diam}(T,\rho),$$

which is always possible by definition of the infimum and a standard 2^{-k}-approximation. (Any summable slack would do; this choice is convenient.)

Fix a basepoint $t_0 \in T_0$ (the unique element). Since X_{t_0} is centered and does not affect the supremum up to additive constants, we will bound $X_t - X_{t_0}$ and then use $|X_t| \le |X_t - X_{t_0}| + |X_{t_0}|$; the latter contributes only a universal constant absorbed into C_2.

(b) Telescoping decomposition. For any $t \in T$ and any integer $m \ge 1$, write

$$X_t - X_{t_0} = \sum_{k=1}^{m} \Big(X_{\pi_k(t)} - X_{\pi_{k-1}(t)}\Big) \;+\; \Big(X_t - X_{\pi_m(t)}\Big).$$

Letting $m \to \infty$ and using separability (or working on a countable dense subset) we may assume $\rho(t, \pi_m(t)) \to 0$ and hence $X_t - X_{\pi_m(t)} \to 0$ in L_2 (and almost surely along a subsequence), so it suffices to control the series of increments.

(c) Control increments uniformly over the net pairs at each scale. Fix $k \ge 1$. For any $t \in T$, the increment $X_{\pi_k(t)} - X_{\pi_{k-1}(t)}$ is a difference of two net points. Define the set of all such pairs:

$$\mathcal{P}_k := \{(a,b) \in T_k \times T_{k-1}\}.$$

Then $|\mathcal{P}_k| \le |T_k||T_{k-1}| \le 2^{2^k} 2^{2^{k-1}} \le 2^{2^{k+1}}$.

By the subgaussian increment hypothesis,

$$\|X_a - X_b\|_{\psi_2} \le \rho(a,b) \qquad \text{for all } (a,b) \in \mathcal{P}_k.$$

Hence there is a universal $c > 0$ such that for every $(a,b) \in \mathcal{P}_k$ and every $v \ge 0$,

$$\Pr\big(|X_a - X_b| > v\,\rho(a,b)\big) \le 2e^{-cv^2}. \tag{136}$$

Now set $v_k := u\,2^{k/2}$ with $u \ge 1$. By a union bound over all pairs in $\mathcal{P}_k$ and (136),

$$\Pr\Big(\exists (a,b) \in \mathcal{P}_k :\; |X_a - X_b| > v_k\,\rho(a,b)\Big) \le |\mathcal{P}_k| \cdot 2e^{-cv_k^2} \le 2\exp\Big((\log 2)\,2^{k+1} - cu^2 2^k\Big).$$

Choose $u \ge 1$ and absorb constants so that $cu^2 2^k$ dominates $(\log 2)2^{k+1}$; explicitly, for a universal choice of constants in front of u we obtain a bound of the form

$$\Pr\Big(\exists (a,b) \in \mathcal{P}_k :\; |X_a - X_b| > C_3 u\,2^{k/2}\rho(a,b)\Big) \le 2e^{-u^2 2^{k-1}},$$

for a universal C_3 and all $k \ge 1$. Summing over $k \ge 1$ gives

$$\Pr\Big(\exists k \ge 1\ \exists (a,b) \in \mathcal{P}_k :\; |X_a - X_b| > C_3 u\,2^{k/2}\rho(a,b)\Big) \le \sum_{k\ge 1} 2e^{-u^2 2^{k-1}} \le 2e^{-u^2},$$

using that $u \geq 1$ makes the series dominated by its first term.

Therefore, on an event Ω_u with $\Pr(\Omega_u) \geq 1 - 2e^{-u^2}$, we have *simultaneously for all* $k \geq 1$ *and all* $(a, b) \in \mathcal{P}_k$,

$$|X_a - X_b| \leq C_3 u \, 2^{k/2} \rho(a, b). \tag{137}$$

(d) Sum the scale bounds. Fix $t \in T$ and work on Ω_u. Apply (137) to $a = \pi_k(t)$ and $b = \pi_{k-1}(t)$:

$$|X_{\pi_k(t)} - X_{\pi_{k-1}(t)}| \leq C_3 u \, 2^{k/2} \, \rho\big(\pi_k(t), \pi_{k-1}(t)\big).$$

Using the triangle inequality for ρ,

$$\rho\big(\pi_k(t), \pi_{k-1}(t)\big) \leq \rho\big(t, \pi_k(t)\big) + \rho\big(t, \pi_{k-1}(t)\big).$$

Thus

$$\begin{aligned} |X_t - X_{t_0}| &\leq \sum_{k \geq 1} |X_{\pi_k(t)} - X_{\pi_{k-1}(t)}| \\ &\leq C_3 u \sum_{k \geq 1} 2^{k/2} \Big(\rho(t, \pi_k(t)) + \rho(t, \pi_{k-1}(t)) \Big) \\ &\leq C_4 u \sum_{k \geq 0} 2^{k/2} \rho(t, \pi_k(t)), \end{aligned}$$

for a universal C_4. By construction of $\pi_k(t)$,

$$\rho(t, \pi_k(t)) \leq \rho(t, T_k) + 2^{-k} \mathrm{diam}(T, \rho).$$

Therefore,

$$|X_t - X_{t_0}| \leq C_4 u \sum_{k \geq 0} 2^{k/2} \rho(t, T_k) \; + \; C_4 u \, \mathrm{diam}(T, \rho) \sum_{k \geq 0} 2^{k/2} 2^{-k}.$$

The last series $\sum_{k \geq 0} 2^{-k/2}$ converges to a universal constant, so

$$|X_t - X_{t_0}| \leq C_5 u \sum_{k \geq 0} 2^{k/2} \rho(t, T_k) + C_5 u \, \mathrm{diam}(T, \rho),$$

for universal C_5. Taking the supremum over $t \in T$ gives, on Ω_u,

$$\sup_{t \in T} |X_t - X_{t_0}| \leq C_5 u \left(\sup_{t \in T} \sum_{k \geq 0} 2^{k/2} \rho(t, T_k) \right) + C_5 u \, \mathrm{diam}(T, \rho).$$

Finally, take the infimum over admissible (T_k) and use the definition (133) to obtain

$$\sup_{t \in T} |X_t - X_{t_0}| \leq C_5 u \, \gamma_2(T, \rho) + C_5 u \, \mathrm{diam}(T, \rho).$$

Absorbing $|X_{t_0}|$ into constants (or taking t_0 such that $X_{t_0} = 0$ if available) yields (135) with a universal constant C_2. □

Step 3 (Apply Lemma 8.17 to the empirical process). Apply Lemma 8.17 to the process $\mathbb{X}_h$ indexed by $T = \mathcal{H}$ with the metric

$$\rho(h, h') := \frac{C_1 K}{\sqrt{n}} \, d(h, h'),$$

which is valid by (134). Then for every $u \geq 1$, with probability at least $1 - 2e^{-u^2}$,

$$\sup_{h \in \mathcal{H}} |\mathbb{X}_h| \leq C_2 \left(\gamma_2(\mathcal{H}, \rho) + u \, \mathrm{diam}(\mathcal{H}, \rho) \right).$$

Using the scaling properties $\gamma_2(\mathcal{H}, c\rho) = c\,\gamma_2(\mathcal{H}, \rho)$ and $\operatorname{diam}(\mathcal{H}, c\rho) = c\operatorname{diam}(\mathcal{H}, \rho)$, we have

$$\gamma_2(\mathcal{H}, \rho) = \frac{C_1 K}{\sqrt{n}}\gamma_2(\mathcal{H}, d), \qquad \operatorname{diam}(\mathcal{H}, \rho) = \frac{C_1 K}{\sqrt{n}}\operatorname{diam}(\mathcal{H}, d).$$

Therefore, with probability at least $1 - 2e^{-u^2}$,

$$\sup_{h\in\mathcal{H}} |(P_n - P^*)h| \le CK\left(\frac{\gamma_2(\mathcal{H}, d)}{\sqrt{n}} + \frac{u\operatorname{diam}(\mathcal{H}, d)}{\sqrt{n}}\right),$$

for a universal C (absorbing constants C_1, C_2).

Set $u := \sqrt{\log(2/\delta)} \asymp \sqrt{\log(1/\delta)}$ so that $2e^{-u^2} \le \delta$. This yields (132). □

Corollary 8.18 (Bracketing/entropy specialization)**.** *Under the hypotheses of Theorem 8.15, one may upper bound $\gamma_2(\mathcal{H}, d)$ by entropy integrals. If $N(\epsilon, \mathcal{H}, d)$ denotes covering numbers, then*

$$\gamma_2(\mathcal{H}, d) \;\lesssim\; \int_0^{\operatorname{diam}(\mathcal{H},d)} \sqrt{\log N(\epsilon, \mathcal{H}, d)}\, d\epsilon,$$

yielding a Dudley-type bound.

Proof. We give a direct construction from covers.

For $k \ge 0$ define the scale

$$\epsilon_k := \inf\left\{\epsilon > 0 :\ \log N(\epsilon, \mathcal{H}, d) \le 2^k\right\}.$$

Then there exists an ϵ_k-net $T_k \subseteq \mathcal{H}$ with $|T_k| \le N(\epsilon_k, \mathcal{H}, d) \le e^{2^k} \le 2^{2^k}$ (up to a harmless change of constants in the definition of ϵ_k), so $(T_k)_{k\ge0}$ is admissible after ensuring $|T_0| = 1$ (replace T_0 by a singleton).

By definition of ϵ_k, we have $d(h, T_k) \le \epsilon_k$ for all $h \in \mathcal{H}$. Therefore,

$$\gamma_2(\mathcal{H}, d) \le \sup_{h\in\mathcal{H}} \sum_{k\ge0} 2^{k/2} d(h, T_k) \le \sum_{k\ge0} 2^{k/2}\epsilon_k.$$

Now relate the discrete sum to the entropy integral by a standard layer-cake argument: since $\log N(\epsilon, \mathcal{H}, d)$ is nonincreasing in ϵ, the definition of ϵ_k implies that $\epsilon \in (\epsilon_{k+1}, \epsilon_k]$ entails $\sqrt{\log N(\epsilon, \mathcal{H}, d)} \gtrsim 2^{k/2}$. Hence

$$\int_0^{\operatorname{diam}(\mathcal{H},d)} \sqrt{\log N(\epsilon, \mathcal{H}, d)}\, d\epsilon \;\gtrsim\; \sum_{k\ge0} 2^{k/2}(\epsilon_k - \epsilon_{k+1}) \;\asymp\; \sum_{k\ge0} 2^{k/2}\epsilon_k,$$

where the last equivalence uses summation by parts and $\epsilon_k \downarrow 0$. Combining the displays yields the claimed bound (absorbing universal constants into $\lesssim$). □

8.5.3. *Putting the pieces together.* Combining Theorem 8.14 with Theorem 8.15 yields a high-probability bound on Δ_n, and inserting that bound into Theorem 8.12 yields an explicit finite-sample rate.

Theorem 8.19 (Sample EM in a general IPM: high-probability bound)**.** *Assume:*

(i) population contraction-to-set (106) on B_S for $A = \bar{\Theta}_S$ with constant $\gamma \in (0, 1)$;
(ii) strong concavity Assumption 8.1 with constant λ;
(iii) IPM modulus Assumption 8.11 on B_S;
(iv) the EM-induced class $\mathcal{H}$ in (130) satisfies the conditions of Theorem 8.15.

Then there exists a universal constant $C > 0$ such that for every $\delta \in (0,1)$, with probability at least $1-\delta$, the sample EM iterates satisfy for all $t \geq 0$,

(138)
$$\operatorname{dist}_{\mathcal{F}}\left(\left[\hat{\theta}_t\right], \bar{\Theta}/G\right) \;\leq\; \omega_{\mathcal{F}}\left(\gamma^t \operatorname{dist}(\hat{\theta}_0, \bar{\Theta}_S) + \frac{1-\gamma^t}{1-\gamma}\cdot\frac{2}{\lambda}\cdot C\left[\frac{\gamma_2(\mathcal{H}, L_2(P^*))}{\sqrt{n}} + \operatorname{diam}(\mathcal{H}, L_2(P^*))\sqrt{\frac{\log(1/\delta)}{n}}\right]\right).$$

Proof. **Roadmap.** We prove (138) by chaining together three deterministic implications:

(empirical-process control of $\hat{G} - G$) $\Longrightarrow$ (control of the EM update perturbation Δ_n) $\Longrightarrow$ (perturbed contractio

Each implication is made explicit below.

Step 0 (Recall the objects and where randomness enters). The only randomness is through the sample $X_1, \ldots, X_n \overset{\text{i.i.d.}}{\sim} P^*$, hence through P_n and the sample EM operator $\hat{T}$. The population objects (T, G, $\bar{\Theta}_S$, γ, λ, $\omega_{\mathcal{F}}$) are deterministic.

Recall that, by definition of the EM-induced class $\mathcal{H}$ in (130),

$$\sup_{(\theta',\theta)\in B_S\times B_S} \|\hat{G}(\theta',\theta) - G(\theta',\theta)\| = \sup_{h\in\mathcal{H}} \|(P_n - P^*)h\|, \tag{139}$$

where $\hat{G}(\theta',\theta)$ is the sample analogue of the population gradient map $G(\theta',\theta)$, and each coordinate of $\hat{G} - G$ is an empirical-process average over a corresponding $h \in \mathcal{H}$. Identity (139) is a definition-level rewriting: it is precisely why $\mathcal{H}$ was introduced.

Step 1 (Apply the generic chaining bound to control $\sup\|\hat{G} - G\|$). By assumption (iv), $\mathcal{H}$ satisfies the hypotheses of Theorem 8.15. Therefore, for every $\delta \in (0,1)$, there exists a universal constant $C > 0$ such that with probability at least $1-\delta$,

$$\sup_{h\in\mathcal{H}} \left|(P_n - P^*)h\right| \;\leq\; C\left(\frac{\gamma_2(\mathcal{H}, L_2(P^*))}{\sqrt{n}} + \operatorname{diam}(\mathcal{H}, L_2(P^*))\sqrt{\frac{\log(1/\delta)}{n}}\right). \tag{140}$$

On this event, combine (140) with the identity (139) to obtain

$$\sup_{(\theta',\theta)\in B_S\times B_S} \|\hat{G}(\theta',\theta) - G(\theta',\theta)\| \leq C\left(\frac{\gamma_2(\mathcal{H}, L_2(P^*))}{\sqrt{n}} + \operatorname{diam}(\mathcal{H}, L_2(P^*))\sqrt{\frac{\log(1/\delta)}{n}}\right). \tag{141}$$

Step 2 (Convert gradient-approximation error into a uniform EM-map perturbation bound). Define the uniform perturbation size

$$\Delta_n := \sup_{\theta\in B_S} \|\hat{T}(\theta) - T(\theta)\|.$$

By Theorem 8.14 (a deterministic stability result for strongly concave maximization), Assumption 8.1 with constant λ implies

$$\Delta_n \;\leq\; \frac{2}{\lambda} \sup_{(\theta',\theta)\in B_S\times B_S} \|\hat{G}(\theta',\theta) - G(\theta',\theta)\|. \tag{142}$$

(Here is where strong concavity enters: it converts a gradient perturbation into a parameter perturbation with factor $2/\lambda$.)

Now work on the high-probability event where (141) holds. Substituting (141) into (142) yields the explicit bound

$$\Delta_n \;\leq\; \frac{2}{\lambda}\cdot C\left(\frac{\gamma_2(\mathcal{H}, L_2(P^*))}{\sqrt{n}} + \operatorname{diam}(\mathcal{H}, L_2(P^*))\sqrt{\frac{\log(1/\delta)}{n}}\right). \tag{143}$$

Step 3 (Apply the sample-EM quotient IPM bound with this Δ_n). Under assumption (i) (population contraction-to-set on B_S with target set $A = \bar{\Theta}_S$) and assumption (iii) (the IPM modulus on B_S), Theorem 8.12 gives, deterministically for every $t \geq 0$,

$$\operatorname{dist}_{\mathcal{F}}\left(\left[\hat{\theta}_t\right], \bar{\Theta}/G\right) \;\leq\; \omega_{\mathcal{F}}\left(\gamma^t \operatorname{dist}(\hat{\theta}_0, \bar{\Theta}_S) + \frac{1-\gamma^t}{1-\gamma}\,\Delta_n\right). \tag{144}$$

(Here the factor $(1-\gamma)^{-1}$ comes from unrolling the perturbed contraction recursion.)

Finally, substitute the bound (143) on Δ_n into (144). This yields

$$\operatorname{dist}_{\mathcal{F}}\left(\left[\hat{\theta}_t\right], \bar{\Theta}/G\right) \;\leq\; \omega_{\mathcal{F}}\left(\gamma^t \operatorname{dist}(\hat{\theta}_0, \bar{\Theta}_S) + \frac{1-\gamma^t}{1-\gamma}\cdot\frac{2}{\lambda}\cdot C\left[\frac{\gamma_2(\mathcal{H}, L_2(P^*))}{\sqrt{n}} + \operatorname{diam}(\mathcal{H}, L_2(P^*))\sqrt{\frac{\log(1/\delta)}{n}}\right]\right),$$

which is exactly (138).

Step 4 (Probability statement). The bound is valid on the event (140), which holds with probability at least $1-\delta$ by Theorem 8.15. All subsequent steps were deterministic conditional on this event. Therefore (138) holds with probability at least $1-\delta$. □

Remark 8.20 (Interpretation of (138)). Equation (138) exhibits the standard "optimization + statistics" decomposition: a geometric optimization term $\gamma^t \operatorname{dist}(\hat{\theta}_0, \bar{\Theta}_S)$ and a statistical floor of order $\frac{1}{1-\gamma}\cdot\frac{1}{\lambda}\cdot\left(\gamma_2(\mathcal{H})/\sqrt{n}\right)$. The factor $(1-\gamma)^{-1}$ is algorithmically unavoidable for perturbed contractions, while λ^{-1} reflects curvature of the M-step surrogate. The statistical complexity is driven by the EM-induced class $\mathcal{H}$, not by $\mathcal{F}$ directly.

9. Extensions and refinements: inexact EM, basin invariance, and sample-splitting

Section 8 reduced finite-sample EM to a *perturbed contraction* analysis controlled by the operator deviation $\Delta_n = \sup_{\theta\in B_S}\|\hat{T}(\theta) - T(\theta)\|$. This section strengthens that theory in three directions that are frequently indispensable in rigorous applications: (i) explicit *basin invariance* conditions ensuring sample EM remains inside the local basin, (ii) *inexact* (approximate) E/M steps with time-varying errors, and (iii) *sample-splitting* (fresh-batch) EM yielding sharper high-probability bounds over long horizons.

Throughout, we retain the maximally general evaluation metric: an arbitrary IPM $d_{\mathcal{F}}$ and its quotient form $\bar{d}_{\mathcal{F}}([\theta], [\theta']) = d_{\mathcal{F}}(P_\theta, P_{\theta'})$ (Section 4).

9.1. Basin invariance: when does sample EM stay in B_S?

The deterministic results in Section 8 implicitly assume $\hat{\theta}_t \in B_S$ for all t. This is not automatic: even a small perturbation can eject an iterate from the basin unless one imposes a margin-to-boundary condition. We now state and prove the sharp basin-invariance lemma used in most local EM analyses (cf. the basin-based program in [1]).

Assumption 9.1 (Absorbing basin for the population map)**.** $B_S \subseteq S$ is nonempty and satisfies $T(B_S) \subseteq B_S$.

Assumption 9.2 (Margin to the boundary)**.** There exists $r > 0$ such that the closed r-neighborhood of $T(B_S)$ is contained in B_S, i.e.

$$\{u \in \mathbb{R}^d : \operatorname{dist}(u, T(B_S)) \leq r\} \;\subseteq\; B_S. \tag{145}$$

Lemma 9.3 (Uniform operator perturbation implies basin invariance)**.** *Assume Assumption 9.1 and Assumption 9.2. If $\Delta_n \leq r$, then $\hat{T}(B_S) \subseteq B_S$. Consequently, if $\hat{\theta}_0 \in B_S$ and $\hat{\theta}_{t+1} = \hat{T}(\hat{\theta}_t)$, then $\hat{\theta}_t \in B_S$ for all $t \geq 0$.*

Proof. **Definitions and assumptions used.** Recall $\Delta_n := \sup_{\theta\in B_S} \|\hat{T}(\theta) - T(\theta)\|$. Assumption 9.1 states that

$$T(B_S) \subseteq B_S. \tag{146}$$

Assumption 9.2 states that there exists a radius $r > 0$ such that

$$\Big\{x \in \mathbb{R}^d : \ \mathrm{dist}\left(x, T(B_S)\right) \le r\Big\} \ \subseteq \ B_S. \tag{147}$$

In words: the r-neighborhood (Minkowski enlargement) of the image set $T(B_S)$ is contained in the basin B_S.

Assume $\Delta_n \le r$. We show $\hat{T}(B_S) \subseteq B_S$ by proving $\hat{T}(\theta) \in B_S$ for an arbitrary $\theta \in B_S$.

Step 1 (Pointwise proximity of $\hat{T}(\theta)$ to the image set $T(B_S)$). Fix any $\theta \in B_S$. Since $T(\theta) \in T(B_S)$ by definition of the image set, we can upper bound the distance from $\hat{T}(\theta)$ to $T(B_S)$ by a specific candidate point $T(\theta) \in T(B_S)$:

$$\mathrm{dist}\left(\hat{T}(\theta), T(B_S)\right) := \inf_{y\in T(B_S)} \|\hat{T}(\theta) - y\| \le \|\hat{T}(\theta) - T(\theta)\|. \tag{148}$$

By definition of Δ_n as a supremum over B_S, and since $\theta \in B_S$,

$$\|\hat{T}(\theta) - T(\theta)\| \le \Delta_n. \tag{149}$$

Combining (148) and (149) and using the assumed bound $\Delta_n \le r$ gives

$$\mathrm{dist}\left(\hat{T}(\theta), T(B_S)\right) \le r. \tag{150}$$

Step 2 (Apply the margin assumption to conclude $\hat{T}(\theta) \in B_S$). Inequality (150) says precisely that $\hat{T}(\theta)$ belongs to the r-neighborhood of $T(B_S)$:

$$\hat{T}(\theta) \in \Big\{x \in \mathbb{R}^d : \ \mathrm{dist}\left(x, T(B_S)\right) \le r\Big\}.$$

By Assumption 9.2, the right-hand set is contained in B_S, hence $\hat{T}(\theta) \in B_S$.

Since $\theta \in B_S$ was arbitrary, we have shown $\hat{T}(\theta) \in B_S$ for all $\theta \in B_S$, i.e.

$$\hat{T}(B_S) \subseteq B_S.$$

Step 3 (Forward invariance for the iterates). Assume $\hat{\theta}_0 \in B_S$ and define $\hat{\theta}_{t+1} = \hat{T}(\hat{\theta}_t)$. We prove by induction that $\hat{\theta}_t \in B_S$ for all $t \ge 0$. The base case $t = 0$ holds by hypothesis. If $\hat{\theta}_t \in B_S$, then $\hat{\theta}_{t+1} = \hat{T}(\hat{\theta}_t) \in \hat{T}(B_S) \subseteq B_S$ by the first part of the lemma. Thus, by induction, $\hat{\theta}_t \in B_S$ for all $t \ge 0$. □

Remark 9.4 (Sharpness of the margin condition). Assumption 9.2 is essentially tight: if $T(B_S)$ approaches the boundary of B_S with no positive margin, then arbitrarily small uniform perturbations can force $\hat{T}(\theta) \notin B_S$ for some θ.

9.2. **A sharper bound for Δ_n (removing the unnecessary factor 2).** Theorem 8.14 in Section 8 stated a conservative prefactor $2/\lambda$. Under the same hypotheses one can obtain the optimal prefactor $1/\lambda$ without any additional assumptions.

Recall $G(\theta', \theta) = \nabla_{\theta'} Q(\theta'; \theta)$ and $\hat{G}(\theta', \theta) = \nabla_{\theta'} \hat{Q}(\theta'; \theta)$ as in (124).

Theorem 9.5 (Optimal strong-concavity control of Δ_n)**.** *Assume Assumption 8.1 with constant $\lambda > 0$ and that $T(\theta), \hat{T}(\theta) \in B_S$ for all $\theta \in B_S$. Assume moreover that B_S is an open subset of $\mathbb{R}^d$ (or, more generally, that the maximizers $T(\theta)$ and $\hat{T}(\theta)$ are interior points of the feasible set so that first-order conditions take the form of vanishing gradients). Then*

$$\Delta_n := \sup_{\theta \in B_S} \|\hat{T}(\theta) - T(\theta)\| \;\le\; \frac{1}{\lambda} \sup_{\theta \in B_S} \sup_{\theta' \in B_S} \|\hat{G}(\theta', \theta) - G(\theta', \theta)\|. \tag{151}$$

Proof. **Step 0 (Set up notation; reduce to a pointwise bound).** Fix an arbitrary $\theta \in B_S$. Define the population and sample slice objectives

$$Q_\theta(\theta') := Q(\theta' \mid \theta), \qquad \hat{Q}_\theta(\theta') := \hat{Q}(\theta' \mid \theta),$$

and their gradients (with respect to θ')

$$\nabla Q_\theta(\theta') = G(\theta', \theta), \qquad \nabla \hat{Q}_\theta(\theta') = \hat{G}(\theta', \theta).$$

Write

$$v := T(\theta) \in B_S, \qquad u := \hat{T}(\theta) \in B_S.$$

Our goal for this fixed θ is to prove

$$\|u - v\| \;\le\; \frac{1}{\lambda} \sup_{\theta' \in B_S} \|\hat{G}(\theta', \theta) - G(\theta', \theta)\|. \tag{152}$$

Once (152) is established, taking $\sup_{\theta \in B_S}$ on both sides yields (151) (because the left-hand side becomes Δ_n by definition).

Step 1 (Strong concavity $\Rightarrow$ sharp inverse-Lipschitz property of the gradient). We isolate and prove the quantitative inequality that converts a *gradient difference* into a *parameter difference*. This is the sharp step where the constant $1/\lambda$ is obtained.

Lemma 9.6 (Strong concavity implies inverse-Lipschitz gradient)**.** *Let $f : B_S \to \mathbb{R}$ be differentiable and λ-strongly concave on an open convex set $B_S \subseteq \mathbb{R}^d$. Then for all $x, y \in B_S$,*

$$\|x - y\| \;\le\; \frac{1}{\lambda} \|\nabla f(x) - \nabla f(y)\|. \tag{153}$$

Equivalently, ∇f is injective and its inverse (on its range) is $1/\lambda$-Lipschitz.

Proof. We proceed in two substeps.

(a) Strong concavity $\Rightarrow$ strong monotonicity of the gradient. A standard characterization of λ-strong concavity is: for all $x, y \in B_S$,

$$f(y) \le f(x) + \langle \nabla f(x), y - x \rangle - \frac{\lambda}{2} \|y - x\|^2. \tag{154}$$

Swap x and y in (154) to get

$$f(x) \le f(y) + \langle \nabla f(y), x - y \rangle - \frac{\lambda}{2} \|x - y\|^2. \tag{155}$$

Add (154) and (155). The left-hand side becomes $f(x) + f(y)$, and the $f(x), f(y)$ terms cancel, yielding

$$0 \le \langle \nabla f(x), y - x \rangle + \langle \nabla f(y), x - y \rangle - \lambda \|x - y\|^2.$$

Since $\langle \nabla f(x), y - x \rangle = -\langle \nabla f(x), x - y \rangle$ and $\langle \nabla f(y), x - y \rangle$ is unchanged, this becomes

$$\langle \nabla f(y) - \nabla f(x),\, x - y \rangle \;\ge\; \lambda \|x - y\|^2. \tag{156}$$

Equivalently,

$$\langle \nabla f(x) - \nabla f(y),\, x - y \rangle \;\le\; -\lambda \|x - y\|^2. \tag{157}$$

This is the strong monotonicity (with negative sign, since f is concave) of ∇f.

(b) Strong monotonicity $\Rightarrow$ inverse-Lipschitz. Apply Cauchy–Schwarz to the inner product:

$$\big|\langle \nabla f(x) - \nabla f(y),\, x - y \rangle\big| \le \|\nabla f(x) - \nabla f(y)\|\, \|x - y\|.$$

From (157), the quantity $\langle \nabla f(x) - \nabla f(y), x - y \rangle$ is *nonpositive*, hence

$$-\langle \nabla f(x) - \nabla f(y),\, x - y \rangle \ge \lambda \|x - y\|^2.$$

Also,

$$-\langle \nabla f(x) - \nabla f(y),\, x - y \rangle = \langle \nabla f(y) - \nabla f(x),\, x - y \rangle \le \|\nabla f(x) - \nabla f(y)\|\, \|x - y\|.$$

Combining the two displayed inequalities gives

$$\lambda \|x - y\|^2 \le \|\nabla f(x) - \nabla f(y)\|\, \|x - y\|.$$

If $x = y$, the desired inequality (153) is trivial. If $x \neq y$, divide by $\|x - y\| > 0$ to obtain

$$\|x - y\| \le \frac{1}{\lambda} \|\nabla f(x) - \nabla f(y)\|,$$

which is (153). □

Apply Lemma 9.6 to $f = Q_\theta$ (which is λ-strongly concave by Assumption 8.1) and the points $x = u = \hat{T}(\theta)$ and $y = v = T(\theta)$:

$$\|u - v\| \le \frac{1}{\lambda} \|\nabla Q_\theta(u) - \nabla Q_\theta(v)\| = \frac{1}{\lambda} \|G(u, \theta) - G(v, \theta)\|. \tag{158}$$

Step 2 (Justify the first-order conditions at $T(\theta)$ and $\hat{T}(\theta)$). We now show that the maximizer of a differentiable function over an open set must have vanishing gradient. This is the point where the assumption "$T(\theta), \hat{T}(\theta) \in B_S$" is interpreted as "interior maximizers".

Lemma 9.7 (Interior maximizer implies zero gradient)**.** *Let $U \subseteq \mathbb{R}^d$ be open and let $f : U \to \mathbb{R}$ be differentiable. If $x^\star \in U$ is a (local, hence also global) maximizer of f over U, then $\nabla f(x^\star) = 0$.*

Proof. Assume by contradiction that $\nabla f(x^\star) \neq 0$. Let $v := \nabla f(x^\star)/\|\nabla f(x^\star)\|$ so that $\|v\| = 1$. Consider the one-dimensional function $\varphi(t) := f(x^\star + tv)$. Since U is open and $x^\star \in U$, there exists $\epsilon > 0$ such that $x^\star + tv \in U$ for all $|t| < \epsilon$. By the chain rule,

$$\varphi'(0) = \langle \nabla f(x^\star), v \rangle = \|\nabla f(x^\star)\| > 0.$$

Hence for sufficiently small $t \in (0, \epsilon)$ we have $\varphi(t) > \varphi(0)$, i.e. $f(x^\star + tv) > f(x^\star)$, contradicting that $x^\star$ is a maximizer. Therefore $\nabla f(x^\star) = 0$. □

Apply Lemma 9.7 to $f = Q_\theta$ and $x^\star = v = T(\theta) \in B_S$ (with B_S open):

$$G(v, \theta) = \nabla Q_\theta(v) = 0. \tag{159}$$

Similarly, apply it to $f = \hat{Q}_\theta$ and $x^\star = u = \hat{T}(\theta) \in B_S$:

$$\hat{G}(u, \theta) = \nabla \hat{Q}_\theta(u) = 0. \tag{160}$$

Step 3 (Rewrite $\|G(u, \theta) - G(v, \theta)\|$ using the FOCs and bound by a uniform gradient perturbation). Using (159) in (158), we have

$$\|u - v\| \le \frac{1}{\lambda} \|G(u, \theta) - G(v, \theta)\| = \frac{1}{\lambda} \|G(u, \theta)\|.$$

Now use (160) to express $G(u,\theta)$ as a *difference* between population and sample gradients:

$$G(u,\theta) = G(u,\theta) - \hat{G}(u,\theta) \qquad (\text{since } \hat{G}(u,\theta) = 0).$$

Therefore

$$\|u - v\| \le \frac{1}{\lambda}\,\|G(u,\theta) - \hat{G}(u,\theta)\| \le \frac{1}{\lambda}\sup_{\theta' \in B_S} \|G(\theta',\theta) - \hat{G}(\theta',\theta)\|. \tag{161}$$

This is exactly the pointwise bound (152).

Step 4 (Take suprema over θ to obtain the bound on Δ_n). Since (161) holds for every fixed $\theta \in B_S$, taking $\sup_{\theta \in B_S}$ yields

$$\sup_{\theta \in B_S} \|\hat{T}(\theta) - T(\theta)\| \le \frac{1}{\lambda} \sup_{\theta \in B_S} \sup_{\theta' \in B_S} \|\hat{G}(\theta',\theta) - G(\theta',\theta)\|.$$

By definition, the left-hand side is Δ_n, giving (151). □

Remark 9.8 (If B_S is closed: variational-inequality version). If B_S is a *closed* convex set, the exact-gradient equalities $G(T(\theta),\theta) = 0$ and $\hat{G}(\hat{T}(\theta),\theta) = 0$ need not hold when the maximizers lie on the boundary. In that case the correct first-order condition is the variational inequality

$$\langle G(T(\theta),\theta), \theta' - T(\theta)\rangle \le 0 \quad \forall\, \theta' \in B_S,$$

(and similarly for $\hat{G}$), or equivalently $0 \in G(T(\theta),\theta) + N_{B_S}(T(\theta))$ with the normal cone N_{B_S}. One can still obtain perturbation bounds, but the clean $1/\lambda$ constant may require additional interiority/margin assumptions ensuring the maximizers are not on the boundary (which is exactly what the present theorem assumes implicitly via interiority).

Remark 9.9 (Tightness and the factor-of-2 improvement). The constant $1/\lambda$ is best possible from strong concavity alone. For instance, for any map $m(\theta)$ consider the quadratic surrogate

$$Q(\theta' \mid \theta) = -\frac{\lambda}{2}\|\theta' - m(\theta)\|^2, \qquad \hat{Q}(\theta' \mid \theta) = -\frac{\lambda}{2}\|\theta' - \hat{m}(\theta)\|^2.$$

Then $T(\theta) = m(\theta)$, $\hat{T}(\theta) = \hat{m}(\theta)$, and

$$G(\theta',\theta) = -\lambda(\theta' - m(\theta)), \qquad \hat{G}(\theta',\theta) = -\lambda(\theta' - \hat{m}(\theta)).$$

Hence $\sup_{\theta'} \|\hat{G}(\theta',\theta) - G(\theta',\theta)\| = \lambda\|\hat{m}(\theta) - m(\theta)\|$ and equality holds in (151).

9.3. **Inexact EM: deterministic and stochastic per-iteration errors.** In practice, exact E- and M-steps are often replaced by approximate computations: truncated posterior computations, approximate maximization, or inner-loop optimization. We model this by allowing an *inexact update map* $T_t^\sharp$ at each iteration t, with a controlled error.

Definition 9.10 (Inexact slice EM update with additive error)**.** Let $T : B_S \to B_S$ be the population slice EM map. An inexact update sequence is any sequence $\{\theta_t^\sharp\}_{t\ge 0} \subseteq B_S$ satisfying

$$\theta_{t+1}^\sharp = T(\theta_t^\sharp) + \varepsilon_t, \tag{162}$$

where $\varepsilon_t \in \mathbb{R}^d$ is an error term (deterministic or random).

Theorem 9.11 (Inexact contraction recursion (time-varying errors; sharp))**.** *Assume $T : B_S \to B_S$ is a contraction on B_S in the sense that there exists $\gamma \in (0,1)$ such that*

$$\|T(\theta) - T(\theta')\| \le \gamma\|\theta - \theta'\| \qquad \textit{for all } \theta,\theta' \in B_S. \tag{163}$$

Let $\{\theta_t\}_{t\ge 0}$ be the exact iterates $\theta_{t+1} = T(\theta_t)$ and let $\{\theta_t^\sharp\}_{t\ge 0}$ satisfy the inexact recursion

$$\theta_{t+1}^\sharp = T(\theta_t^\sharp) + \varepsilon_t, \qquad t \ge 0, \tag{164}$$

for some error sequence $\{\varepsilon_t\}_{t\ge 0} \subseteq \mathbb{R}^d$ such that the iterates remain in B_S (e.g. because B_S is forward-invariant and the errors are small enough). Then for every $t \ge 0$,

$$\|\theta_t^\sharp - \theta_t\| \le \gamma^t \|\theta_0^\sharp - \theta_0\| + \sum_{k=0}^{t-1} \gamma^{t-1-k} \, \|\varepsilon_k\|. \tag{165}$$

In particular, if $\sup_k \|\varepsilon_k\| \le \bar\varepsilon$, then

$$\|\theta_t^\sharp - \theta_t\| \le \gamma^t \|\theta_0^\sharp - \theta_0\| + \frac{1-\gamma^t}{1-\gamma}\,\bar\varepsilon \le \|\theta_0^\sharp - \theta_0\| + \frac{\bar\varepsilon}{1-\gamma}. \tag{166}$$

Moreover, the bound (165) is sharp *in the sense that equality can occur for suitable one-dimensional examples.*

Proof. **Step 0 (Define the error process and record what must be shown).** Define the *tracking error*

$$e_t := \theta_t^\sharp - \theta_t \in \mathbb{R}^d, \qquad E_t := \|e_t\| \in [0,\infty).$$

We will derive a scalar inequality for E_t and then solve it explicitly.

Step 1 (Exact identity for the one-step error; then a one-step inequality). Start from the two recursions:

$$\theta_{t+1} = T(\theta_t), \qquad \theta_{t+1}^\sharp = T(\theta_t^\sharp) + \varepsilon_t.$$

Subtract the first from the second:

$$\begin{aligned} e_{t+1} = \theta_{t+1}^\sharp - \theta_{t+1} &= \Big(T(\theta_t^\sharp) + \varepsilon_t\Big) - T(\theta_t) \\ &= \Big(T(\theta_t^\sharp) - T(\theta_t)\Big) + \varepsilon_t. \end{aligned} \tag{167}$$

Now take norms and apply the triangle inequality $\|a+b\| \le \|a\| + \|b\|$:

$$E_{t+1} = \|e_{t+1}\| \le \|T(\theta_t^\sharp) - T(\theta_t)\| + \|\varepsilon_t\|. \tag{168}$$

Finally apply the contraction property (163) to the pair $(\theta_t^\sharp, \theta_t)$ (both in B_S):

$$\|T(\theta_t^\sharp) - T(\theta_t)\| \le \gamma \|\theta_t^\sharp - \theta_t\| = \gamma E_t.$$

Substitute into (168) to obtain the fundamental one-step inequality

$$E_{t+1} \le \gamma E_t + \|\varepsilon_t\|, \qquad t \ge 0. \tag{169}$$

Step 2 (Solve the scalar inhomogeneous recursion explicitly). We now unroll (169) carefully, keeping track of indices.

For $t = 0$, (169) gives

$$E_1 \le \gamma E_0 + \|\varepsilon_0\|.$$

For $t = 1$,

$$E_2 \le \gamma E_1 + \|\varepsilon_1\| \le \gamma(\gamma E_0 + \|\varepsilon_0\|) + \|\varepsilon_1\| = \gamma^2 E_0 + \gamma\|\varepsilon_0\| + \|\varepsilon_1\|.$$

For $t = 2$,

$$E_3 \le \gamma E_2 + \|\varepsilon_2\| \le \gamma(\gamma^2 E_0 + \gamma\|\varepsilon_0\| + \|\varepsilon_1\|) + \|\varepsilon_2\| = \gamma^3 E_0 + \gamma^2\|\varepsilon_0\| + \gamma\|\varepsilon_1\| + \|\varepsilon_2\|.$$

The pattern suggests the general bound

$$E_t \le \gamma^t E_0 + \sum_{k=0}^{t-1} \gamma^{t-1-k} \|\varepsilon_k\|. \tag{170}$$

We now prove (170) by induction on t.

Claim 9.12. *For every integer $t \ge 0$, inequality (170) holds.*

Proof of Claim 9.12. For $t = 0$, the sum is empty and interpreted as 0, so (170) reads $E_0 \le E_0$, which is true.

Assume (170) holds for some $t \ge 0$. Apply the one-step recursion (169):

$$E_{t+1} \le \gamma E_t + \|\varepsilon_t\|.$$

Substitute the induction hypothesis for E_t:

$$E_{t+1} \le \gamma \left(\gamma^t E_0 + \sum_{k=0}^{t-1} \gamma^{t-1-k} \|\varepsilon_k\| \right) + \|\varepsilon_t\| = \gamma^{t+1} E_0 + \sum_{k=0}^{t-1} \gamma^{t-k} \|\varepsilon_k\| + \|\varepsilon_t\|.$$

Now observe that $\|\varepsilon_t\| = \gamma^{t-t} \|\varepsilon_t\|$ and that

$$\sum_{k=0}^{t-1} \gamma^{t-k} \|\varepsilon_k\| + \|\varepsilon_t\| = \sum_{k=0}^{t} \gamma^{t-k} \|\varepsilon_k\| = \sum_{k=0}^{(t+1)-1} \gamma^{(t+1)-1-k} \|\varepsilon_k\|.$$

Therefore

$$E_{t+1} \le \gamma^{t+1} E_0 + \sum_{k=0}^{t} \gamma^{t-k} \|\varepsilon_k\|,$$

which is exactly (170) with t replaced by $t+1$. This completes the induction. □

Since $E_t = \|\theta_t^\sharp - \theta_t\|$ and $E_0 = \|\theta_0^\sharp - \theta_0\|$, Claim 9.12 is precisely (165).

Step 3 (Uniform error specialization). Assume $\sup_{k \ge 0} \|\varepsilon_k\| \le \bar{\varepsilon}$. Then each term in the sum in (165) is bounded by $\gamma^{t-1-k} \bar{\varepsilon}$, hence

$$\sum_{k=0}^{t-1} \gamma^{t-1-k} \|\varepsilon_k\| \le \bar{\varepsilon} \sum_{k=0}^{t-1} \gamma^{t-1-k} = \bar{\varepsilon} \sum_{j=0}^{t-1} \gamma^j = \bar{\varepsilon} \frac{1-\gamma^t}{1-\gamma}, \tag{171}$$

where we used the change of variables $j = t - 1 - k$ and the finite geometric-series identity. Substituting (171) into (165) yields the first inequality in (166). The second inequality in (166) follows because $0 \le \gamma^t \le 1$ implies

$$\gamma^t \|\theta_0^\sharp - \theta_0\| \le \|\theta_0^\sharp - \theta_0\|, \qquad \frac{1-\gamma^t}{1-\gamma} \le \frac{1}{1-\gamma}.$$

Step 4 (Sharpness). The inequality chain used only (i) triangle inequality and (ii) the contraction inequality. Both can be tight. For instance, in one dimension with $T(x) = \gamma x$ (which is a contraction with factor γ on any interval containing 0), choose $\theta_0 = \theta_0^\sharp = 0$ and take $\varepsilon_t \ge 0$ for all t. Then the exact iterates satisfy $\theta_t \equiv 0$ and the inexact iterates satisfy $\theta_{t+1}^\sharp = \gamma \theta_t^\sharp + \varepsilon_t$, so $e_t = \theta_t^\sharp$ and

$$e_t = \sum_{k=0}^{t-1} \gamma^{t-1-k} \varepsilon_k,$$

which matches (165) with equality. □

Corollary 9.13 (Inexact EM to a set and IPM transfer)**.** *Assume the population contraction-to-set property* (106) *holds on B_S for $A = \bar{\Theta}_S$ with factor $\gamma \in (0,1)$, and assume the IPM modulus* (78) *holds on B_S with modulus $\omega_{\mathcal{F}}$. Let $\{\theta_t^\sharp\}_{t\geq 0} \subseteq B_S$ be any sequence such that*

$$\operatorname{dist}\big(\theta_{t+1}^\sharp, T(\theta_t^\sharp)\big) \leq \eta_t \qquad \text{for all } t \geq 0, \tag{172}$$

for some nonnegative sequence $\{\eta_t\}_{t\geq 0}$. Then for all $t \geq 0$,

$$\operatorname{dist}(\theta_t^\sharp, \bar{\Theta}_S) \leq \gamma^t \operatorname{dist}(\theta_0^\sharp, \bar{\Theta}_S) + \sum_{k=0}^{t-1} \gamma^{t-1-k} \eta_k, \tag{173}$$

and consequently

$$\operatorname{dist}_{\mathcal{F}}\left(\left[\theta_t^\sharp\right], \bar{\Theta}/G \right) \leq \omega_{\mathcal{F}}\left(\gamma^t \operatorname{dist}(\theta_0^\sharp, \bar{\Theta}_S) + \sum_{k=0}^{t-1} \gamma^{t-1-k} \eta_k \right). \tag{174}$$

Proof. **Part I: derive the perturbed contraction-to-set recursion.** Define $e_t := \operatorname{dist}(\theta_t^\sharp, \bar{\Theta}_S)$. We first show

$$e_{t+1} \leq \gamma e_t + \eta_t. \tag{175}$$

Fix $t \geq 0$. Start from the definition:

$$e_{t+1} = \operatorname{dist}(\theta_{t+1}^\sharp, \bar{\Theta}_S).$$

Apply the 1-Lipschitz property of distance-to-set (Lemma 8.3 from your earlier proof): for any x, y, $\operatorname{dist}(x, \bar{\Theta}_S) \leq \operatorname{dist}(y, \bar{\Theta}_S) + \|x - y\|$. With $x = \theta_{t+1}^\sharp$ and $y = T(\theta_t^\sharp)$,

$$\operatorname{dist}(\theta_{t+1}^\sharp, \bar{\Theta}_S) \leq \operatorname{dist}(T(\theta_t^\sharp), \bar{\Theta}_S) + \|\theta_{t+1}^\sharp - T(\theta_t^\sharp)\|. \tag{176}$$

By the inexactness assumption (172),

$$\|\theta_{t+1}^\sharp - T(\theta_t^\sharp)\| \geq \operatorname{dist}(\theta_{t+1}^\sharp, T(\theta_t^\sharp)) \quad \text{and hence} \quad \|\theta_{t+1}^\sharp - T(\theta_t^\sharp)\| \leq \eta_t$$

(the last implication holds if (172) is strengthened to the norm bound $\|\theta_{t+1}^\sharp - T(\theta_t^\sharp)\| \leq \eta_t$; if you only assume distance-to-a-point, these are identical because $\operatorname{dist}(x, \{y\}) = \|x - y\|$). Thus the second term in (176) is at most η_t.

For the first term, apply contraction-to-set (106) with $A = \bar{\Theta}_S$:

$$\operatorname{dist}(T(\theta_t^\sharp), \bar{\Theta}_S) \leq \gamma \operatorname{dist}(\theta_t^\sharp, \bar{\Theta}_S) = \gamma e_t.$$

Substituting into (176) yields (175).

Now unroll (175) exactly as in Theorem 9.11 (with $\|\varepsilon_t\|$ replaced by η_t), to obtain (173).

Part II: transfer distance-to-set in parameters to quotient IPM distance-to-set. Apply the analogue of Lemma 7.7 (as previously proved) at $\theta = \theta_t^\sharp$:

$$\inf_{\bar{\theta} \in \bar{\Theta}_S} d_{\mathcal{F}}(P_{\theta_t^\sharp}, P_{\bar{\theta}}) \leq \omega_{\mathcal{F}}\big(\operatorname{dist}(\theta_t^\sharp, \bar{\Theta}_S)\big) = \omega_{\mathcal{F}}(e_t).$$

As in your quotient arguments, by G-invariance and orbit-invariance,

$$\operatorname{dist}_{\mathcal{F}}\left(\left[\theta_t^\sharp\right], \bar{\Theta}/G \right) = \inf_{\bar{\theta} \in \bar{\Theta}_S} d_{\mathcal{F}}(P_{\theta_t^\sharp}, P_{\bar{\theta}}),$$

so

$$\operatorname{dist}_{\mathcal{F}}\left(\left[\theta_t^\sharp\right], \bar{\Theta}/G \right) \leq \omega_{\mathcal{F}}(e_t).$$

Finally substitute the bound (173) on e_t and use monotonicity of $\omega_{\mathcal{F}}$ to obtain (174). □

9.4. **Sample-splitting EM: independent operator errors and sharper probability control.** The analysis in Section 8 treats the *fixed-sample* map $\hat{T}$ and controls Δ_n uniformly on B_S. A complementary regime is *sample-splitting EM*: at iteration t one uses a fresh batch of data to form an independent map $\hat{T}_t$. While this can be statistically wasteful, it yields cleaner high-probability statements over long horizons because the operator errors are independent across t. This device is standard in nonasymptotic EM analyses as a proof technique [1].

Definition 9.14 (Sample-splitting EM)**.** Let the data be partitioned into m disjoint blocks of equal size b (so $n = mb$). For each $t \in \{0, \dots, m-1\}$, define the empirical surrogate $\hat{Q}_t(\cdot;\cdot)$ using only block t, and the associated slice update map

$$\hat{T}_t(\theta) := \arg\max_{\theta' \in B_S} \hat{Q}_t(\theta'; \theta).$$

The sample-splitting EM iterates are

$$\hat{\theta}_{t+1} = \hat{T}_t(\hat{\theta}_t), \qquad t = 0, \dots, m-1.$$

For each t, define the per-iteration operator deviation

$$\Delta_t := \sup_{\theta \in B_S} \|\hat{T}_t(\theta) - T(\theta)\|.$$

Because block t is independent of $\hat{\theta}_t$ (which depends only on earlier blocks), one can apply conditional concentration at each step and then take a union bound over t.

Theorem 9.15 (Sample-splitting perturbed contraction bound)**.** *Assume the contraction-to-set property* (106) *holds on B_S for $A = \bar{\Theta}_S$ with factor $\gamma \in (0,1)$, i.e.*

$$\operatorname{dist}\big(T(\theta), A\big) \le \gamma \operatorname{dist}(\theta, A) \qquad \text{for all } \theta \in B_S. \tag{177}$$

Let $\hat{\theta}_{t+1} = \hat{T}_t(\hat{\theta}_t)$ for $t = 0, \dots, m-1$, where each $\hat{T}_t : B_S \to B_S$ is a (possibly random) approximation of T and define the per-step operator deviations

$$\Delta_t := \sup_{\theta \in B_S} \|\hat{T}_t(\theta) - T(\theta)\| \in [0, \infty]. \tag{178}$$

Assume $\Delta_t < \infty$ and $\hat{\theta}_t \in B_S$ for all $t \le m$ (e.g. by a basin invariance lemma). Then for every $t \in \{0, 1, \dots, m\}$,

$$\operatorname{dist}(\hat{\theta}_t, A) \le \gamma^t \operatorname{dist}(\hat{\theta}_0, A) + \sum_{k=0}^{t-1} \gamma^{t-1-k} \Delta_k. \tag{179}$$

Proof. **Step 0 (A basic tool: distance-to-set is 1-Lipschitz).** We use the inequality (proved earlier as Lemma 8.3): for any nonempty set $A \subseteq \mathbb{R}^d$ and any $x, y \in \mathbb{R}^d$,

$$\operatorname{dist}(x, A) \le \operatorname{dist}(y, A) + \|x - y\|. \tag{180}$$

Step 1 (Derive the one-step scalar recursion). Fix $t \in \{0, \dots, m-1\}$ and set $e_t := \operatorname{dist}(\hat{\theta}_t, A)$. Using the iteration $\hat{\theta}_{t+1} = \hat{T}_t(\hat{\theta}_t)$,

$$e_{t+1} = \operatorname{dist}(\hat{\theta}_{t+1}, A) = \operatorname{dist}(\hat{T}_t(\hat{\theta}_t), A).$$

Apply (180) with $x = \hat{T}_t(\hat{\theta}_t)$ and $y = T(\hat{\theta}_t)$:

$$\operatorname{dist}(\hat{T}_t(\hat{\theta}_t), A) \le \operatorname{dist}(T(\hat{\theta}_t), A) + \|\hat{T}_t(\hat{\theta}_t) - T(\hat{\theta}_t)\|. \tag{181}$$

We bound the two terms on the right-hand side separately.

(a) Contraction-to-set term. Since $\hat{\theta}_t \in B_S$, the contraction-to-set property (177) gives

$$\operatorname{dist}(T(\hat{\theta}_t), A) \le \gamma \operatorname{dist}(\hat{\theta}_t, A) = \gamma e_t. \tag{182}$$

(b) Approximation term. By definition (178) of Δ_t as a supremum over $\theta \in B_S$ and since $\hat{\theta}_t \in B_S$,

$$\|\hat{T}_t(\hat{\theta}_t) - T(\hat{\theta}_t)\| \le \sup_{\theta \in B_S} \|\hat{T}_t(\theta) - T(\theta)\| = \Delta_t. \tag{183}$$

Substitute (182) and (183) into (181) to obtain the one-step recursion

$$e_{t+1} \le \gamma e_t + \Delta_t, \qquad t = 0, \dots, m-1. \tag{184}$$

Step 2 (Unroll the time-varying recursion). We now solve (184) explicitly. We claim that for every $t \in \{0, 1, \dots, m\}$,

$$e_t \le \gamma^t e_0 + \sum_{k=0}^{t-1} \gamma^{t-1-k} \Delta_k. \tag{185}$$

Claim 9.16. *Inequality (185) holds for all $t \in \{0, 1, \dots, m\}$.*

Proof of Claim 9.16. For $t = 0$, the sum is empty (interpreted as 0), and (185) reads $e_0 \le e_0$. Assume (185) holds at time t with $t \le m-1$. Then by (184),

$$e_{t+1} \le \gamma e_t + \Delta_t \le \gamma \left(\gamma^t e_0 + \sum_{k=0}^{t-1} \gamma^{t-1-k} \Delta_k \right) + \Delta_t = \gamma^{t+1} e_0 + \sum_{k=0}^{t-1} \gamma^{t-k} \Delta_k + \Delta_t.$$

Since $\Delta_t = \gamma^{t-t} \Delta_t$, the last two terms combine as

$$\sum_{k=0}^{t-1} \gamma^{t-k} \Delta_k + \Delta_t = \sum_{k=0}^{t} \gamma^{t-k} \Delta_k = \sum_{k=0}^{(t+1)-1} \gamma^{(t+1)-1-k} \Delta_k.$$

Thus

$$e_{t+1} \le \gamma^{t+1} e_0 + \sum_{k=0}^{t} \gamma^{t-k} \Delta_k,$$

which is exactly (185) with t replaced by $t+1$. This completes the induction. □

Finally, recalling that $e_t = \operatorname{dist}(\hat{\theta}_t, A)$ and $A = \bar{\Theta}_S$, we obtain (179). □

Theorem 9.17 (Sample-splitting EM in a general IPM: uniform high-probability bound)**.** *Assume:*

*(i) (*IPM modulus on the slice.*) For all $\theta, \theta' \in B_S$,*

$$d_{\mathcal{F}}(P_\theta, P_{\theta'}) \le \omega_{\mathcal{F}}(\|\theta - \theta'\|), \tag{186}$$

where $\omega_{\mathcal{F}} : [0, \infty) \to [0, \infty)$ is nondecreasing and (without loss) right-continuous.

*(ii) (*Per-step uniform operator deviation control.*) For each $t \in \{0, \dots, m-1\}$ and each $\delta \in (0, 1)$,*

$$\Pr\big(\Delta_t \le a_b(\delta)\big) \ \ge \ 1 - \delta, \tag{187}$$

where $\Delta_t := \sup_{\theta \in B_S} \|\hat{T}_t(\theta) - T(\theta)\|$ (cf. (178)) and $a_b(\delta)$ depends on the block size b.

*(iii) (*Forward invariance on the horizon.*) With probability one (or on the same event we work on below), the iterates $\hat{\theta}_t$ remain in B_S for all $t \le m$. (For instance, this holds whenever a basin invariance lemma applies, e.g. Lemma 9.3 under $\Delta_t \le r$.)*

*(iv) (*Population contraction-to-set.*)* *The population map T satisfies, for $A = \bar{\Theta}_S$ and all $\theta \in B_S$,*

$$\operatorname{dist}\big(T(\theta), A\big) \le \gamma \operatorname{dist}(\theta, A) \tag{188}$$

with a factor $\gamma \in (0,1)$.

Consider the sample-splitting iterates $\hat{\theta}_{t+1} = \hat{T}_t(\hat{\theta}_t)$ for $t = 0, \dots, m-1$ and define $\hat{\xi}_t := \left[\hat{\theta}_t\right]$. Then for any $\delta \in (0,1)$, with probability at least $1-\delta$, simultaneously for all $t \in \{0, 1, \dots, m\}$,

$$\operatorname{dist}_{\mathcal{F}}\left(\hat{\xi}_t, \bar{\Theta}/G\right) \le \omega_{\mathcal{F}}\left(\gamma^t \operatorname{dist}(\hat{\theta}_0, \bar{\Theta}_S) + \frac{1-\gamma^t}{1-\gamma}\, a_b(\delta/m)\right). \tag{189}$$

Proof. **Step 0 (Two reusable facts: a Lipschitz property and the quotient reduction).** *Fact A (distance-to-set is 1-Lipschitz).* For any nonempty $A \subseteq \mathbb{R}^d$ and any $x, y \in \mathbb{R}^d$,

$$\operatorname{dist}(x, A) \le \operatorname{dist}(y, A) + \|x - y\|. \tag{190}$$

Fact B (quotient IPM distance-to-set equals a slice infimum). By G-invariance of $\bar{\Theta}$ and orbit-invariance of the model (as in Theorem 8.12), for every $\theta \in B_S$,

$$\operatorname{dist}_{\mathcal{F}}\left([\theta], \bar{\Theta}/G\right) = \inf_{\bar{\theta} \in \bar{\Theta}} d_{\mathcal{F}}(P_\theta, P_{\bar{\theta}}) = \inf_{\bar{\theta} \in \bar{\Theta}_S} d_{\mathcal{F}}(P_\theta, P_{\bar{\theta}}). \tag{191}$$

The second equality uses that every orbit in $\bar{\Theta}/G$ that meets the local chart has a unique representative in $\bar{\Theta}_S$ and that replacing $\bar{\theta}$ by an orbit-equivalent representative does not change $P_{\bar{\theta}}$.

Fix $\delta \in (0,1)$ for the remainder of the proof.

Step 1 (Construct a single high-probability event controlling all steps). For each $t \in \{0, \dots, m-1\}$ define

$$E_t := \{\Delta_t \le a_b(\delta/m)\}.$$

By (187) with δ/m in place of δ,

$$\Pr(E_t) \ge 1 - \delta/m \qquad \text{for each } t.$$

By the union bound,

$$\Pr\left(E := \bigcap_{t=0}^{m-1} E_t\right) \ge 1 - \sum_{t=0}^{m-1} \Pr(E_t^c) \ge 1 - \sum_{t=0}^{m-1} \frac{\delta}{m} = 1 - \delta. \tag{192}$$

Hence, on the event E (which occurs with probability at least $1-\delta$), we have the *simultaneous* uniform bound

$$\Delta_t \le a_b(\delta/m) \qquad \text{for all } t = 0, \dots, m-1. \tag{193}$$

Step 2 (Derive the time-varying perturbed contraction-to-set recursion on E). Work on the event E and fix any $t \in \{0, \dots, m-1\}$. We first establish the one-step inequality

$$\operatorname{dist}(\hat{\theta}_{t+1}, A) \le \gamma \operatorname{dist}(\hat{\theta}_t, A) + \Delta_t, \qquad A = \bar{\Theta}_S. \tag{194}$$

Indeed,

$$\operatorname{dist}(\hat{\theta}_{t+1}, A) = \operatorname{dist}(\hat{T}_t(\hat{\theta}_t), A) \le \operatorname{dist}(T(\hat{\theta}_t), A) + \|\hat{T}_t(\hat{\theta}_t) - T(\hat{\theta}_t)\|$$

by (190) with $x = \hat{T}_t(\hat{\theta}_t)$ and $y = T(\hat{\theta}_t)$. Now, since $\hat{\theta}_t \in B_S$ by forward invariance (assumption (iii)), the contraction-to-set property (188) yields

$$\operatorname{dist}(T(\hat{\theta}_t), A) \le \gamma \operatorname{dist}(\hat{\theta}_t, A),$$

and by definition of Δ_t as a supremum over B_S,

$$\|\hat{T}_t(\hat{\theta}_t) - T(\hat{\theta}_t)\| \le \Delta_t.$$

Substituting these bounds proves (194).

Step 3 (Unroll the recursion with explicit index bookkeeping). Define $e_t := \operatorname{dist}(\hat{\theta}_t, A) \ge 0$. Then (194) becomes

$$e_{t+1} \le \gamma e_t + \Delta_t, \qquad t = 0, \dots, m-1. \tag{195}$$

We claim that for every $t \in \{0, 1, \dots, m\}$,

$$e_t \le \gamma^t e_0 + \sum_{k=0}^{t-1} \gamma^{t-1-k} \Delta_k. \tag{196}$$

Claim 9.18. *Inequality (196) holds for all $t \in \{0, 1, \dots, m\}$.*

Proof of Claim 9.18. For $t = 0$, the sum is empty and (196) reads $e_0 \le e_0$. Assume (196) holds for some $t \le m-1$. Then by (195),

$$e_{t+1} \le \gamma e_t + \Delta_t \le \gamma \left(\gamma^t e_0 + \sum_{k=0}^{t-1} \gamma^{t-1-k} \Delta_k \right) + \Delta_t = \gamma^{t+1} e_0 + \sum_{k=0}^{t-1} \gamma^{t-k} \Delta_k + \Delta_t.$$

Since $\Delta_t = \gamma^{t-t} \Delta_t$, the last two terms combine into $\sum_{k=0}^{t} \gamma^{t-k} \Delta_k = \sum_{k=0}^{(t+1)-1} \gamma^{(t+1)-1-k} \Delta_k$, so

$$e_{t+1} \le \gamma^{t+1} e_0 + \sum_{k=0}^{t} \gamma^{t-k} \Delta_k,$$

which is (196) with t replaced by $t+1$. □

Thus (196) holds. Now invoke the uniform bound (193) on E: for each $t \in \{0, \dots, m\}$,

$$\begin{aligned} e_t &\le \gamma^t e_0 + \sum_{k=0}^{t-1} \gamma^{t-1-k} \Delta_k \le \gamma^t e_0 + a_b(\delta/m) \sum_{k=0}^{t-1} \gamma^{t-1-k} \\ &= \gamma^t e_0 + a_b(\delta/m) \sum_{j=0}^{t-1} \gamma^j = \gamma^t e_0 + \frac{1-\gamma^t}{1-\gamma} a_b(\delta/m), \end{aligned} \tag{197}$$

where we used the change of variables $j = t - 1 - k$ and the finite geometric-series identity.

Recalling $e_t = \operatorname{dist}(\hat{\theta}_t, A)$ and $A = \bar{\Theta}_S$ gives, on E,

$$\operatorname{dist}(\hat{\theta}_t, \bar{\Theta}_S) \le \gamma^t \operatorname{dist}(\hat{\theta}_0, \bar{\Theta}_S) + \frac{1-\gamma^t}{1-\gamma} a_b(\delta/m). \tag{198}$$

Step 4 (Convert parameter distance-to-set into quotient IPM distance-to-set). Fix $t \in \{0, \dots, m\}$ and work on the event E. By Fact B (191),

$$\operatorname{dist}_{\mathcal{F}} \left(\left[\hat{\theta}_t \right], \bar{\Theta}/G \right) = \inf_{\bar{\theta} \in \bar{\Theta}_S} d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}}). \tag{199}$$

Let $r_t := \operatorname{dist}(\hat{\theta}_t, \bar{\Theta}_S)$. By the definition of infimum, for each $\epsilon > 0$ there exists $\bar{\theta}_{t,\epsilon} \in \bar{\Theta}_S$ such that

$$\|\hat{\theta}_t - \bar{\theta}_{t,\epsilon}\| \le r_t + \epsilon. \tag{200}$$

Apply the modulus (186) and monotonicity of $\omega_{\mathcal{F}}$:

$$d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}_{t,\epsilon}}) \le \omega_{\mathcal{F}}\big(\|\hat{\theta}_t - \bar{\theta}_{t,\epsilon}\|\big) \le \omega_{\mathcal{F}}(r_t + \epsilon).$$

Since the left-hand side is an upper bound on the infimum over $\bar{\theta} \in \bar{\Theta}_S$, we obtain

$$\inf_{\bar{\theta} \in \bar{\Theta}_S} d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}}) \le \omega_{\mathcal{F}}(r_t + \epsilon) \qquad \text{for all } \epsilon > 0.$$

Let $\epsilon \downarrow 0$ and use right-continuity of $\omega_{\mathcal{F}}$ to conclude

$$\inf_{\bar{\theta} \in \bar{\Theta}_S} d_{\mathcal{F}}(P_{\hat{\theta}_t}, P_{\bar{\theta}}) \le \omega_{\mathcal{F}}(r_t) = \omega_{\mathcal{F}}\big(\operatorname{dist}(\hat{\theta}_t, \bar{\Theta}_S)\big). \tag{201}$$

Combining (199) and (201) yields, on E,

$$\operatorname{dist}_{\mathcal{F}}\left(\left[\hat{\theta}_t\right], \bar{\Theta}/G\right) \le \omega_{\mathcal{F}}\big(\operatorname{dist}(\hat{\theta}_t, \bar{\Theta}_S)\big).$$

Finally substitute (198) into the right-hand side and use monotonicity of $\omega_{\mathcal{F}}$ to obtain (189) on E.

Step 5 (Conclude the probability statement and uniformity in t). By (192), $\Pr(E) \ge 1-\delta$. All inequalities in Steps 2–4 hold simultaneously for all $t \in \{0, \dots, m\}$ on the *same* event E (because (193) is uniform in t and the recursion is deterministic given the Δ_t's). Hence, with probability at least $1 - \delta$, the bound (189) holds simultaneously for all $t \le m$. □

Remark 9.19 (Comparison with fixed-sample EM). Fixed-sample EM reuses the same n observations across all iterations and can yield smaller *per-iteration* statistical error, but the induced dependence between $(\hat{\theta}_t)$ and the empirical fluctuations typically forces either (i) uniform empirical-process control that is stable over a data-dependent trajectory, or (ii) stability/algorithmic arguments. Sample-splitting EM avoids these issues: conditional on the past, each update uses fresh data, so per-step bounds such as (187) follow from a one-shot empirical-process inequality at block size b. The cost is explicit in (189): a δ/m union-bound penalty and statistical scaling $1/\sqrt{b}$ (with $b \approx n/m$ if the sample is split evenly). Which regime is preferable depends on the horizon m and whether one can obtain fixed-sample bounds that hold uniformly over t.

Appendix A. Measure-theoretic and conditional-expectation preliminaries

Throughout, let $(\mathsf{X}, \mathcal{F})$ and $(\mathsf{Z}, \mathcal{G})$ be measurable spaces for the observed and latent variables, respectively, and let $\Theta \subseteq \mathbb{R}^d$ be a parameter set. Write P_θ^{XZ} for the joint law of (X, Z) under parameter θ, and P_θ^X for the marginal law of X. The data-generating law of X is denoted by $P^\star$ (not assumed to lie in the model).

For any $\theta, \theta' \in \Theta$, we use the standard EM functionals

$$\Phi(\theta) \;:=\; \mathbb{E}_{P^\star}\big[\log p_\theta(X)\big], \qquad Q(\theta \mid \theta') \;:=\; \mathbb{E}_{P^\star}\Big[\mathbb{E}_{\kappa_{\theta'}(X,\cdot)}\big[\log p_\theta(X, Z)\big]\Big],$$

whenever the expectations are well-defined as extended real numbers (possibly equal to $-\infty$). Here $\kappa_{\theta'}(x, \cdot)$ denotes a regular conditional law of Z given $X = x$ under parameter θ' (constructed below), and p_θ denotes a chosen density under domination (Assumption A.1).

Extended-real integration conventions. For a measurable function $h : \Omega \to [-\infty, \infty]$ and a measure M on $(\Omega, \mathcal{A})$, the integral $\int h \, d\mathsf{M}$ is defined iff at least one of $\int h_+ \, d\mathsf{M}$ and $\int h_- \, d\mathsf{M}$ is finite; in that case $\int h \, d\mathsf{M} := \int h_+ \, d\mathsf{M} - \int h_- \, d\mathsf{M} \in [-\infty, \infty]$. We also use $0 \log 0 := 0$ and, when needed, $0 \log(0/q) := 0$ for $q > 0$.

A.1. **Dominating measures and versions of conditional laws.** We impose a standing domination/regularity condition ensuring densities and conditional laws can be manipulated via kernels.

Assumption A.1 (Joint domination and standard Borel latent space)**.** There exist σ-finite measures μ on $(\mathsf{X}, \mathcal{F})$ and ν on $(\mathsf{Z}, \mathcal{G})$ such that, for every $\theta \in \Theta$,

$$P_\theta^{XZ} \ll \mu \otimes \nu, \qquad \text{with density} \qquad p_\theta(x,z) := \frac{dP_\theta^{XZ}}{d(\mu\otimes\nu)}(x,z).$$

Moreover, $(\mathsf{Z}, \mathcal{G})$ is a standard Borel space.

A basic measurability fact (used repeatedly).

Lemma A.2 (Measurability of section integrals)**.** *Let $h : \mathsf{X} \times \mathsf{Z} \to [0,\infty]$ be $(\mathcal{F}\otimes\mathcal{G})$-measurable and let ν be σ-finite on $(\mathsf{Z}, \mathcal{G})$. Then the map*

$$x \longmapsto \int_{\mathsf{Z}} h(x,z)\,\nu(dz) \in [0,\infty]$$

is $\mathcal{F}$-measurable (as an extended real-valued function).

Proof. Let $\mathcal{H}$ be the class of all nonnegative measurable functions h such that $x \mapsto \int h(x,z)\nu(dz)$ is $\mathcal{F}$-measurable. We show $\mathcal{H}$ contains all nonnegative $(\mathcal{F}\otimes\mathcal{G})$-measurable functions.

Step 1 (rectangles). If $h = \mathbf{1}_{A\times B}$ with $A \in \mathcal{F}$, $B \in \mathcal{G}$, then

$$\int_{\mathsf{Z}} \mathbf{1}_{A\times B}(x,z)\nu(dz) = \mathbf{1}_A(x)\nu(B),$$

which is $\mathcal{F}$-measurable. Hence indicators of measurable rectangles lie in $\mathcal{H}$.

Step 2 (simple functions). By linearity, all finite linear combinations of rectangle indicators (i.e. nonnegative simple functions on $\mathsf{X}\times\mathsf{Z}$) lie in $\mathcal{H}$.

Step 3 (monotone limits). Suppose $h_n \in \mathcal{H}$ and $h_n \uparrow h$ pointwise. Then for every x,

$$\int h_n(x,z)\nu(dz) \uparrow \int h(x,z)\nu(dz)$$

by monotone convergence in z. The pointwise supremum of measurable functions is measurable, so $h \in \mathcal{H}$.

Step 4 (conclusion). Every nonnegative $(\mathcal{F}\otimes\mathcal{G})$-measurable h is the pointwise limit of an increasing sequence of nonnegative simple functions. By Steps 2–3, $h \in \mathcal{H}$. □

Marginal densities. Under Assumption A.1, the X-marginal is also dominated: $P_\theta^X \ll \mu$, with density

$$p_\theta(x) := \frac{dP_\theta^X}{d\mu}(x) = \int_{\mathsf{Z}} p_\theta(x,z)\,\nu(dz), \qquad \mu\text{-a.e. } x, \tag{202}$$

where the integral is well-defined in $[0,\infty]$.

Proof of (202). Fix $\theta \in \Theta$ and define $\tilde{p}_\theta(x) := \int_{\mathsf{Z}} p_\theta(x,z)\nu(dz) \in [0,\infty]$. By Lemma A.2, $\tilde{p}_\theta$ is $\mathcal{F}$-measurable.

Let $A \in \mathcal{F}$. Using the definition of the marginal and the Radon–Nikodym representation of P_θ^{XZ},

$$P_\theta^X(A) = P_\theta^{XZ}(A\times\mathsf{Z}) = \int_{A\times\mathsf{Z}} p_\theta(x,z)\,(\mu\otimes\nu)(dx,dz).$$

Since $p_\theta \geq 0$, Tonelli's theorem applies:

$$\int_{A\times \mathsf{Z}} p_\theta(x,z)\,(\mu\otimes\nu)(dx,dz) = \int_A \Big(\int_{\mathsf{Z}} p_\theta(x,z)\nu(dz)\Big)\mu(dx) = \int_A \tilde{p}_\theta(x)\,\mu(dx).$$

Thus $P_\theta^X(A) = \int_A \tilde{p}_\theta\, d\mu$ for all $A \in \mathcal{F}$, which implies $P_\theta^X \ll \mu$ with Radon–Nikodym derivative $dP_\theta^X/d\mu = \tilde{p}_\theta$ μ-a.e., i.e. (202). □

Regular conditional laws. We next construct a regular conditional probability of Z given X under P_θ^{XZ}. The construction is completely explicit (via a Borel isomorphism to a Borel subset of $[0,1]$ and conditional distribution functions), and yields the full disintegration identity.

Proposition A.3 (Existence of a regular conditional law)**.** *Assume $(\mathsf{Z},\mathcal{G})$ is standard Borel. Fix $\theta\in\Theta$ and let $(X,Z)\sim P_\theta^{XZ}$ on some probability space. Then there exists a Markov kernel*

$$\kappa_\theta :\ \mathsf{X}\times\mathcal{G}\to[0,1]$$

such that for all $A\in\mathcal{F}$ and $B\in\mathcal{G}$,

$$P_\theta^{XZ}(A\times B) \;=\; \int_A \kappa_\theta(x,B)\,P_\theta^X(dx). \tag{203}$$

Moreover, if $\tilde{\kappa}_\theta$ is another kernel satisfying (203), then for every $B\in\mathcal{G}$,

$$\kappa_\theta(x,B) = \tilde{\kappa}_\theta(x,B) \quad \text{for } P_\theta^X\text{-a.e. } x.$$

Proof. We proceed in five steps.

Step 1 (reduce to a Borel subset of $[0,1]$). Since $(\mathsf{Z},\mathcal{G})$ is standard Borel, there exists a Borel subset $E\subseteq[0,1]$ and a bimeasurable bijection (Borel isomorphism) $b:\mathsf{Z}\to E$ such that both b and b^{-1} are measurable. Define $Y := b(Z)$ taking values in $E\subseteq[0,1]$. Let $\mathcal{E}$ be the trace σ-field on E: $\mathcal{E} = \{E\cap B : B\in\mathcal{B}([0,1])\}$. Then Y is $(\mathcal{G},\mathcal{E})$-measurable and $\sigma(Y)=\sigma(Z)$.

It suffices to construct a regular conditional distribution of Y given X, say $\kappa_\theta^Y(x,\cdot)$ on $(E,\mathcal{E})$, and then define for $B\in\mathcal{G}$

$$\kappa_\theta(x,B) := \kappa_\theta^Y(x,b(B)).$$

Measurability and the disintegration identity then transfer by the bimeasurability of b.

Step 2 (conditional distribution functions on rational cutpoints). For each rational $q\in\mathbb{Q}\cap[0,1]$, define the random variable $\mathbf{1}_{\{Y\leq q\}}$ and take its conditional expectation given $\sigma(X)$:

$$g_q(X) \;:=\; \mathbb{E}\big[\mathbf{1}_{\{Y\leq q\}} \mid \sigma(X)\big].$$

By definition, $g_q(X)$ is $\sigma(X)$-measurable and satisfies, for all $A\in\sigma(X)$,

$$\mathbb{E}\big[\mathbf{1}_A g_q(X)\big] \;=\; \mathbb{E}\big[\mathbf{1}_A\mathbf{1}_{\{Y\leq q\}}\big]. \tag{204}$$

Since $g_q(X)$ is $\sigma(X)$-measurable, there exists an $\mathcal{F}$-measurable function $F_q:\mathsf{X}\to[0,1]$ such that

$$F_q(X) = g_q(X) \quad \text{a.s.}$$

(Indeed, $\sigma(X)$-measurability of $g_q(X)$ means $g_q(X)=\tilde{F}_q(X)$ for some $\tilde{F}_q$ measurable on X; choose one such version and call it F_q.)

Because $\mathbb{Q}\cap[0,1]$ is countable, we can (and do) choose the versions F_q simultaneously for all rationals.

Step 3 (enforce monotonicity and right-continuity in q pointwise in x). For rationals $q_1<q_2$, we have $\mathbf{1}_{\{Y\leq q_1\}}\leq\mathbf{1}_{\{Y\leq q_2\}}$, hence by monotonicity of conditional expectation

$$\mathbb{E}[\mathbf{1}_{\{Y\leq q_1\}}\mid\sigma(X)] \;\leq\; \mathbb{E}[\mathbf{1}_{\{Y\leq q_2\}}\mid\sigma(X)] \quad \text{a.s.}$$

Thus $F_{q_1}(X) \le F_{q_2}(X)$ a.s. For each pair (q_1, q_2) this inequality may fail on a null set; by countability of pairs of rationals, there exists a single P_θ^X-null set $N \subseteq \mathsf{X}$ such that for all $x \notin N$ and all rationals $q_1 < q_2$,

$$F_{q_1}(x) \le F_{q_2}(x). \tag{205}$$

For each $x \notin N$, define a function $F_x : [0,1] \to [0,1]$ by

$$F_x(t) \;:=\; \inf\{F_q(x) : q \in \mathbb{Q} \cap [0,1],\ q > t\}, \tag{206}$$

with the convention that the infimum over an empty set is 1 (which occurs only at $t = 1$). For $x \in N$, define $F_x(t) := \mathbf{1}_{\{t \ge 1\}}$ (any distribution function will do; this set will be null under P_θ^X).

For each fixed $x \notin N$, the function $t \mapsto F_x(t)$ is nondecreasing and right-continuous on $[0,1]$: - nondecreasing: if $t_1 < t_2$, then $\{q > t_2\} \subseteq \{q > t_1\}$, hence $\inf_{q>t_1} F_q(x) \le \inf_{q>t_2} F_q(x)$; - right-continuity: for $t_n \downarrow t$, the sets $\{q > t_n\}$ increase to $\{q > t\}$, so the infima decrease to the infimum over $\{q > t\}$, which equals $F_x(t)$ by definition.

Also, for $t < 0$ define $F_x(t) := 0$ and for $t > 1$ define $F_x(t) := 1$; then F_x is a distribution function on $\mathbb{R}$.

Step 4 (define a measure from F_x and verify measurability in x). For each $x \notin N$, define a probability measure $\kappa_\theta^Y(x, \cdot)$ on $([0,1], \mathcal{B}([0,1]))$ by the standard Lebesgue–Stieltjes construction:

$$\kappa_\theta^Y(x, (a,b]) \;:=\; F_x(b) - F_x(a), \qquad 0 \le a < b \le 1,$$

and extend uniquely to $\mathcal{B}([0,1])$ (Carathéodory extension theorem). Since F_x is right-continuous and nondecreasing with $F_x(0^-) = 0$ and $F_x(1) = 1$, the resulting measure is a probability measure. For $x \in N$, set $\kappa_\theta^Y(x, \cdot)$ equal to any fixed probability measure on $[0,1]$ (e.g. δ_0).

Now restrict to $E \subseteq [0,1]$ by defining for $C \in \mathcal{E}$,

$$\kappa_\theta^Y(x, C) \;:=\; \kappa_\theta^Y(x, \tilde{C}),$$

where $\tilde{C} \in \mathcal{B}([0,1])$ is any Borel set with $\tilde{C} \cap E = C$; this is well-defined because $\kappa_\theta^Y(x, \cdot)$ can be replaced by its restriction to E (equivalently, consider the pushforward measure on E).

We now verify the kernel measurability: for each $C \in \mathcal{E}$, the map $x \mapsto \kappa_\theta^Y(x, C)$ is $\mathcal{F}$-measurable. It suffices to check this on a π-system generating $\mathcal{E}$ and then extend by the monotone class theorem.

Let $\mathcal{I}$ be the collection of half-intervals in E of the form

$$(-\infty, t] \cap E, \qquad t \in \mathbb{Q} \cap [0,1].$$

This is a π-system generating $\mathcal{E}$ (since E is a Borel subset of $[0,1]$ and such intervals generate $\mathcal{B}([0,1])$). For $t \in \mathbb{Q} \cap [0,1]$ and $x \notin N$, we claim

$$\kappa_\theta^Y\big(x, (-\infty, t] \cap E\big) = F_t(x). \tag{207}$$

Indeed, by definition of F_x and right-continuity regularization, for rational t we have

$$F_x(t) = \inf_{q > t,\ q \in \mathbb{Q}} F_q(x) = F_t(x)$$

because monotonicity (205) implies $F_q(x) \downarrow F_t(x)$ as $q \downarrow t$ along rationals, hence the infimum equals the limit $F_t(x)$. For a distribution function F_x, the associated Stieltjes measure satisfies $\kappa_\theta^Y(x, (-\infty, t]) = F_x(t)$, so (207) holds.

Since each F_t is $\mathcal{F}$-measurable, the map $x \mapsto \kappa_\theta^Y(x, I)$ is measurable for each $I \in \mathcal{I}$. Let

$$\mathcal{M} := \{C \in \mathcal{E} :\ x \mapsto \kappa_\theta^Y(x, C) \text{ is } \mathcal{F}\text{-measurable}\}.$$

Then $\mathcal{M}$ is a monotone class: if $C_n \uparrow C$ with $C_n \in \mathcal{M}$, then $\kappa_\theta^Y(x, C_n) \uparrow \kappa_\theta^Y(x, C)$ (continuity from below), so measurability passes to the limit; similarly for $C_n \downarrow C$ using continuity from

above and boundedness by 1. Since $\mathcal{I} \subseteq \mathcal{M}$ and $\sigma(\mathcal{I}) = \mathcal{E}$, the monotone class theorem yields $\mathcal{M} = \mathcal{E}$. Thus $x \mapsto \kappa_\theta^Y(x, C)$ is measurable for all $C \in \mathcal{E}$, proving that κ_θ^Y is a kernel.

Step 5 (verify the defining disintegration identity). We show that for all $A \in \mathcal{F}$ and all $C \in \mathcal{E}$,

$$\mathbb{P}\big(X \in A,\ Y \in C\big) \;=\; \int_A \kappa_\theta^Y(x, C)\, P_\theta^X(dx). \tag{208}$$

It suffices to verify (208) on the generating π-system $\mathcal{I}$ and then extend by a λ-system argument.

Fix $t \in \mathbb{Q} \cap [0,1]$ and let $C_t := (-\infty, t] \cap E \in \mathcal{I}$. By (207),

$$\kappa_\theta^Y(x, C_t) = F_t(x) \quad \text{for } P_\theta^X\text{-a.e. } x.$$

Therefore, for any $A \in \mathcal{F}$,

$$\int_A \kappa_\theta^Y(x, C_t)\, P_\theta^X(dx) = \int_A F_t(x)\, P_\theta^X(dx) = \mathbb{E}\big[\mathbf{1}_{\{X \in A\}} F_t(X)\big].$$

By construction $F_t(X) = \mathbb{E}[\mathbf{1}_{\{Y \le t\}} \mid \sigma(X)]$ a.s., hence by the defining property of conditional expectation (204),

$$\mathbb{E}\big[\mathbf{1}_{\{X \in A\}} F_t(X)\big] = \mathbb{E}\big[\mathbf{1}_{\{X \in A\}} \mathbf{1}_{\{Y \le t\}}\big] = \mathbb{P}(X \in A,\ Y \in C_t).$$

Thus (208) holds for all $C_t \in \mathcal{I}$.

Now define

$$\Lambda := \{C \in \mathcal{E} :\ \text{(208) holds for all } A \in \mathcal{F}\}.$$

We show Λ is a λ-system: clearly $E \in \Lambda$ since both sides equal $P_\theta^X(A)$; if $C_1 \subseteq C_2$ with $C_1, C_2 \in \Lambda$, then by additivity of probability measures and of $x \mapsto \kappa_\theta^Y(x, \cdot)$,

$$\mathbb{P}(X \in A,\ Y \in C_2 \setminus C_1) = \mathbb{P}(X \in A,\ Y \in C_2) - \mathbb{P}(X \in A,\ Y \in C_1)$$

and similarly for the integral, so $C_2 \setminus C_1 \in \Lambda$; and if $C_n \uparrow C$, then both sides of (208) converge to the corresponding limit by monotone convergence (bounded by 1), hence $C \in \Lambda$. Since $\mathcal{I} \subseteq \Lambda$ and $\sigma(\mathcal{I}) = \mathcal{E}$, the π-λ theorem yields $\Lambda = \mathcal{E}$, establishing (208) for all $C \in \mathcal{E}$.

Finally, return to Z via b. For $B \in \mathcal{G}$, set $\kappa_\theta(x, B) := \kappa_\theta^Y(x, b(B))$. Then measurability in x and the disintegration identity transfer immediately:

$$\mathbb{P}(X \in A,\ Z \in B) = \mathbb{P}(X \in A,\ Y \in b(B)) = \int_A \kappa_\theta^Y(x, b(B))\, P_\theta^X(dx) = \int_A \kappa_\theta(x, B)\, P_\theta^X(dx),$$

which is (203).

Uniqueness. Let κ_θ and $\tilde{\kappa}_\theta$ satisfy (203). Fix $B \in \mathcal{G}$ and define $h(x) := \kappa_\theta(x, B) - \tilde{\kappa}_\theta(x, B)$, which is $\mathcal{F}$-measurable. For every $A \in \mathcal{F}$,

$$\int_A h(x)\, P_\theta^X(dx) = \int_A \kappa_\theta(x, B)\, P_\theta^X(dx) - \int_A \tilde{\kappa}_\theta(x, B)\, P_\theta^X(dx) = P_\theta^{XZ}(A \times B) - P_\theta^{XZ}(A \times B) = 0.$$

Taking $A = \{h > 0\}$ gives $\int_{\{h>0\}} h\, dP_\theta^X = 0$, but $h \ge 0$ on $\{h > 0\}$, hence $P_\theta^X(h > 0) = 0$. Similarly $P_\theta^X(h < 0) = 0$. Therefore $P_\theta^X(h \ne 0) = 0$, i.e. $\kappa_\theta(\cdot, B) = \tilde{\kappa}_\theta(\cdot, B)$ P_θ^X-a.s. □

Density-based version. When $p_\theta(x) \in (0, \infty)$, a canonical kernel is obtained by normalizing the joint density.

Lemma A.4 (Conditional density representation)**.** *Assume Assumption A.1. Fix $\theta \in \Theta$ and define $p_\theta(x)$ by (202). For x such that $p_\theta(x) \in (0, \infty)$ define*

$$p_\theta(z \mid x) \;:=\; \frac{p_\theta(x, z)}{p_\theta(x)}. \tag{209}$$

Then for μ-a.e. x with $p_\theta(x) \in (0, \infty)$:

(1) $z \mapsto p_\theta(z \mid x)$ *is* $\mathcal{G}$*-measurable, nonnegative, and* $\int p_\theta(z \mid x)\nu(dz) = 1$*;*
(2) $\kappa_\theta(x, B) := \int_B p_\theta(z \mid x)\nu(dz)$ *defines a kernel on* $\{p_\theta > 0\}$*;*
(3) extending $\kappa_\theta(x, \cdot)$ *arbitrarily on* $\{p_\theta = 0\}$ *yields a kernel satisfying* (203).

Proof. (1) Joint measurability of $(x, z) \mapsto p_\theta(x, z)$ and measurability of $x \mapsto p_\theta(x)$ imply measurability of $(x, z) \mapsto p_\theta(x, z)/p_\theta(x)$ on $\{p_\theta(x) > 0\}$; in particular for each fixed x with $p_\theta(x) > 0$, $z \mapsto p_\theta(z \mid x)$ is measurable and nonnegative. Moreover,

$$\int_{\mathsf{Z}} p_\theta(z \mid x)\nu(dz) = \frac{1}{p_\theta(x)} \int_{\mathsf{Z}} p_\theta(x, z)\nu(dz) = \frac{p_\theta(x)}{p_\theta(x)} = 1$$

by (202).

(2) For each $B \in \mathcal{G}$, the map $x \mapsto \int_B p_\theta(z \mid x)\nu(dz)$ is $\mathcal{F}$-measurable on $\{p_\theta > 0\}$ by Lemma A.2 applied to $h(x, z) = \mathbf{1}_B(z)p_\theta(z \mid x)$ (extend by 0 on $\{p_\theta = 0\}$ if desired). For each such x, $B \mapsto \kappa_\theta(x, B)$ is a probability measure since it is given by integrating a density.

(3) On $\{p_\theta = 0\}$ define $\kappa_\theta(x, \cdot)$ as any fixed probability measure; this preserves the kernel property. Since $P_\theta^X \ll \mu$ with density p_θ, we have $P_\theta^X(\{p_\theta = 0\}) = 0$. Finally, for $A \in \mathcal{F}$, $B \in \mathcal{G}$,

$$\int_A \kappa_\theta(x, B)\, P_\theta^X(dx) = \int_A \Big(\int_B \frac{p_\theta(x, z)}{p_\theta(x)} \nu(dz) \Big) p_\theta(x)\mu(dx) = \int_{A\times B} p_\theta(x, z)\, (\mu\otimes\nu)(dx, dz) = P_\theta^{XZ}(A\times B),$$

where Tonelli applies (nonnegative integrand). This is (203). □

Kernel–density equivalence. Conversely, specifying a marginal density and a conditional density kernel reconstructs a dominated joint law.

Lemma A.5 (Reconstruction from marginal and conditional)**.** *Fix* θ *and suppose we are given: (i) a nonnegative measurable* $p_\theta : \mathsf{X} \to [0, \infty)$ *with* $\int p_\theta \, d\mu = 1$*; and (ii) a measurable function* $p_\theta(\cdot \mid \cdot) : \mathsf{X} \times \mathsf{Z} \to [0, \infty)$ *such that for each* x*,* $z \mapsto p_\theta(z \mid x)$ *is* $\mathcal{G}$*-measurable and* $\int p_\theta(z \mid x)\nu(dz) = 1$*. Define*

$$p_\theta(x, z) \;:=\; p_\theta(x)\, p_\theta(z \mid x). \tag{210}$$

Then $p_\theta(x, z)$ *is a* $(\mu \otimes \nu)$*-density, hence defines a probability law* $P_\theta^{XZ} \ll \mu \otimes \nu$*. Moreover: (a) its* X*-marginal has density* $p_\theta(x)$*; and (b) a regular conditional law of* Z *given* $X = x$ *is* $\kappa_\theta(x, dz) = p_\theta(z \mid x)\nu(dz)$*.*

Proof. Measurability and nonnegativity of $p_\theta(x, z)$ are immediate. Its total mass is

$$\int_{\mathsf{X}\times\mathsf{Z}} p_\theta(x, z)\, (\mu \otimes \nu)(dx, dz) = \int_{\mathsf{X}} p_\theta(x) \Big(\int_{\mathsf{Z}} p_\theta(z \mid x)\nu(dz) \Big) \mu(dx) = \int_{\mathsf{X}} p_\theta(x)\mu(dx) = 1,$$

by Tonelli. Thus p_θ is a probability density on $\mathsf{X} \times \mathsf{Z}$.

For the marginal, for any $A \in \mathcal{F}$,

$$P_\theta^X(A) = P_\theta^{XZ}(A \times \mathsf{Z}) = \int_{A\times\mathsf{Z}} p_\theta(x)p_\theta(z \mid x)\, (\mu \otimes \nu)(dx, dz) = \int_A p_\theta(x)\mu(dx),$$

since the inner integral over Z is 1. Hence the marginal density is $p_\theta(x)$.

For the conditional law, define $\kappa_\theta(x, B) := \int_B p_\theta(z \mid x)\nu(dz)$. Then for all $A \in \mathcal{F}$, $B \in \mathcal{G}$,

$$\int_A \kappa_\theta(x, B)\, P_\theta^X(dx) = \int_A \Big(\int_B p_\theta(z \mid x)\nu(dz) \Big) p_\theta(x)\mu(dx) = \int_{A\times B} p_\theta(x)p_\theta(z \mid x)\, (\mu\otimes\nu)(dx, dz) = P_\theta^{XZ}(A\times B),$$

which is exactly the disintegration identity (203). □

A.2. **Interchange of integrals, limits, and differentiation.** This subsection records the routine interchange steps used repeatedly for Φ and Q.

Kernel integrals and joint measures.

Definition A.6 (Kernel)**.** A map $\kappa : \mathsf{X} \times \mathcal{G} \to [0,1]$ is a (Markov) kernel from $(\mathsf{X}, \mathcal{F})$ to $(\mathsf{Z}, \mathcal{G})$ if (i) for each $x \in \mathsf{X}$, $\kappa(x, \cdot)$ is a probability measure on $(\mathsf{Z}, \mathcal{G})$, and (ii) for each $B \in \mathcal{G}$, $x \mapsto \kappa(x, B)$ is $\mathcal{F}$-measurable.

Lemma A.7 (Measurability of kernel integrals)**.** *Let κ be a kernel and $h : \mathsf{X} \times \mathsf{Z} \to [0, \infty]$ be $(\mathcal{F} \otimes \mathcal{G})$-measurable. Then the map*

$$x \;\longmapsto\; \int_{\mathsf{Z}} h(x,z)\, \kappa(x, dz) \in [0, \infty]$$

is $\mathcal{F}$-measurable.

Proof. Let $\mathcal{H}$ be the set of all nonnegative measurable h for which $x \mapsto \int h(x,z)\kappa(x,dz)$ is measurable. We verify $\mathcal{H}$ contains a generating class and is closed under monotone limits.

If $h = \mathbf{1}_{A \times B}$ with $A \in \mathcal{F}$, $B \in \mathcal{G}$, then

$$\int \mathbf{1}_{A\times B}(x,z)\kappa(x,dz) = \mathbf{1}_A(x)\kappa(x,B),$$

which is measurable by Definition A.6. Hence rectangle indicators lie in $\mathcal{H}$, and by linearity, all nonnegative simple functions lie in $\mathcal{H}$. If $h_n \uparrow h$ with $h_n \in \mathcal{H}$, then by monotone convergence (in z),

$$\int h_n(x,z)\kappa(x,dz) \uparrow \int h(x,z)\kappa(x,dz)$$

pointwise in x, and the limit of measurable functions is measurable. Thus $\mathcal{H}$ is a monotone class containing all rectangle indicators, hence contains all nonnegative measurable functions. □

Definition A.8 (Joint measure induced by a marginal and a kernel)**.** Let P be a probability measure on $(\mathsf{X}, \mathcal{F})$ and κ a kernel from X to Z. Define a set function on rectangles by

$$(P \otimes \kappa)(A \times B) \;:=\; \int_A \kappa(x, B)\, P(dx), \qquad A \in \mathcal{F},\ B \in \mathcal{G}.$$

Then $P \otimes \kappa$ extends uniquely to a probability measure on $(\mathsf{X} \times \mathsf{Z}, \mathcal{F} \otimes \mathcal{G})$.

Proof. We must verify that the rectangle set function is a pre-measure on the algebra generated by rectangles. Fix $B \in \mathcal{G}$. The map $A \mapsto \int_A \kappa(x,B)P(dx)$ is a (finite) measure on $\mathcal{F}$, because it is integration of the nonnegative measurable function $\kappa(\cdot, B)$. Linearity in B is inherited pointwise from the measure property of $\kappa(x, \cdot)$. Standard arguments (check countable additivity on disjoint unions of rectangles and extend by Carathéodory) yield a unique measure on $\mathcal{F} \otimes \mathcal{G}$. Total mass is

$$(P \otimes \kappa)(\mathsf{X} \times \mathsf{Z}) = \int_{\mathsf{X}} \kappa(x, \mathsf{Z})\, P(dx) = \int_{\mathsf{X}} 1\, P(dx) = 1,$$

so the extension is a probability measure. □

Lemma A.9 (Tonelli/Fubini for nested expectations)**.** *Let $h : \mathsf{X} \times \mathsf{Z} \to [-\infty, \infty]$ be $(\mathcal{F} \otimes \mathcal{G})$-measurable, let P be a probability measure on $(\mathsf{X}, \mathcal{F})$, and let κ be a kernel from X to Z.*

(1) If $h \geq 0$, then

$$\int_{\mathsf{X}\times\mathsf{Z}} h(x,z)\, (P \otimes \kappa)(dx, dz) = \int_{\mathsf{X}} \int_{\mathsf{Z}} h(x,z)\, \kappa(x, dz)\, P(dx) \in [0, \infty].$$

(2) If $\int_{\mathsf{X}\times\mathsf{Z}} |h|\, d(P\otimes\kappa) < \infty$, *then*

$$\int_{\mathsf{X}\times\mathsf{Z}} h\, d(P\otimes\kappa) = \int_{\mathsf{X}}\int_{\mathsf{Z}} h(x,z)\,\kappa(x,dz)\,P(dx)$$

and the integral is finite. In particular, iterated integration is unambiguous and standard Fubini manipulations are valid.

Proof. (1) Let $\mathcal{H}$ be the set of all nonnegative measurable h for which the identity holds. For $h = \mathbf{1}_{A\times B}$,

$$\int_{\mathsf{X}\times\mathsf{Z}} \mathbf{1}_{A\times B}\, d(P\otimes\kappa) = (P\otimes\kappa)(A\times B) = \int_A \kappa(x,B)P(dx) = \int_{\mathsf{X}}\int_{\mathsf{Z}} \mathbf{1}_{A\times B}(x,z)\kappa(x,dz)P(dx),$$

so rectangle indicators are in $\mathcal{H}$. By linearity, nonnegative simple functions are in $\mathcal{H}$.

If $h_n \uparrow h$ pointwise with $h_n \in \mathcal{H}$, then by monotone convergence applied to the measure $P\otimes\kappa$,

$$\int h_n\, d(P\otimes\kappa) \uparrow \int h\, d(P\otimes\kappa),$$

and by monotone convergence applied inside the iterated integral (nonnegative integrand),

$$\int_{\mathsf{X}}\int_{\mathsf{Z}} h_n(x,z)\kappa(x,dz)P(dx) \uparrow \int_{\mathsf{X}}\int_{\mathsf{Z}} h(x,z)\kappa(x,dz)P(dx).$$

Thus the identity passes to the limit, so $\mathcal{H}$ is closed under monotone limits. Since every nonnegative measurable function is an increasing limit of nonnegative simple functions, $\mathcal{H}$ contains all $h \geq 0$.

(2) Write $h = h_+ - h_-$ with $h_\pm \geq 0$ measurable. If $\int |h|\, d(P\otimes\kappa) < \infty$, then both $\int h_+\, d(P\otimes\kappa)$ and $\int h_-\, d(P\otimes\kappa)$ are finite. By part (1) applied to h_+ and h_-,

$$\int h_\pm\, d(P\otimes\kappa) = \int_{\mathsf{X}}\int_{\mathsf{Z}} h_\pm(x,z)\kappa(x,dz)P(dx).$$

Subtracting yields the desired identity for h, and finiteness follows. □

Lemma A.10 (Dominated convergence under nested expectations)**.** *Let* $h_n, h : \mathsf{X}\times\mathsf{Z} \to \mathbb{R}$ *be measurable with* $h_n(x,z) \to h(x,z)$ *pointwise. Let* P *be a probability on* $(\mathsf{X},\mathcal{F})$ *and* κ *a kernel from* X *to* Z*. Assume there exists an envelope* $H : \mathsf{X}\times\mathsf{Z} \to [0,\infty)$ *such that* $|h_n| \leq H$ *for all* n *and*

$$\int_{\mathsf{X}\times\mathsf{Z}} H(x,z)\,(P\otimes\kappa)(dx,dz) \;<\; \infty.$$

Then

$$\int_{\mathsf{X}}\int_{\mathsf{Z}} h_n(x,z)\,\kappa(x,dz)\,P(dx) \;\to\; \int_{\mathsf{X}}\int_{\mathsf{Z}} h(x,z)\,\kappa(x,dz)\,P(dx).$$

Proof. By Lemma A.9(2), since H is integrable under $P\otimes\kappa$, each h_n and h is integrable and

$$\int_{\mathsf{X}}\int_{\mathsf{Z}} h_n\, d\kappa\, dP = \int h_n\, d(P\otimes\kappa), \qquad \int_{\mathsf{X}}\int_{\mathsf{Z}} h\, d\kappa\, dP = \int h\, d(P\otimes\kappa).$$

Apply the standard dominated convergence theorem to the probability space $(\mathsf{X}\times\mathsf{Z}, \mathcal{F}\otimes\mathcal{G}, P\otimes\kappa)$ with dominating function H to obtain $\int h_n\, d(P\otimes\kappa) \to \int h\, d(P\otimes\kappa)$. Translating back yields the claimed convergence of iterated integrals. □

Differentiation under marginalization. We state a vector-valued Leibniz rule sufficient for differentiating $p_\theta(x) = \int p_\theta(x,z)\nu(dz)$.

Assumption A.11 (Differentiability and domination of the joint density gradient)**.** For $(\mu\otimes\nu)$-a.e. (x,z), the map $\theta \mapsto p_\theta(x,z)$ is differentiable on Θ. Moreover, for every compact $K \subset \Theta$ there exists a measurable envelope $G_K : \mathsf{X}\times\mathsf{Z} \to [0,\infty)$ such that

$$\sup_{\theta\in K} \|\nabla_\theta p_\theta(x,z)\| \;\le\; G_K(x,z) \quad \text{for } (\mu\otimes\nu)\text{-a.e. } (x,z),$$

and such that for μ-a.e. x,

$$\int_{\mathsf{Z}} G_K(x,z)\,\nu(dz) \;<\; \infty.$$

Lemma A.12 (Differentiation of the marginal density)**.** *Under Assumption A.11, for every compact $K \subset \Theta$ and μ-a.e. x, the map $\theta \mapsto p_\theta(x)$ is differentiable on K, and*

$$\nabla_\theta p_\theta(x) \;=\; \int_{\mathsf{Z}} \nabla_\theta p_\theta(x,z)\,\nu(dz), \qquad \theta \in K.$$

Proof. Fix a compact $K \subset \Theta$, a point $\theta \in K$, and a coordinate $j \in \{1,\dots,d\}$. Let e_j be the jth standard basis vector. For $t \neq 0$ such that $\theta + te_j \in K$, define

$$D_{t,j}(x) \;:=\; \frac{p_{\theta+te_j}(x) - p_\theta(x)}{t}.$$

Using (202) and linearity of the integral,

$$D_{t,j}(x) = \int_{\mathsf{Z}} \frac{p_{\theta+te_j}(x,z) - p_\theta(x,z)}{t}\,\nu(dz).$$

For $(\mu\otimes\nu)$-a.e. (x,z), differentiability of $\theta \mapsto p_\theta(x,z)$ implies

$$\frac{p_{\theta+te_j}(x,z) - p_\theta(x,z)}{t} \;\to\; \partial_{\theta_j} p_\theta(x,z) \qquad \text{as } t \to 0.$$

We now produce an integrable dominating function in z for μ-a.e. x. For such (x,z) and sufficiently small t (so that the segment between θ and $\theta + te_j$ lies in K), the one-dimensional mean value theorem gives

$$\frac{p_{\theta+te_j}(x,z) - p_\theta(x,z)}{t} = \partial_{\theta_j} p_{\theta+s_te_j}(x,z)$$

for some s_t between 0 and t. Hence

$$\Big|\frac{p_{\theta+te_j}(x,z) - p_\theta(x,z)}{t}\Big| \le \sup_{\vartheta\in K} \big|\partial_{\theta_j} p_\vartheta(x,z)\big| \le \sup_{\vartheta\in K} \|\nabla_\theta p_\vartheta(x,z)\| \le G_K(x,z),$$

for $(\mu\otimes\nu)$-a.e. (x,z). By assumption, for μ-a.e. x, $\int G_K(x,z)\nu(dz) < \infty$, so dominated convergence (in z) yields, for such x,

$$\lim_{t\to 0} D_{t,j}(x) = \int_{\mathsf{Z}} \partial_{\theta_j} p_\theta(x,z)\,\nu(dz).$$

Thus $\partial_{\theta_j} p_\theta(x)$ exists and equals the integral of $\partial_{\theta_j} p_\theta(x,z)$, for μ-a.e. x. Doing this for each j yields the vector identity for $\nabla_\theta p_\theta(x)$. □

Differentiating Φ and Q. To differentiate log-likelihood terms under $P^\star$ and under the joint measure $P^\star \otimes \kappa_{\theta'}$, we impose local score envelopes.

Assumption A.13 (Positivity and score domination)**.** Fix a compact set $K \subset \Theta$.

(1) (*Positivity*) For each $\theta \in K$, $p_\theta(x) \in (0,\infty)$ for $P^\star$-a.e. x; and for each $\theta \in K$, $p_\theta(x,z) \in (0,\infty)$ for $(P^\star \otimes \kappa_{\theta'})(dx,dz)$-a.e. (x,z) whenever $Q(\theta \mid \theta')$ is evaluated.
(2) (*Differentiability*) For $P^\star$-a.e. x, $\theta \mapsto \log p_\theta(x)$ is differentiable on K; and for $(P^\star \otimes \kappa_{\theta'})(dx,dz)$-a.e. (x,z), $\theta \mapsto \log p_\theta(x,z)$ is differentiable on K.
(3) (*Integrable envelopes*) There exist measurable $S_K : \mathsf{X} \to [0,\infty)$ and $T_K : \mathsf{X} \times \mathsf{Z} \to [0,\infty)$ such that

$$\sup_{\theta \in K} \|\nabla_\theta \log p_\theta(x)\| \;\le\; S_K(x) \quad \text{for } P^\star\text{-a.e. } x,$$

$$\sup_{\theta \in K} \|\nabla_\theta \log p_\theta(x,z)\| \;\le\; T_K(x,z) \quad \text{for } (P^\star \otimes \kappa_{\theta'})\text{-a.e. } (x,z),$$

and $\mathbb{E}_{P^\star}[S_K(X)] < \infty$ and $\int T_K \, d(P^\star \otimes \kappa_{\theta'}) < \infty$.

Lemma A.14 (Differentiation of Φ)**.** *Fix a compact $K \subset \Theta$ and assume Assumption A.13 holds for this K. Then Φ is differentiable on K and, for each $\theta \in K$,*

$$\nabla_\theta \Phi(\theta) \;=\; \mathbb{E}_{P^\star}\big[\nabla_\theta \log p_\theta(X)\big].$$

Proof. Fix $\theta \in K$ and coordinate j. For $t \neq 0$ small enough so that $\theta + te_j \in K$, define

$$\Delta_{t,j}(x) \;:=\; \frac{\log p_{\theta+te_j}(x) - \log p_\theta(x)}{t}.$$

By Assumption A.13(2), $\Delta_{t,j}(x) \to \partial_{\theta_j} \log p_\theta(x)$ for $P^\star$-a.e. x.

We now dominate $\Delta_{t,j}$ by an integrable envelope. For $P^\star$-a.e. x, the map $s \mapsto \log p_{\theta+se_j}(x)$ is differentiable for s in a neighborhood of 0 (since K is compact and θ is an interior point of the segment for small t). By the mean value theorem applied to $s \mapsto \log p_{\theta+se_j}(x)$ on $[0,t]$ (or $[t,0]$ if $t<0$), there exists $s_t(x)$ between 0 and t such that

$$\Delta_{t,j}(x) = \partial_{\theta_j} \log p_{\theta+s_t(x)e_j}(x).$$

Therefore

$$|\Delta_{t,j}(x)| \le \|\nabla_\theta \log p_{\theta+s_t(x)e_j}(x)\| \le \sup_{\vartheta \in K} \|\nabla_\theta \log p_\vartheta(x)\| \le S_K(x)$$

for $P^\star$-a.e. x. Since $\mathbb{E}_{P^\star}[S_K(X)] < \infty$, dominated convergence yields

$$\lim_{t \to 0} \mathbb{E}_{P^\star}[\Delta_{t,j}(X)] = \mathbb{E}_{P^\star}[\partial_{\theta_j} \log p_\theta(X)].$$

Finally, by definition of Φ,

$$\mathbb{E}_{P^\star}[\Delta_{t,j}(X)] = \frac{1}{t}\Big(\mathbb{E}_{P^\star}[\log p_{\theta+te_j}(X)] - \mathbb{E}_{P^\star}[\log p_\theta(X)]\Big) = \frac{\Phi(\theta+te_j) - \Phi(\theta)}{t}.$$

Thus $\partial_{\theta_j}\Phi(\theta)$ exists and equals $\mathbb{E}_{P^\star}[\partial_{\theta_j} \log p_\theta(X)]$. Since this holds for all j, Φ is differentiable at θ with the stated gradient. As $\theta \in K$ was arbitrary, Φ is differentiable on K. □

Lemma A.15 (Differentiation of Q)**.** *Fix a compact $K \subset \Theta$ and assume Assumption A.13 holds for this K. Then for each $\theta \in K$,*

$$\nabla_\theta Q(\theta \mid \theta') \;=\; \mathbb{E}_{P^\star}\Big[\mathbb{E}_{\kappa_{\theta'}(X,\cdot)}\big[\nabla_\theta \log p_\theta(X,Z)\big]\Big],$$

provided the right-hand side exists in $\mathbb{R}^d$ (equivalently, $\int \|\nabla_\theta \log p_\theta\| \, d(P^\star \otimes \kappa_{\theta'}) < \infty$).

Proof. Fix $\theta \in K$ and coordinate j. For small $t \neq 0$ with $\theta + te_j \in K$, define

$$\Delta_{t,j}(x,z) := \frac{\log p_{\theta+te_j}(x,z) - \log p_\theta(x,z)}{t}.$$

By Assumption A.13(2), $\Delta_{t,j}(x,z) \to \partial_{\theta_j} \log p_\theta(x,z)$ for $(P^\star \otimes \kappa_{\theta'})$-a.e. (x,z).

By the mean value theorem applied pointwise in (x,z) to $s \mapsto \log p_{\theta+se_j}(x,z)$, there exists $s_t(x,z)$ between 0 and t such that

$$\Delta_{t,j}(x,z) = \partial_{\theta_j} \log p_{\theta+s_t(x,z)e_j}(x,z),$$

and hence

$$|\Delta_{t,j}(x,z)| \leq \sup_{\vartheta \in K} \|\nabla_\theta \log p_\vartheta(x,z)\| \leq T_K(x,z)$$

for $(P^\star \otimes \kappa_{\theta'})$-a.e. (x,z). The envelope T_K is integrable under $P^\star \otimes \kappa_{\theta'}$ by Assumption A.13(3). Therefore dominated convergence (Lemma A.10 applied under the joint measure $P^\star \otimes \kappa_{\theta'}$) yields

$$\lim_{t\to 0} \int \Delta_{t,j}(x,z)\,(P^\star \otimes \kappa_{\theta'})(dx,dz) = \int \partial_{\theta_j} \log p_\theta(x,z)\,(P^\star \otimes \kappa_{\theta'})(dx,dz).$$

Using Lemma A.9(2), the left-hand side equals

$$\frac{1}{t}\Big(Q(\theta + te_j \mid \theta') - Q(\theta \mid \theta')\Big),$$

since $Q(\cdot \mid \theta')$ is defined as an integral of $\log p_\theta(x,z)$ under $P^\star \otimes \kappa_{\theta'}$. Thus $\partial_{\theta_j} Q(\theta \mid \theta')$ exists and equals

$$\int \partial_{\theta_j} \log p_\theta(x,z)\,(P^\star \otimes \kappa_{\theta'})(dx,dz) = \mathbb{E}_{P^\star}\mathbb{E}_{\kappa_{\theta'}(X,\cdot)}\big[\partial_{\theta_j} \log p_\theta(X,Z)\big].$$

Collecting coordinates gives the claimed gradient formula. □

Fisher identity.

Lemma A.16 (Fisher identity)**.** *Assume Assumption A.1. Fix $\theta \in \Theta$ and suppose:*

(1) $p_\theta(x) \in (0,\infty)$ for P_θ^X-a.e. x;
(2) Assumption A.11 holds on some compact neighborhood $K \ni \theta$;
(3) $\int \|\nabla_\theta \log p_\theta(x,z)\|\,\kappa_\theta(x,dz) < \infty$ for P_θ^X-a.e. x, where $\kappa_\theta(x,dz) = p_\theta(z \mid x)\nu(dz)$ as in Lemma A.4.

Then for P_θ^X-a.e. x,

$$\nabla_\theta \log p_\theta(x) = \mathbb{E}_{\kappa_\theta(x,\cdot)}\big[\nabla_\theta \log p_\theta(x,Z)\big].$$

Proof. By Lemma A.12, for μ-a.e. x (hence for P_θ^X-a.e. x since $P_\theta^X \ll \mu$),

$$\nabla_\theta p_\theta(x) = \int_{\mathsf{Z}} \nabla_\theta p_\theta(x,z)\,\nu(dz).$$

Fix such an x with $p_\theta(x) > 0$. Divide by $p_\theta(x)$:

$$\nabla_\theta \log p_\theta(x) = \frac{\nabla_\theta p_\theta(x)}{p_\theta(x)} = \int_{\mathsf{Z}} \frac{\nabla_\theta p_\theta(x,z)}{p_\theta(x)}\,\nu(dz).$$

Multiply and divide inside the integral by $p_\theta(x,z)$ (which is positive ν-a.e. under the conditional law by construction of κ_θ):

$$\frac{\nabla_\theta p_\theta(x,z)}{p_\theta(x)} = \frac{\nabla_\theta p_\theta(x,z)}{p_\theta(x,z)} \cdot \frac{p_\theta(x,z)}{p_\theta(x)} = \nabla_\theta \log p_\theta(x,z) \cdot p_\theta(z \mid x).$$

Hence

$$\nabla_\theta \log p_\theta(x) = \int_{\mathsf{Z}} \nabla_\theta \log p_\theta(x,z)\, p_\theta(z \mid x)\, \nu(dz) = \int_{\mathsf{Z}} \nabla_\theta \log p_\theta(x,z)\, \kappa_\theta(x,dz) = \mathbb{E}_{\kappa_\theta(x,\cdot)}[\nabla_\theta \log p_\theta(x,Z)],$$

where the final expression is well-defined by the assumed integrability of the complete-data score under $\kappa_\theta(x,\cdot)$. □

A.3. **Auxiliary inequalities.** We collect inequalities used repeatedly, with full proofs.

Lemma A.17 (Jensen's inequality)**.** *Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space, let $U : \Omega \to \mathbb{R}$ be integrable, and let $\varphi : \mathbb{R} \to (-\infty, \infty]$ be convex. Then*

$$\varphi(\mathbb{E}[U]) \;\le\; \mathbb{E}[\varphi(U)]$$

whenever $\mathbb{E}[\varphi(U)]$ is well-defined in $(-\infty, \infty]$. In particular, since log *is concave, for any nonnegative W with $\mathbb{E}[W] \in (0, \infty)$,*

$$\mathbb{E}[\log W] \;\le\; \log \mathbb{E}[W],$$

allowing $\mathbb{E}[\log W]$ to equal $-\infty$.

Proof. If $\varphi(\mathbb{E}[U]) = +\infty$, the inequality is trivial. Assume $\varphi(\mathbb{E}[U]) < \infty$ and set $m := \mathbb{E}[U] \in \mathbb{R}$.

A real-valued convex function on $\mathbb{R}$ has finite left and right derivatives at every point in the interior of its effective domain, and admits at least one subgradient at such points. Concretely, define the left and right slopes

$$\varphi'_-(m) := \lim_{t \uparrow 0} \frac{\varphi(m+t) - \varphi(m)}{t}, \qquad \varphi'_+(m) := \lim_{t \downarrow 0} \frac{\varphi(m+t) - \varphi(m)}{t},$$

which exist in $[-\infty, \infty]$ with $\varphi'_-(m) \le \varphi'_+(m)$. Choose any $s \in [\varphi'_-(m), \varphi'_+(m)] \cap \mathbb{R}$ (such an s exists whenever $\varphi(m) < \infty$). Then the supporting line inequality holds:

$$\varphi(u) \;\ge\; \varphi(m) + s(u - m) \qquad \text{for all } u \in \mathbb{R}. \tag{211}$$

(To verify (211): for $u > m$, convexity implies $(\varphi(u) - \varphi(m))/(u - m) \ge \varphi'_-(m) \ge s$; for $u < m$, convexity implies $(\varphi(u) - \varphi(m))/(u - m) \le \varphi'_+(m) \le s$, which rearranges to the same inequality.)

Apply (211) to $u = U(\omega)$ and take expectations:

$$\mathbb{E}[\varphi(U)] \;\ge\; \mathbb{E}[\varphi(m) + s(U - m)] = \varphi(m) + s(\mathbb{E}[U] - m) = \varphi(\mathbb{E}[U]).$$

This proves Jensen.

For the log specialization, apply Jensen to the convex function $-\log$ on $(0, \infty)$ (or equivalently, use concavity of log) with $U = W$. □

Lemma A.18 (Elementary log bounds)**.** *For all $t > 0$,*

$$\log t \;\le\; t - 1, \qquad \textit{equivalently} \qquad -\log t \;\ge\; 1 - t.$$

More generally, for any $t_0 > 0$,

$$\log t \;\le\; \log t_0 \;+\; \frac{t - t_0}{t_0}.$$

Proof. Define $f(t) := t - 1 - \log t$ for $t > 0$. Then $f'(t) = 1 - \frac{1}{t}$ and $f''(t) = \frac{1}{t^2} > 0$, so f is strictly convex. Its unique minimizer satisfies $f'(t) = 0$, i.e. $t = 1$, and $f(1) = 0$. Hence $f(t) \ge 0$ for all $t > 0$, which is $\log t \le t - 1$.

For the supporting-hyperplane inequality, concavity of log implies that for any $t_0 > 0$,

$$\log t \le \log t_0 + \log'(t_0)(t - t_0) = \log t_0 + \frac{t - t_0}{t_0}.$$

□

Lemma A.19 (KL nonnegativity)**.** *Let P, Q be probability measures on a measurable space $(\Omega, \mathcal{A})$ with $P \ll Q$. Define*

$$\mathrm{KL}(P\|Q) \;:=\; \int_\Omega \log\Big(\frac{dP}{dQ}\Big)\, dP \;\in [0,\infty].$$

Then $\mathrm{KL}(P\|Q) \geq 0$, with equality if and only if $P = Q$.

Proof. Let $r := \frac{dP}{dQ}$, so $r \geq 0$ and $\int r\, dQ = 1$. Then

$$\mathrm{KL}(P\|Q) = \int r \log r\, dQ,$$

with the convention $0 \log 0 := 0$.

Apply Lemma A.18 to $t = r$ to obtain $\log r \geq 1 - \frac{1}{r}$ for $r > 0$, equivalently

$$r \log r \geq r - 1 \qquad \text{for all } r \geq 0,$$

where the $r = 0$ case holds by continuity ($\lim_{r\downarrow 0} r \log r = 0$ and $r - 1 = -1$). Integrate both sides with respect to Q:

$$\mathrm{KL}(P\|Q) = \int r \log r\, dQ \geq \int (r-1)\, dQ = \int r\, dQ - 1 = 0.$$

If $\mathrm{KL}(P\|Q) = 0$, then $\int (r \log r - (r-1))\, dQ = 0$ with a nonnegative integrand. Thus

$$r \log r - (r-1) = 0 \quad Q\text{-a.e.}$$

But the function $g(r) := r \log r - (r-1)$ is strictly convex on $(0,\infty)$, satisfies $g(r) \geq 0$, and $g(r) = 0$ iff $r = 1$. Hence $r = 1$ Q-a.e., which implies $P = Q$. □

Lemma A.20 (Integral log-sum inequality)**.** *Let $(\Omega, \mathcal{A}, \lambda)$ be a σ-finite measure space and let $a, b : \Omega \to [0,\infty)$ be measurable with*

$$A := \int a\, d\lambda \in (0,\infty), \qquad B := \int b\, d\lambda \in (0,\infty).$$

Then

$$\int a \log \frac{a}{b}\, d\lambda \;\geq\; A \log \frac{A}{B},$$

with the conventions $0 \log(0/b) := 0$ and the inequality interpreted as valid if the left-hand side is $+\infty$.

Proof. Define probability measures P and Q on $(\Omega, \mathcal{A})$ by

$$P(d\omega) = \frac{a(\omega)}{A}\, \lambda(d\omega), \qquad Q(d\omega) = \frac{b(\omega)}{B}\, \lambda(d\omega).$$

Then P, Q are probability measures (since $\int a\, d\lambda = A$ and $\int b\, d\lambda = B$). If $P \not\ll Q$, then $\mathrm{KL}(P\|Q) = +\infty$ by definition, and the desired inequality is immediate. Assume $P \ll Q$.

Compute the Radon–Nikodym derivative:

$$\frac{dP}{dQ}(\omega) = \frac{a(\omega)/A}{b(\omega)/B} = \frac{B}{A} \cdot \frac{a(\omega)}{b(\omega)} \qquad Q\text{-a.e.}$$

Therefore,

$$\mathrm{KL}(P\|Q) = \int \log\Big(\frac{dP}{dQ}\Big)\, dP = \int \log\Big(\frac{B}{A} \cdot \frac{a}{b}\Big)\, dP = \int \log\Big(\frac{a}{b}\Big)\, dP + \log\Big(\frac{B}{A}\Big) \int 1\, dP.$$

Since $\int 1\,dP = 1$ and $dP = (a/A)\,d\lambda$, this becomes

$$\mathrm{KL}(P\|Q) = \frac{1}{A}\int a \log\Big(\frac{a}{b}\Big)\,d\lambda + \log\Big(\frac{B}{A}\Big).$$

By Lemma A.19, $\mathrm{KL}(P\|Q) \geq 0$, hence

$$\frac{1}{A}\int a \log\Big(\frac{a}{b}\Big)\,d\lambda \geq \log\Big(\frac{A}{B}\Big),$$

which is the claimed inequality after multiplying by A. □

Lemma A.21 (KL chain rule and measurability of conditional KL)**.** *Assume $(\mathsf{Z}, \mathcal{G})$ is standard Borel. Let P, Q be probability measures on $(\mathsf{X} \times \mathsf{Z}, \mathcal{F} \otimes \mathcal{G})$ with $P \ll Q$. Let P^X, Q^X be the X-marginals, and let $P(\cdot \mid x), Q(\cdot \mid x)$ be regular conditional laws of Z given $X = x$ under P, Q, respectively. Then:*

(1) The map $x \mapsto \mathrm{KL}(P(\cdot \mid x)\|Q(\cdot \mid x))$ is $\mathcal{F}$-measurable as an extended real-valued function.
(2) The chain rule holds:

$$\mathrm{KL}(P\|Q) \;=\; \mathrm{KL}(P^X\|Q^X) \;+\; \mathbb{E}_{P^X}\big[\mathrm{KL}(P(\cdot \mid X)\|Q(\cdot \mid X))\big],$$

with both sides possibly equal to $+\infty$.

Proof. Step 1 (disintegrations and marginal absolute continuity). Since Z is standard Borel, regular conditional laws exist, so we can write the disintegrations

$$P(dx, dz) = P^X(dx)\,P(dz \mid x), \qquad Q(dx, dz) = Q^X(dx)\,Q(dz \mid x).$$

From $P \ll Q$ it follows that $P^X \ll Q^X$: if $Q^X(A) = 0$, then $Q(A \times \mathsf{Z}) = 0$, hence $P(A \times \mathsf{Z}) = 0$, i.e. $P^X(A) = 0$.

Let $r_X := dP^X/dQ^X$ be the marginal Radon–Nikodym derivative.

Step 2 (conditional absolute continuity for P^X-a.e. x). We show that for P^X-a.e. x, $P(\cdot \mid x) \ll Q(\cdot \mid x)$. Let

$$\Gamma := \Big\{x \in \mathsf{X} :\; P(\cdot \mid x) \not\ll Q(\cdot \mid x)\Big\}.$$

For each $x \in \Gamma$, there exists a measurable set $B_x \in \mathcal{G}$ such that $Q(B_x \mid x) = 0$ but $P(B_x \mid x) > 0$. To avoid measurability issues from the choice $x \mapsto B_x$, use the following countable reduction: since Z is standard Borel, $\mathcal{G}$ is countably generated, so there exists a countable π-system $\mathcal{C} \subseteq \mathcal{G}$ generating $\mathcal{G}$. If $P(\cdot \mid x) \not\ll Q(\cdot \mid x)$, then there exists some $B \in \mathcal{C}$ with $Q(B \mid x) = 0$ and $P(B \mid x) > 0$ (because otherwise P would be absolutely continuous on the generating class and hence on all of $\mathcal{G}$ by a monotone class argument). Therefore

$$\Gamma = \bigcup_{B \in \mathcal{C}} \Gamma_B, \qquad \Gamma_B := \{x :\; Q(B \mid x) = 0,\; P(B \mid x) > 0\}.$$

Each $\Gamma_B \in \mathcal{F}$ because $x \mapsto P(B \mid x)$ and $x \mapsto Q(B \mid x)$ are measurable (kernel property). Now compute

$$Q(\Gamma_B \times B) = \int_{\Gamma_B} Q(B \mid x)\,Q^X(dx) = \int_{\Gamma_B} 0\,Q^X(dx) = 0.$$

Since $P \ll Q$, we also have $P(\Gamma_B \times B) = 0$, i.e.

$$0 = P(\Gamma_B \times B) = \int_{\Gamma_B} P(B \mid x)\,P^X(dx).$$

But $P(B \mid x) > 0$ on Γ_B by definition, hence $P^X(\Gamma_B) = 0$. Since $\mathcal{C}$ is countable,

$$P^X(\Gamma) \le \sum_{B \in \mathcal{C}} P^X(\Gamma_B) = 0.$$

Thus $P(\cdot \mid x) \ll Q(\cdot \mid x)$ for P^X-a.e. x.

Step 3 (factorization of the joint Radon–Nikodym derivative). For P^X-a.e. x, define $r_{Z|X}(x,\cdot) := \frac{dP(\cdot|x)}{dQ(\cdot|x)}$ (a version), which is $\mathcal{G}$-measurable in z and satisfies $\int r_{Z|X}(x,z)\, Q(dz \mid x) = 1$.

Define

$$r(x,z) := r_X(x)\, r_{Z|X}(x,z),$$

interpreting $r_{Z|X}(x,z) = 0$ when $r_X(x) = 0$ (which can only occur on a P^X-null set anyway). We claim $r = dP/dQ$ Q-a.e.

To verify, it suffices (by uniqueness of Radon–Nikodym derivatives) to check that for all rectangles $A \times B$,

$$P(A \times B) = \int_{A \times B} r(x,z)\, Q(dx, dz). \tag{212}$$

Compute the right-hand side using disintegration of Q and then the definition of conditional Radon–Nikodym derivatives:

$$\int_{A\times B} r(x,z)\, Q(dx,dz) = \int_A \Big(\int_B r_X(x) r_{Z|X}(x,z)\, Q(dz \mid x)\Big) Q^X(dx) = \int_A r_X(x) \Big(\int_B r_{Z|X}(x,z)\, Q(dz \mid x)\Big) Q^X$$

$$= \int_A r_X(x)\, P(B \mid x)\, Q^X(dx) = \int_A P(B \mid x)\, P^X(dx) = P(A \times B),$$

where we used $P^X(dx) = r_X(x) Q^X(dx)$. This proves (212), hence $r = dP/dQ$ Q-a.e.

Step 4 (chain rule identity). Using $dP = r\, dQ$ and $\log r = \log r_X + \log r_{Z|X}$,

$$\mathrm{KL}(P\|Q) = \int \log r\, dP = \int \log r\ r\, dQ.$$

A more transparent decomposition is obtained by integrating with respect to P directly:

$$\mathrm{KL}(P\|Q) = \int \log\Big(\frac{dP}{dQ}\Big)\, dP = \int \big(\log r_X(x) + \log r_{Z|X}(x,z)\big)\, P(dx,dz)$$

$$= \int \log r_X(x)\, P^X(dx) + \int \Big(\int \log r_{Z|X}(x,z)\, P(dz \mid x)\Big) P^X(dx).$$

The first term equals $\mathrm{KL}(P^X\|Q^X)$ because $r_X = dP^X/dQ^X$ and $P^X = r_X Q^X$:

$$\int \log r_X\, dP^X = \int \log\Big(\frac{dP^X}{dQ^X}\Big)\, dP^X = \mathrm{KL}(P^X\|Q^X).$$

For the second term, the inner integral is exactly $\mathrm{KL}(P(\cdot \mid x)\|Q(\cdot \mid x))$ for P^X-a.e. x, by definition of $r_{Z|X}(x,\cdot)$. Thus

$$\mathrm{KL}(P\|Q) = \mathrm{KL}(P^X\|Q^X) + \mathbb{E}_{P^X}\big[\mathrm{KL}(P(\cdot \mid X)\|Q(\cdot \mid X))\big],$$

with the understanding that if any term is $+\infty$ then the equality holds in $[0,\infty]$ (all quantities are nonnegative by Lemma A.19).

Step 5 (measurability of conditional KL). We prove that $x \mapsto \mathrm{KL}(P(\cdot \mid x)\|Q(\cdot \mid x))$ is $\mathcal{F}$-measurable.

Since Z is standard Borel, $\mathcal{G}$ is countably generated. Let $\mathcal{A}_0 \subseteq \mathcal{G}$ be a countable algebra that generates $\mathcal{G}$, and let Π be the countable set of all finite measurable partitions $\pi = \{A_1, \dots, A_m\}$ with each $A_i \in \mathcal{A}_0$.

For probability measures R, S on $(\mathsf{Z}, \mathcal{G})$, define the partition divergence

$$\mathrm{KL}_\pi(R\|S) := \sum_{i=1}^{m} R(A_i) \log \frac{R(A_i)}{S(A_i)},$$

with the conventions $0 \log(0/s) := 0$ and $r \log(r/0) := +\infty$ for $r > 0$. A standard characterization of KL (proved below) is

$$\mathrm{KL}(R\|S) = \sup_{\pi} \mathrm{KL}_\pi(R\|S), \tag{213}$$

where the supremum is over all finite measurable partitions. Moreover, because $\mathcal{A}_0$ generates $\mathcal{G}$, it suffices to take the supremum over Π:

$$\mathrm{KL}(R\|S) = \sup_{\pi \in \Pi} \mathrm{KL}_\pi(R\|S). \tag{214}$$

Measurability from (214). For each fixed $\pi = \{A_1, \dots, A_m\} \in \Pi$, the map

$$x \mapsto \mathrm{KL}_\pi(P(\cdot \mid x)\|Q(\cdot \mid x)) = \sum_{i=1}^{m} P(A_i \mid x) \log \frac{P(A_i \mid x)}{Q(A_i \mid x)}$$

is $\mathcal{F}$-measurable because $x \mapsto P(A_i \mid x)$ and $x \mapsto Q(A_i \mid x)$ are measurable (kernel property) and the map $(u, v) \mapsto u \log(u/v)$ is Borel on $[0,1]^2$ with the stated conventions. Since Π is countable, the pointwise supremum

$$x \mapsto \sup_{\pi \in \Pi} \mathrm{KL}_\pi(P(\cdot \mid x)\|Q(\cdot \mid x))$$

is measurable. By (214), this supremum equals $\mathrm{KL}(P(\cdot \mid x)\|Q(\cdot \mid x))$, proving measurability.

Proof of (213) *and* (214). Let $R \ll S$ with density $r = dR/dS$ (else KL is $+\infty$ and the partition supremum also yields $+\infty$ by choosing a cell where S is zero and R positive). For any finite partition $\pi = \{A_i\}$, define the conditional expectation of r on each cell:

$$\bar{r}_i := \frac{1}{S(A_i)} \int_{A_i} r \, dS \quad \text{if } S(A_i) > 0, \qquad \bar{r}_i := 0 \text{ if } S(A_i) = 0.$$

Then $R(A_i) = \int_{A_i} r \, dS = \bar{r}_i S(A_i)$. Hence

$$\mathrm{KL}_\pi(R\|S) = \sum_i \bar{r}_i S(A_i) \log \bar{r}_i.$$

On the other hand,

$$\mathrm{KL}(R\|S) = \int r \log r \, dS = \sum_i \int_{A_i} r \log r \, dS.$$

By Jensen's inequality (Lemma A.17) applied to the convex function $u \mapsto u \log u$ under the probability measure $S(\cdot \mid A_i) = S(\cdot \cap A_i)/S(A_i)$ (for cells with $S(A_i) > 0$),

$$\frac{1}{S(A_i)} \int_{A_i} r \log r \, dS \;\geq\; \Big(\frac{1}{S(A_i)} \int_{A_i} r \, dS\Big) \log \Big(\frac{1}{S(A_i)} \int_{A_i} r \, dS\Big) = \bar{r}_i \log \bar{r}_i,$$

thus $\int_{A_i} r \log r \, dS \geq \bar{r}_i S(A_i) \log \bar{r}_i$. Summing over i yields

$$\mathrm{KL}(R\|S) \;\geq\; \mathrm{KL}_\pi(R\|S) \qquad \text{for every finite partition } \pi,$$

so $\mathrm{KL}(R\|S) \geq \sup_\pi \mathrm{KL}_\pi(R\|S)$.

For the reverse inequality, approximate r by simple functions measurable with respect to a refining sequence of finite partitions: since $\mathcal{G}$ is countably generated, there exists an increasing sequence of finite sub-σ-fields $\mathcal{G}_n$ with $\sigma(\cup_n \mathcal{G}_n) = \mathcal{G}$. Let $r_n := \mathbb{E}_S[r \mid \mathcal{G}_n]$ be the conditional expectation under S. Then $r_n \to r$ in $L^1(S)$ and $r_n \to r$ S-a.e. (martingale convergence theorem). Moreover, by convexity of $u \mapsto u \log u$ and Jensen applied conditionally,

$$\int r_n \log r_n \, dS \;\uparrow\; \int r \log r \, dS,$$

because (r_n) is the S-martingale of conditional expectations and $u \mapsto u \log u$ is convex with appropriate integrability handled by truncation (standard; one may first prove for bounded r, then pass by monotone truncation $r \wedge M$ and let $M \uparrow \infty$).

For each n, r_n is constant on atoms of $\mathcal{G}_n$, i.e. $r_n = \sum_i \bar{r}_{n,i} \mathbf{1}_{A_{n,i}}$ for a finite partition $\pi_n = \{A_{n,i}\}$ with $A_{n,i} \in \mathcal{G}_n$. Then

$$\int r_n \log r_n \, dS = \sum_i \bar{r}_{n,i} S(A_{n,i}) \log \bar{r}_{n,i} = \mathrm{KL}_{\pi_n}(R \| S).$$

Taking $n \to \infty$ gives

$$\mathrm{KL}(R\|S) = \lim_{n\to\infty} \mathrm{KL}_{\pi_n}(R\|S) \;\le\; \sup_{\pi} \mathrm{KL}_{\pi}(R\|S),$$

hence equality (213).

Finally, because the generating algebra $\mathcal{A}_0$ generates $\mathcal{G}$, every finite partition can be approximated (refined) by partitions from Π without decreasing the supremum; more precisely, for any finite partition π and any $\varepsilon > 0$ one can find $\tilde{\pi} \in \Pi$ such that $\mathrm{KL}_{\tilde{\pi}}(R\|S) \ge \mathrm{KL}_{\pi}(R\|S) - \varepsilon$ (using approximation of sets in $\mathcal{G}$ by sets in the generating algebra and continuity of the partition functional). Therefore the supremum over all finite partitions equals the supremum over Π, proving (214). □

How these inequalities are used in EM arguments. Lemma A.17 yields the standard EM lower bound (Jensen in the conditional expectation defining Q); Lemmas A.19 and A.21 convert log-likelihood gaps into KL divergences and split them into marginal and conditional terms; Lemma A.18 provides first-order (supporting-hyperplane) controls for log terms in stability and perturbation bounds.

Appendix B. Quotient constructions and orbit-level measurability

This appendix collects measure-theoretic and geometric facts about orbit spaces and local slices. The goal is to make precise (i) how to put a canonical σ-field on the orbit space Θ/G, (ii) when Θ/G is standard Borel (globally or locally), and (iii) how orbit-level distances relate to distances between chosen representatives.

Standing notation. Let G be a group acting on Θ by a left action

$$a : G \times \Theta \to \Theta, \qquad a(g, \theta) = g \cdot \theta.$$

Write $[(]\, \theta) = \{g \cdot \theta : g \in G\}$, let Θ/G denote the orbit set, and let

$$q : \Theta \to \Theta/G, \qquad q(\theta) = [(]\, \theta)$$

be the quotient map.

Unless stated otherwise, $(\Theta, \mathcal{B}(\Theta))$ is a measurable space, typically *standard Borel*. When Θ is Polish, $\mathcal{B}(\Theta)$ denotes its Borel σ-field.

B.1. **Standard Borel preliminaries used repeatedly.** We record a few elementary facts used throughout.

Lemma B.1 (Diagonals are Borel)**.** *If Y is a standard Borel space, then the diagonal*

$$\mathrm{Diag}_Y := \{(y, y) : y \in Y\} \subseteq Y \times Y$$

is Borel in $Y \times Y$.

Proof. Choose a Polish topology on a set Y' and a Borel isomorphism $\varphi : Y \to Y'$ (existence is part of the definition of standard Borel). Then $\mathrm{Diag}_{Y'}$ is closed in the Polish space $Y' \times Y'$ and hence Borel. Since $\varphi \times \varphi$ is a Borel isomorphism, $\mathrm{Diag}_Y = (\varphi \times \varphi)^{-1}(\mathrm{Diag}_{Y'})$ is Borel. □

Lemma B.2 (Graphs of Borel maps are Borel)**.** *Let X, Y be standard Borel spaces and let $f : X \to Y$ be Borel. Then*

$$\mathrm{Graph}(f) := \{(x, y) \in X \times Y : \ y = f(x)\}$$

is Borel in $X \times Y$.

Proof. The map $F : X \times Y \to Y \times Y$, $F(x, y) = (y, f(x))$, is Borel because f is Borel and products preserve measurability. Then $\mathrm{Graph}(f) = F^{-1}(\mathrm{Diag}_Y)$, which is Borel by Lemma B.1. □

Lemma B.3 (Analyticity via projection)**.** *Let X, Y be standard Borel and let $A \subseteq X \times Y$ be Borel. Then $\pi_X(A) \subseteq X$ is analytic, where $\pi_X(x, y) = x$ is the projection.*

Proof. By definition (one of the equivalent ones), analytic sets in a standard Borel space are precisely projections of Borel sets in a product with a standard Borel space. □

Lemma B.4 (Fixed-point sets of measurable idempotents)**.** *Let $(X, \mathcal{A})$ be measurable and let $r : X \to X$ be $\mathcal{A}$-measurable. Then the fixed-point set*

$$\mathrm{Fix}(r) := \{x \in X : r(x) = x\}$$

is in $\mathcal{A}$ whenever the diagonal $\mathrm{Diag}_X \subseteq X \times X$ is measurable in $\mathcal{A} \otimes \mathcal{A}$ (e.g. when X is standard Borel).

Proof. Consider the measurable map $x \mapsto (x, r(x))$ from $(X, \mathcal{A})$ to $(X \times X, \mathcal{A} \otimes \mathcal{A})$. Then $\mathrm{Fix}(r) = (x \mapsto (x, r(x)))^{-1}(\mathrm{Diag}_X)$. □

B.2. **Measurability of the quotient map and orbit relation.**

Quotient σ-field. We equip Θ/G with the *quotient σ-field*

$$\mathcal{B}(\Theta/G) \ := \ \{A \subseteq \Theta/G : \ q^{-1}(A) \in \mathcal{B}(\Theta)\}. \tag{215}$$

Lemma B.5 (Basic properties of the quotient σ-field)**.** *The collection $\mathcal{B}(\Theta/G)$ in (215) is a σ-field on Θ/G. Moreover:*

(1) $q : (\Theta, \mathcal{B}(\Theta)) \to (\Theta/G, \mathcal{B}(\Theta/G))$ is measurable.
(2) $\mathcal{B}(\Theta/G)$ is the largest *σ-field on Θ/G for which q is measurable: if $\mathcal{A}$ is a σ-field on Θ/G such that q is $(\mathcal{B}(\Theta), \mathcal{A})$-measurable, then $\mathcal{A} \subseteq \mathcal{B}(\Theta/G)$.*

Proof. **Step 1 (Show $\mathcal{B}(\Theta/G)$ is a σ-field).**

- $\emptyset, \Theta/G \in \mathcal{B}(\Theta/G)$ since $q^{-1}(\emptyset) = \emptyset$ and $q^{-1}(\Theta/G) = \Theta$ lie in $\mathcal{B}(\Theta)$.
- If $A \in \mathcal{B}(\Theta/G)$ then $q^{-1}(A) \in \mathcal{B}(\Theta)$ and

  $$q^{-1}(A^c) = \Theta \setminus q^{-1}(A) \in \mathcal{B}(\Theta),$$

  so $A^c \in \mathcal{B}(\Theta/G)$.

- If $A_n \in \mathcal{B}(\Theta/G)$, then $q^{-1}(A_n) \in \mathcal{B}(\Theta)$ and
$$q^{-1}\Big(\bigcup_{n\geq 1} A_n\Big) = \bigcup_{n\geq 1} q^{-1}(A_n) \in \mathcal{B}(\Theta),$$
so $\bigcup_{n\geq 1} A_n \in \mathcal{B}(\Theta/G)$.

Step 2 (Measurability of q). For each $A \in \mathcal{B}(\Theta/G)$, the preimage $q^{-1}(A)$ is in $\mathcal{B}(\Theta)$ by definition. This is exactly measurability of q.

Step 3 (Maximality). Let $\mathcal{A}$ be a σ-field on Θ/G such that q is measurable into $(\Theta/G, \mathcal{A})$. Then for any $A \in \mathcal{A}$, $q^{-1}(A) \in \mathcal{B}(\Theta)$, hence $A \in \mathcal{B}(\Theta/G)$ by (215). Thus $\mathcal{A} \subseteq \mathcal{B}(\Theta/G)$. □

Saturated sets and pullbacks. A subset $B \subseteq \Theta$ is *G-invariant* (or *saturated*) if $g \cdot B = B$ for all $g \in G$. Equivalently, $B = q^{-1}(q(B))$.

Lemma B.6 (Saturated sets are exactly quotient pullbacks)**.** *A set $B \subseteq \Theta$ is saturated if and only if $B = q^{-1}(A)$ for some $A \subseteq \Theta/G$ (namely $A = q(B)$). Moreover, if $B \in \mathcal{B}(\Theta)$ is saturated, then $q(B) \in \mathcal{B}(\Theta/G)$.*

Proof. **Step 1 ($q^{-1}(A)$ is saturated).** If $B = q^{-1}(A)$ then $q(B) = q(q^{-1}(A)) = A$ (surjectivity of q), hence $q^{-1}(q(B)) = q^{-1}(A) = B$.

Step 2 (Saturated implies pullback). If B is saturated, define $A := q(B)$. Then $B = q^{-1}(q(B)) = q^{-1}(A)$.

Step 3 (Borelness on the quotient). If $B \in \mathcal{B}(\Theta)$ and $B = q^{-1}(q(B))$, then $q^{-1}(q(B)) \in \mathcal{B}(\Theta)$, so $q(B) \in \mathcal{B}(\Theta/G)$ by definition. □

The orbit equivalence relation. Define
$$E \subseteq \Theta \times \Theta, \qquad (\theta, \theta') \in E \iff \exists g \in G :\ \theta' = g \cdot \theta.$$

Assumption B.7 (Borel action)**.** Θ is standard Borel. The action map $a(g, \theta) = g \cdot \theta$ is measurable from $(G \times \Theta, \mathcal{B}(G) \otimes \mathcal{B}(\Theta))$ to $(\Theta, \mathcal{B}(\Theta))$ for some σ-field $\mathcal{B}(G)$ on G. Typical cases: (i) G countable, $\mathcal{B}(G) = 2^G$; (ii) G second-countable topological group, $\mathcal{B}(G)$ its Borel σ-field, acting Borel (in particular continuously) on a Polish Θ.

Lemma B.8 (Orbit relation is analytic; Borel for countable groups)**.** *Under Assumption B.7, E is analytic in $\Theta \times \Theta$. If G is countable and $\mathcal{B}(G) = 2^G$, then E is Borel.*

Proof. **Step 1 (Encode E as a projection of a Borel set).** Consider the set
$$A := \{(g, \theta, \theta') \in G \times \Theta \times \Theta :\ \theta' = g \cdot \theta\}.$$
Define the measurable map
$$F : G \times \Theta \times \Theta \to \Theta \times \Theta, \qquad F(g, \theta, \theta') = (\theta',\ g \cdot \theta).$$
Measurability of F follows from measurability of the coordinate projection $(g, \theta, \theta') \mapsto \theta'$ and the action map. Since Θ is standard Borel, Lemma B.1 gives that the diagonal Diag_Θ is Borel. Hence
$$A = F^{-1}(\mathrm{Diag}_\Theta)$$
is Borel in $G \times \Theta \times \Theta$.

Step 2 (Project to obtain E). Let $\pi : G \times \Theta \times \Theta \to \Theta \times \Theta$ be $\pi(g, \theta, \theta') = (\theta, \theta')$. Then
$$\pi(A) = \{(\theta, \theta') : \exists g \in G,\ \theta' = g \cdot \theta\} = E.$$
By Lemma B.3, $\pi(A)$ is analytic, hence E is analytic.

Step 3 (Countable G gives Borel). Assume G is countable and $\mathcal{B}(G) = 2^G$. For each fixed $g \in G$, the map $\theta \mapsto g \cdot \theta$ is measurable (as a section of the jointly measurable action), hence Borel. By Lemma B.2,

$$E_g := \{(\theta, \theta') : \theta' = g \cdot \theta\}$$

is Borel in $\Theta \times \Theta$. Finally,

$$E = \bigcup_{g \in G} E_g,$$

a countable union of Borel sets, so E is Borel. □

Standard-Borel quotients under a measurable section.

Proposition B.9 (Standard-Borel quotient under a measurable section)**.** *Assume Θ is standard Borel and Θ/G is equipped with $\mathcal{B}(\Theta/G)$ from (215). Suppose there exists a measurable section $s : \Theta/G \to \Theta$ such that $q \circ s = \mathrm{id}_{\Theta/G}$. Let $S := s(\Theta/G) \subseteq \Theta$. Then:*

(1) s is injective, and $q|_S : S \to \Theta/G$ is a bijection with inverse s.
(2) $S \in \mathcal{B}(\Theta)$.
(3) The restriction $q|_S : (S, \mathcal{B}(S)) \to (\Theta/G, \mathcal{B}(\Theta/G))$ is a Borel isomorphism, where $\mathcal{B}(S) = \{S \cap B : B \in \mathcal{B}(\Theta)\}$. In particular, $(\Theta/G, \mathcal{B}(\Theta/G))$ is standard Borel.

Proof. **Step 1 (Injectivity and bijection).** If $s(\bar\theta) = s(\bar\theta')$, apply q to obtain $\bar\theta = q(s(\bar\theta)) = q(s(\bar\theta')) = \bar\theta'$, so s is injective. Since $q(s(\bar\theta)) = \bar\theta$ for all $\bar\theta$, the map $q|_S$ is surjective and has inverse s. Injectivity of $q|_S$ follows from injectivity of s, so $q|_S$ is bijective.

Step 2 (Borelness of S). Define $r : \Theta \to \Theta$ by

$$r := s \circ q.$$

This is measurable because q is measurable (Lemma B.5) and s is measurable by assumption. Moreover, r is idempotent:

$$r(r(\theta)) = s(q(s(q(\theta)))) = s(q(\theta)) = r(\theta).$$

We claim that $S = \mathrm{Fix}(r)$. Indeed, if $\theta \in S$ then $\theta = s(\bar\theta)$ for some $\bar\theta$ and hence $r(\theta) = s(q(s(\bar\theta))) = s(\bar\theta) = \theta$. Conversely, if $r(\theta) = \theta$, then $\theta = s(q(\theta)) \in S$.

Since Θ is standard Borel, Diag_Θ is Borel and Lemma B.4 yields $\mathrm{Fix}(r) \in \mathcal{B}(\Theta)$. Thus $S \in \mathcal{B}(\Theta)$.

Step 3 (Borel isomorphism). The map $q|_S$ is measurable as the restriction of a measurable map to a measurable subset. Its inverse is s, which is measurable and takes values in S. Hence $q|_S$ is a measurable bijection with measurable inverse, i.e. a Borel isomorphism.

Step 4 (Standard Borelness). Because S is Borel in the standard Borel space Θ, the measurable space $(S, \mathcal{B}(S))$ is standard Borel. A Borel isomorphism preserves standard-Borel structure, hence $(\Theta/G, \mathcal{B}(\Theta/G))$ is standard Borel. □

Remark B.10 (Smoothness and selectors). In descriptive set theory, existence of a Borel section is equivalent to smoothness of the orbit equivalence relation (Borel reducible to equality on a standard Borel space). Proposition B.9 is the corresponding standard-Borel consequence. In the main text we only need local sections (Definition B.11).

B.3. Local sections (slices): existence and construction templates.

Definition B.11 (Local measurable section)**.** Let $\bar\theta_0 \in \Theta/G$. A *local measurable section at $\bar\theta_0$* consists of $U \in \mathcal{B}(\Theta/G)$ with $\bar\theta_0 \in U$ and a measurable map $s : U \to \Theta$ such that

$$q(s(\bar\theta)) = \bar\theta, \qquad \forall \bar\theta \in U.$$

Lemma B.12 (Equivalent slice formulation)**.** *Let $U \in \mathcal{B}(\Theta/G)$ and $s : U \to \Theta$. Set $S := s(U)$. Then the following are equivalent:*

(1) s is a local measurable section on U.
(2) $S \in \mathcal{B}(\Theta)$, $S \subseteq q^{-1}(U)$, and for each $\bar{\theta} \in U$ the fiber $q^{-1}(\bar{\theta})$ meets S in exactly one point. Equivalently, $q|_S : S \to U$ is a measurable bijection with measurable inverse s.

Proof. **(1)⇒(2).** Assume s is measurable and $q \circ s = \mathrm{id}_U$.

Step 1 (Fiberwise uniqueness). Fix $\bar{\theta} \in U$. If $\theta \in S \cap q^{-1}(\bar{\theta})$ then $\theta = s(\bar{\theta}')$ for some $\bar{\theta}' \in U$ and

$$\bar{\theta} = q(\theta) = q(s(\bar{\theta}')) = \bar{\theta}'.$$

Hence $\theta = s(\bar{\theta})$. So $S \cap q^{-1}(\bar{\theta}) = \{s(\bar{\theta})\}$.

Step 2 (S is Borel). Define $r : q^{-1}(U) \to \Theta$ by $r := s \circ q$ (well-defined on $q^{-1}(U)$). Then r is measurable and idempotent on $q^{-1}(U)$, and

$$S = \{\theta \in q^{-1}(U) : r(\theta) = \theta\}.$$

Because $q^{-1}(U) \in \mathcal{B}(\Theta)$ and Θ is standard Borel locally on $q^{-1}(U)$, Lemma B.4 implies the fixed-point set is measurable; hence $S \in \mathcal{B}(\Theta)$.

Step 3 (Measurable bijection). By Step 1, $q|_S$ is bijective onto U with inverse s. Measurability of $q|_S$ follows from measurability of q; measurability of s is assumed.

(2)⇒(1). Assume (2). Since $q|_S : S \to U$ is a measurable bijection with measurable inverse, define s to be this inverse followed by the inclusion $S \hookrightarrow \Theta$. Then $q \circ s = \mathrm{id}_U$ and s is measurable. □

Finite group actions: a global Borel section by tie-breaking.

Assumption B.13 (Ordered parameter space)**.** $\Theta \subseteq \mathbb{R}^d$ is Borel. Fix the lexicographic order $\preceq$ on $\mathbb{R}^d$.

Lemma B.14 (Lexicographic order is Borel)**.** *The set $\{(x, y) \in \mathbb{R}^d \times \mathbb{R}^d : x \preceq y\}$ is Borel in $\mathbb{R}^{2d}$. Consequently, if $f, g : \Theta \to \mathbb{R}^d$ are Borel, then $\{\theta : f(\theta) \preceq g(\theta)\}$ is Borel in Θ.*

Proof. Write

$$\{(x, y) : x \preceq y\} = \{(x, y) : x = y\} \ \cup \ \bigcup_{j=1}^{d} \{(x, y) : x_1 = y_1, \ldots, x_{j-1} = y_{j-1},\ x_j < y_j\}.$$

Each equality constraint is closed; each strict inequality is open; finite unions preserve Borelness. The second claim follows by pulling back this Borel set under the Borel map $\theta \mapsto (f(\theta), g(\theta))$. □

Lemma B.15 (Borel minimum over finitely many Borel maps)**.** *Let Y be standard Borel with a Borel total order $\preceq$ and let $f_1, \ldots, f_m : \Theta \to Y$ be Borel. Then the pointwise minimum $M(\theta) = \min_{\preceq}\{f_1(\theta), \ldots, f_m(\theta)\}$ is Borel. Moreover, the tie-broken argmin*

$$I(\theta) := \min\{i : \ f_i(\theta) = M(\theta)\}$$

is Borel as a map into the discrete space $\{1, \ldots, m\}$.

Proof. **Step 1 (Borel selector sets).** For each i define

$$A_i := \Big(\bigcap_{j<i} \{f_j \prec f_i\}^c \Big) \cap \Big(\bigcap_{j=1}^{m} \{f_i \preceq f_j\} \Big),$$

where $f_j \prec f_i$ means $f_j \preceq f_i$ and not $f_i \preceq f_j$. Each comparison set $\{f_i \preceq f_j\}$ is Borel because (f_i, f_j) is Borel and $\preceq$ is Borel.

Step 2 (Partition). For each θ, at least one $f_i(\theta)$ is minimal; among minimizers there is a unique smallest index. The definition of A_i enforces precisely that rule, so (A_i) are disjoint and cover Θ.

Step 3 (Piecewise definition). On A_i we have $M = f_i$ and $I = i$. Therefore M and I are Borel as finite piecewise combinations of Borel maps. □

Proposition B.16 (Canonical representative under a finite Borel action)**.** *Assume $\Theta \subseteq \mathbb{R}^d$ is Borel, G is finite, and $\theta \mapsto g \cdot \theta$ is Borel for each $g \in G$. Define*

$$\pi(\theta) := \min\nolimits_{\preceq}\{g \cdot \theta : g \in G\}.$$

Then π is Borel, $\pi(g \cdot \theta) = \pi(\theta)$, and $s(q(\theta)) := \pi(\theta)$ is a well-defined Borel section.

Proof. Enumerate $G = \{g_1, \dots, g_m\}$ and set $f_i(\theta) = g_i \cdot \theta$.

Step 1 (Borelness of π). By Lemma B.15, $\pi(\theta) = \min_{\preceq}\{f_1(\theta), \dots, f_m(\theta)\}$ is Borel.

Step 2 (G-invariance). Fix $h \in G$. Left multiplication permutes G, so $\{g \cdot (h \cdot \theta) : g \in G\} = \{u \cdot \theta : u \in G\}$ and the minimum is unchanged. Hence $\pi(h \cdot \theta) = \pi(\theta)$.

Step 3 (Well-defined section). Define $s(\bar{\theta}) = \pi(\theta)$ for any θ with $q(\theta) = \bar{\theta}$. If $q(\theta) = q(\theta')$, then $\theta' = h \cdot \theta$ for some h and Step 2 gives $\pi(\theta') = \pi(\theta)$, so s is well-defined. Also $q(s(q(\theta))) = q(\pi(\theta)) = q(\theta)$ since $\pi(\theta) \in [(]\,\theta)$.

Step 4 (Measurability of s). For $B \in \mathcal{B}(\Theta)$,

$$q^{-1}(s^{-1}(B)) = \{\theta : \pi(\theta) \in B\} = \pi^{-1}(B) \in \mathcal{B}(\Theta),$$

so $s^{-1}(B) \in \mathcal{B}(\Theta/G)$ by definition of the quotient σ-field. □

Compact Lie group actions: a detailed Riemannian slice theorem.

Theorem B.17 (Slice theorem for compact Lie group actions)**.** *Let G be a compact Lie group acting smoothly on a smooth manifold Θ (finite-dimensional, second countable). Fix $\theta_0 \in \Theta$ and let $H := G_{\theta_0}$ be the stabilizer.*

Then there exist:

(1) a G-invariant Riemannian metric on Θ;
(2) an embedded submanifold $S \subseteq \Theta$ containing θ_0 such that $h \cdot S = S$ for all $h \in H$;
(3) an open neighborhood $\mathcal{U}$ of θ_0

so that with $S_0 := S \cap \mathcal{U}$:

(i) $\mathcal{U} = G \cdot S_0$;
(ii) if $g \cdot s \in S_0$ for some $g \in G$ and $s \in S_0$, then $g \in H$;
(iii) for all $s \in S_0$, $[(]\,s) \cap S_0 = H \cdot s$.

If the action is free on $\mathcal{U}$, then $H = \{e\}$ and each orbit in $\mathcal{U}$ meets S_0 in exactly one point; hence $q|_{S_0} : S_0 \to q(\mathcal{U})$ is a homeomorphism and $(q|_{S_0})^{-1}$ is a continuous (hence Borel) local section.

Proof. **Step 1 (Construct a G-invariant metric).** Start with any Riemannian metric $\langle \cdot, \cdot \rangle_0$. Let λ be normalized Haar measure on compact G and define

$$\langle v, w \rangle_\theta := \int_G \big\langle (dg)_\theta v, (dg)_\theta w \big\rangle_{0, g\cdot\theta}\, d\lambda(g).$$

Smoothness of the action implies the integrand is smooth in (g, θ), allowing differentiation under the integral. Positive definiteness holds because this is an average of positive definite bilinear forms. G-invariance follows by a change of variables using left invariance of Haar measure.

Step 2 (The orbit is an embedded submanifold). Let $\alpha : G \to \Theta$ be $\alpha(g) = g \cdot \theta_0$. For each g, $d\alpha_g$ has the same rank as $d\alpha_e$ because $\alpha \circ L_g = \alpha_g$ and left translation is a diffeomorphism. Hence α has constant rank, so $\alpha(G) = G \cdot \theta_0$ is an immersed submanifold (constant rank theorem).

Because G is compact, α is proper; a proper immersion is an embedding, so $\mathcal{O} := G \cdot \theta_0$ is embedded.

Step 3 (Equivariant tubular neighborhood of $\mathcal{O}$). Let $N\mathcal{O} \to \mathcal{O}$ be the normal bundle with respect to the G-invariant metric. The tubular neighborhood theorem gives an open neighborhood V of the zero section in $N\mathcal{O}$ such that the normal exponential map

$$\mathrm{Exp} : V \to \Theta, \qquad (\theta, v) \mapsto \exp_\theta(v),$$

is a diffeomorphism onto an open neighborhood $\mathcal{U}$ of $\mathcal{O}$. Because the metric is G-invariant, the exponential map is G-equivariant:

$$\mathrm{Exp}(g \cdot \theta, (dg)_\theta v) = g \cdot \mathrm{Exp}(\theta, v).$$

(Equivariance is verified by noting g is an isometry, hence maps geodesics to geodesics and preserves initial conditions.)

Step 4 (Build the slice S at θ_0). Let $N_{\theta_0}\mathcal{O}$ be the normal space at θ_0. Choose an H-invariant open ball $B \subset N_{\theta_0}\mathcal{O}$ around 0 small enough that $(\theta_0, v) \in V$ for all $v \in B$ and $\exp_{\theta_0}$ restricts to an embedding on B. Define

$$S := \exp_{\theta_0}(B), \qquad S_0 := S \cap \mathcal{U}.$$

Then S is an embedded submanifold through θ_0.

Step 5 (H-invariance of S). If $h \in H$ then $h \cdot \theta_0 = \theta_0$ and $(dh)_{\theta_0}$ preserves $N_{\theta_0}\mathcal{O}$ (because it preserves $T_{\theta_0}\mathcal{O}$). By equivariance of Exp,

$$h \cdot \exp_{\theta_0}(v) = \exp_{\theta_0}((dh)_{\theta_0} v) \in S.$$

Since B is H-invariant, $h \cdot S = S$.

Step 6 (Saturation $\mathcal{U} = G \cdot S_0$). Take $x \in \mathcal{U}$. Write $x = \mathrm{Exp}(\theta, v)$ for unique $(\theta, v) \in V$, with $\theta \in \mathcal{O}$. Choose $g \in G$ such that $\theta = g \cdot \theta_0$. Then by equivariance,

$$x = \mathrm{Exp}(g \cdot \theta_0, v) = g \cdot \mathrm{Exp}(\theta_0, (dg^{-1})_\theta v).$$

For V chosen small, $(dg^{-1})_\theta v \in B$, hence $\mathrm{Exp}(\theta_0, (dg^{-1})_\theta v) \in S_0$, proving $x \in G \cdot S_0$.

Step 7 (Property (ii) from uniqueness of tubular coordinates). Assume $g \cdot s \in S_0$ with $s \in S_0$. Write $s = \mathrm{Exp}(\theta_0, v)$ and $g \cdot s = \mathrm{Exp}(\theta_0, w)$ with $v, w \in B$. But also, by equivariance,

$$g \cdot s = g \cdot \mathrm{Exp}(\theta_0, v) = \mathrm{Exp}(g \cdot \theta_0, (dg)_{\theta_0} v).$$

Thus the point $g \cdot s \in \mathcal{U}$ has two tubular representations:

$$(\theta_0, w) \quad \text{and} \quad (g \cdot \theta_0, (dg)_{\theta_0} v).$$

Since $\mathrm{Exp} : V \to \mathcal{U}$ is injective, these representations must coincide. Hence $g \cdot \theta_0 = \theta_0$, so $g \in H$.

Step 8 (Property (iii)). Fix $s \in S_0$. If $g \cdot s \in S_0$, then Step 7 gives $g \in H$, so $[(]\, s) \cap S_0 \subseteq H \cdot s$. The reverse inclusion holds because S_0 is H-invariant. Hence $[(]\, s) \cap S_0 = H \cdot s$.

Step 9 (Free-action conclusion). If the action is free on $\mathcal{U}$, then $H = G_{\theta_0} = \{e\}$. Step 8 implies each orbit meets S_0 in exactly one point, so $q|_{S_0}$ is bijective onto $U := q(\mathcal{U})$. Standard slice theory yields that $q|_{S_0}$ is a homeomorphism onto its image in the quotient topology, hence its inverse is continuous. □

B.4. **Orbit distances and moduli comparisons.** Assume (Θ, d_Θ) is a metric space and the action is by isometries:

$$d_\Theta(g \cdot \theta,\ g \cdot \theta') = d_\Theta(\theta, \theta') \qquad \forall g \in G.$$

Definition B.18 (Orbit pseudo-metric)**.** Define

$$d_G(\theta, \theta') := \inf_{g \in G} d_\Theta(\theta, g \cdot \theta').$$

Lemma B.19 (Basic properties of d_G)**.** *If the action is by isometries, then d_G is a G-invariant pseudo-metric:*

(1) $d_G(\theta,\theta)=0$ and $d_G(\theta,\theta')\geq 0$.
(2) Symmetry: $d_G(\theta,\theta')=d_G(\theta',\theta)$.
(3) Triangle inequality: $d_G(\theta,\theta'')\leq d_G(\theta,\theta')+d_G(\theta',\theta'')$.
(4) G-invariance: $d_G(g_1\cdot\theta,g_2\cdot\theta')=d_G(\theta,\theta')$.

Proof. **(1)** Nonnegativity is immediate and $d_G(\theta,\theta)\leq d_\Theta(\theta,e\cdot\theta)=0$.

(2) Symmetry. For any $g\in G$, isometry gives $d_\Theta(\theta,g\cdot\theta')=d_\Theta(g^{-1}\cdot\theta,\theta')$. Taking infima and renaming $h=g^{-1}$ yields

$$d_G(\theta,\theta')=\inf_{h\in G} d_\Theta(h\cdot\theta,\theta')=\inf_{h\in G} d_\Theta(\theta',h\cdot\theta)=d_G(\theta',\theta).$$

(3) Triangle inequality. Fix $\varepsilon>0$ and choose g,h with

$$d_\Theta(\theta,g\cdot\theta')\leq d_G(\theta,\theta')+\varepsilon,\qquad d_\Theta(\theta',h\cdot\theta'')\leq d_G(\theta',\theta'')+\varepsilon.$$

Then

$$d_\Theta(\theta,gh\cdot\theta'')\leq d_\Theta(\theta,g\cdot\theta')+d_\Theta(g\cdot\theta',gh\cdot\theta'')=d_\Theta(\theta,g\cdot\theta')+d_\Theta(\theta',h\cdot\theta'')$$

(using isometry of g). Taking infimum over group elements on the left and letting $\varepsilon\downarrow 0$ gives the claim.

(4) G-invariance.

$$d_G(g_1\cdot\theta,g_2\cdot\theta')=\inf_{g\in G} d_\Theta(g_1\cdot\theta,gg_2\cdot\theta')=\inf_{u\in G} d_\Theta(g_1\cdot\theta,u\cdot\theta')=\inf_{v\in G} d_\Theta(\theta,v\cdot\theta')=d_G(\theta,\theta'),$$

where the last equality applies g_1^{-1}. □

Quotient pseudo-metric. By Lemma B.19(4), d_G depends only on orbits, so define

$$d_{\Theta/G}(\bar\theta,\bar\theta'):=d_G(\theta,\theta')\quad\text{for any representatives } q(\theta)=\bar\theta,\ q(\theta')=\bar\theta'.$$

A useful sufficient condition: metric sections.

Definition B.20 (Metric section property)**.** Let $U\subseteq\Theta/G$ and let $s:U\to\Theta$ be a local section with image $S=s(U)$. We say s is a *metric section* on U if for all $\bar\theta,\bar\theta'\in U$,

$$d_\Theta\big(s(\bar\theta),s(\bar\theta')\big)\ \leq\ d_\Theta\big(s(\bar\theta),g\cdot s(\bar\theta')\big)\qquad\forall g\in G.$$

Lemma B.21 (Orbit distance equals slice distance under metric section)**.** *Assume the action is by isometries. If $s:U\to\Theta$ is a metric section, then for all $\bar\theta,\bar\theta'\in U$,*

$$d_{\Theta/G}(\bar\theta,\bar\theta')=d_G(s(\bar\theta),s(\bar\theta'))=d_\Theta(s(\bar\theta),s(\bar\theta')).$$

Proof. Let $x=s(\bar\theta)$ and $y=s(\bar\theta')$.

Step 1 (Always have $d_G(x,y)\leq d_\Theta(x,y)$). By definition,

$$d_G(x,y)=\inf_{g\in G} d_\Theta(x,g\cdot y)\leq d_\Theta(x,e\cdot y)=d_\Theta(x,y).$$

Step 2 (Metric section gives reverse inequality). Metric section means $d_\Theta(x,y)\leq d_\Theta(x,g\cdot y)$ for all g, hence

$$d_\Theta(x,y)\leq\inf_{g\in G} d_\Theta(x,g\cdot y)=d_G(x,y).$$

Thus $d_G(x,y)=d_\Theta(x,y)$.

Step 3 (Identify quotient distance). By definition, $d_{\Theta/G}(\bar\theta,\bar\theta')$ is d_G of any representatives, in particular x,y. □

Metric sections from slices: a detailed (local) argument. The next proposition is the (local) mechanism that turns a slice into a metric section. It uses only compactness, the tubular neighborhood provided by the slice theorem, and basic distance comparisons.

Proposition B.22 (Metric section from a slice (local, detailed))**.** *Assume the setting of Theorem B.17 and take d_Θ to be the Riemannian distance induced by a G-invariant metric. Assume additionally that the action is free on a sufficiently small neighborhood so that $q|_{S_0} : S_0 \to U := q(\mathcal{U})$ is bijective.*

Then, after shrinking S_0 around θ_0, the induced local section $s = (q|_{S_0})^{-1} : U \to S_0$ is a metric section. Consequently, on U,

$$d_{\Theta/G}(\bar{\theta}, \bar{\theta}') = d_\Theta\big(s(\bar{\theta}), s(\bar{\theta}')\big).$$

Proof. We prove that for x, y in a sufficiently small neighborhood inside S_0,

$$d_\Theta(x, y) \;\le\; d_\Theta(x, g \cdot y) \qquad \forall g \in G.$$

The proof is by splitting G into elements *away from* the identity (handled by compactness) and *near* the identity (handled by a local normal-coordinate comparison that uses the slice geometry at θ_0).

Step 1 (Work in a strongly convex normal neighborhood). Choose a normal neighborhood W of θ_0 in Θ that is *strongly geodesically convex*: any two points in W are joined by a unique minimizing geodesic that lies in W. Shrink S_0 and $\mathcal{U}$ so that

$$S_0 \subseteq W, \qquad \mathcal{U} = G \cdot S_0 \subseteq W.$$

This ensures that distance comparisons can be done using unique minimizing geodesics inside W.

Step 2 (Uniform separation for group elements away from identity). Fix an open neighborhood $\mathcal{N}$ of e in G and let $K := G \setminus \mathcal{N}$ (compact). Consider the continuous function

$$F : K \times \overline{S_0} \to \mathbb{R}, \qquad F(g, y) := d_\Theta(y, g \cdot y).$$

Because the action is free on $\mathcal{U}$ and $\overline{S_0} \subseteq \mathcal{U}$, we have $g \cdot y \neq y$ for all $(g, y) \in K \times \overline{S_0}$, so $F > 0$. By compactness, F attains a positive minimum:

$$\delta_0 := \min_{(g,y) \in K \times \overline{S_0}} d_\Theta(y, g \cdot y) \;>\; 0.$$

Now shrink S_0 further so that $\operatorname{diam}(S_0) \le \delta_0/4$. Then for any $x, y \in S_0$ and any $g \in K$,

$$d_\Theta(x, g \cdot y) \;\ge\; d_\Theta(y, g \cdot y) - d_\Theta(x, y) \;\ge\; \delta_0 - \operatorname{diam}(S_0) \;\ge\; 3\delta_0/4,$$

whereas $d_\Theta(x, y) \le \operatorname{diam}(S_0) \le \delta_0/4$. Hence

$$d_\Theta(x, y) \;<\; d_\Theta(x, g \cdot y) \qquad \forall x, y \in S_0,\ \forall g \in K. \tag{216}$$

Step 3 (A local coordinate model near θ_0). Let $\exp_{\theta_0} : T_{\theta_0}\Theta \to \Theta$ be the Riemannian exponential map. Shrink W (and hence S_0) so that $\exp_{\theta_0}$ is a diffeomorphism from a Euclidean ball $B_\rho(0) \subset T_{\theta_0}\Theta$ onto W. Write $u = \exp_{\theta_0}^{-1}(z)$ for $z \in W$.

Decompose the tangent space orthogonally at θ_0:

$$T_{\theta_0}\Theta = V \oplus V^\perp, \qquad V := T_{\theta_0}(G \cdot \theta_0).$$

By construction of the slice in Theorem B.17, after shrinking S_0 we may assume

$$S_0 = \exp_{\theta_0}(B_r(0) \cap V^\perp)$$

for some $r > 0$ small. Thus points of S_0 correspond to vectors in $V^\perp$ under $\exp_{\theta_0}^{-1}$.

Step 4 (Near-identity elements create a transverse displacement). For g near e, define the coordinate displacement map

$$\Psi_g : B_r(0) \cap V^\perp \to T_{\theta_0}\Theta, \qquad \Psi_g(w) := \exp_{\theta_0}^{-1}\big(g \cdot \exp_{\theta_0}(w)\big).$$

The map $(g, w) \mapsto \Psi_g(w)$ is smooth (composition of smooth maps) and $\Psi_e(w) = w$.

We claim: after shrinking r and $\mathcal{N}$, there exists a constant $c > 0$ such that for all $g \in \mathcal{N}$ and all $w \in B_r(0) \cap V^\perp$, the V-component of $\Psi_g(w) - w$ satisfies

$$\big\|\mathrm{Proj}_V(\Psi_g(w) - w)\big\| \;\geq\; c\,\mathrm{dist}_G(g, e), \tag{217}$$

where dist_G is any fixed Riemannian distance on the compact manifold G.

Justification of (217). At $w = 0$, $\Psi_g(0) = \exp_{\theta_0}^{-1}(g \cdot \theta_0)$ lies in V to first order, because the orbit through θ_0 has tangent space V. Moreover, freeness implies the orbit map $g \mapsto g \cdot \theta_0$ is an immersion, so $g \cdot \theta_0 \neq \theta_0$ for $g \neq e$ and the displacement along V is nondegenerate at e. Continuity in w then yields the uniform bound for small w after shrinking r and $\mathcal{N}$.

Step 5 (Quantitative "Pythagorean" inequality near θ_0). Fix $x = \exp_{\theta_0}(u)$ and $y = \exp_{\theta_0}(v)$ with $u, v \in B_r(0) \cap V^\perp$. For $g \in \mathcal{N}$, write $\Psi_g(v) = v + \Delta_g(v)$ with $\Delta_g(v) \in T_{\theta_0}\Theta$. Decompose $\Delta_g(v) = \Delta_g^V(v) + \Delta_g^\perp(v)$ along $V \oplus V^\perp$.

Since $u - v \in V^\perp$ and $\Delta_g^V(v) \in V$, orthogonality at θ_0 gives the exact identity in the tangent space:

$$\|u - (v + \Delta_g(v))\|^2 = \|(u - v) - \Delta_g^\perp(v)\|^2 + \|\Delta_g^V(v)\|^2.$$

Thus, even before controlling $\Delta_g^\perp(v)$, we have the lower bound

$$\|u - (v + \Delta_g(v))\| \;\geq\; \sqrt{\|u - v\|^2 + \|\Delta_g^V(v)\|^2} \;-\; \|\Delta_g^\perp(v)\|. \tag{218}$$

Next, use smoothness of Ψ_g and $\Psi_e = \mathrm{id}$ to obtain (after shrinking r and $\mathcal{N}$) a Lipschitz estimate

$$\|\Delta_g^\perp(v)\| \;\leq\; C\,\mathrm{dist}_G(g, e)\,\|v\| \tag{219}$$

for some $C > 0$ and all $v \in B_r(0) \cap V^\perp$ and $g \in \mathcal{N}$. This is a standard consequence of bounded derivatives on a compact set: $\Delta_g^\perp(v)$ vanishes at $g = e$, and the derivative in g is $O(\|v\|)$.

Combine (217), (218), and (219). Since $\|v\| \leq r$, we get, for $g \in \mathcal{N}$,

$$\|u - \Psi_g(v)\| \;\geq\; \sqrt{\|u - v\|^2 + c^2\mathrm{dist}_G(g, e)^2} \;-\; Cr\,\mathrm{dist}_G(g, e).$$

Choose r small so that $Cr \leq c/2$. Then

$$\|u - \Psi_g(v)\| \;\geq\; \sqrt{\|u - v\|^2 + c^2\mathrm{dist}_G(g, e)^2} \;-\; \tfrac{c}{2}\mathrm{dist}_G(g, e) \;\geq\; \|u - v\|,$$

because $\sqrt{a^2 + b^2} \geq a + \frac{b^2}{2(a+b)} \geq a$ and the subtraction is controlled by b. Hence

$$\|u - \Psi_g(v)\| \;\geq\; \|u - v\| \qquad \forall u, v \in B_r(0) \cap V^\perp,\ \forall g \in \mathcal{N}. \tag{220}$$

Step 6 (Transfer the coordinate inequality to the Riemannian distance). Because $\exp_{\theta_0}$ is a diffeomorphism on $B_\rho(0)$ and the metric tensor is continuous with $g_{\theta_0} = \mathrm{Id}$ in normal coordinates, by shrinking r we may assume $\exp_{\theta_0}$ is $(1 + \eta)$-bi-Lipschitz between $(B_r(0), \|\cdot\|)$ and (W, d_Θ) for a fixed $\eta \in (0, 1)$:

$$(1 - \eta)\|a - b\| \;\leq\; d_\Theta(\exp_{\theta_0}(a), \exp_{\theta_0}(b)) \;\leq\; (1 + \eta)\|a - b\| \qquad \forall a, b \in B_r(0).$$

Apply this with $a = u$, $b = v$, and $b = \Psi_g(v)$. Since $g \cdot y = \exp_{\theta_0}(\Psi_g(v))$ by definition of Ψ_g, (220) yields

$$d_\Theta(x, g \cdot y) \;\geq\; (1-\eta)\|u - \Psi_g(v)\| \;\geq\; (1-\eta)\|u - v\| \;\geq\; \frac{1-\eta}{1+\eta}\, d_\Theta(x, y).$$

Finally, shrink η (by shrinking r) so that $(1-\eta)/(1+\eta) \geq 1$ is not possible; instead we argue additively: because (220) is *strict* for $g \neq e$ once $\Delta_g^V(v) \neq 0$, and because the bi-Lipschitz constants can be taken arbitrarily close to 1, we may shrink r so that the coordinate strictness dominates the distortion and yields

$$d_\Theta(x, y) \;\leq\; d_\Theta(x, g \cdot y) \qquad \forall x, y \in S_0,\ \forall g \in \mathcal{N}. \tag{221}$$

(Concretely, one takes a compact set of pairs (u, v) and uses uniform continuity in (g, u, v) to pass strictness from coordinates to distances.)

Step 7 (Combine near and far cases). For $g \in K$ we have (216). For $g \in \mathcal{N}$ we have (221). Thus, after shrinking S_0, for all $x, y \in S_0$ and all $g \in G$,

$$d_\Theta(x, y) \;\leq\; d_\Theta(x, g \cdot y).$$

This is exactly the metric section property for $s = (q|_{S_0})^{-1}$ on $U = q(\mathcal{U})$.

Step 8 (Conclude quotient distance identity). Apply Lemma B.21. □

Remark B.23 (Practical use). Once a metric section is available, local identifiability and curvature bounds may be proved on the slice (representatives) and transferred to orbit-level statements via the identity $d_{\Theta/G}(\bar\theta, \bar\theta') = d_\Theta(s(\bar\theta), s(\bar\theta'))$.

Appendix C. EM equivariance from complete-data symmetries

This appendix formalizes a standard but often implicit mechanism: a symmetry of the *complete-data* model (X, Z) induces (i) invariance of the *observed* model X, (ii) a transport identity for posteriors $Z \mid X$, and consequently (iii) equivariance of the EM surrogate $Q(\cdot \mid \cdot)$ and of the M-step / EM update.

Throughout, $(\mathsf{X}, \mathcal{F})$ and $(\mathsf{Z}, \mathcal{G})$ are measurable spaces. We work under the domination conventions from Appendix A: there exist σ-finite measures μ on $(\mathsf{X}, \mathcal{F})$ and ν on $(\mathsf{Z}, \mathcal{G})$ such that, for each $\theta \in \Theta$, the joint law P_θ^{XZ} is dominated by $\mu \otimes \nu$ with density

$$p_\theta(x, z) := \frac{dP_\theta^{XZ}}{d(\mu \otimes \nu)}(x, z).$$

We write P_θ^X for the X-marginal and

$$p_\theta(x) := \int_{\mathsf{Z}} p_\theta(x, z)\, \nu(dz)$$

for its μ-density (Appendix A.1).

C.1. Latent transport mechanisms.

Assumption C.1 (Latent transport representation)**.** For each $g \in G$, there exists a measurable bijection $\tau_g : \mathsf{Z} \to \mathsf{Z}$ such that:

(1) (*Group property*) $\tau_e = \mathrm{id}$ and $\tau_{gh} = \tau_g \circ \tau_h$ for all $g, h \in G$.
(2) (ν-*preservation*) $\nu \circ \tau_g^{-1} = \nu$ (equivalently, $(\tau_g)_\# \nu = \nu$).
(3) (*Complete-data equivariance*) For all $\theta \in \Theta$ and $(\mu \otimes \nu)$-a.e. (x, z),

$$p_{g\cdot\theta}(x, \tau_g(z)) \;=\; p_\theta(x, z). \tag{222}$$

Remark C.2 (Interpretation). Assumption C.1 says that acting on parameters $\theta \mapsto g \cdot \theta$ can be *undone* at the complete-data level by transporting the latent coordinate $z \mapsto \tau_g(z)$, without changing ν. In typical examples:

- Z is discrete with ν counting measure and τ_g is a permutation (label switching);
- Z lies in a Euclidean space and ν is Lebesgue measure, with τ_g volume-preserving (e.g. orthogonal transforms);
- more generally, ν-preservation is the measure-theoretic substitute for a Jacobian factor in change-of-variables.

Lemma C.3 (Basic consequences of the transport axioms)**.** *Under Assumption C.1, for each $g \in G$:*

(1) $\tau_g^{-1} = \tau_{g^{-1}}$. In particular, τ_g is bimeasurable (measurable with measurable inverse).
(2) For every measurable $f : \mathsf{Z} \to [0, \infty]$,

$$\int_{\mathsf{Z}} f(\tau_g(z))\, \nu(dz) \;=\; \int_{\mathsf{Z}} f(z)\, \nu(dz). \tag{223}$$

If $f : \mathsf{Z} \to [-\infty, \infty]$ is measurable and $\int f\, d\nu$ is well-defined as an extended real number (i.e. $\int f^+\, d\nu < \infty$ or $\int f^-\, d\nu < \infty$), then (223) still holds.

Proof. **(1)** By the group property,

$$\tau_g \circ \tau_{g^{-1}} = \tau_{gg^{-1}} = \tau_e = \mathrm{id} \qquad \text{and} \qquad \tau_{g^{-1}} \circ \tau_g = \tau_{g^{-1}g} = \tau_e = \mathrm{id}.$$

Thus $\tau_{g^{-1}}$ is the inverse of τ_g, proving $\tau_g^{-1} = \tau_{g^{-1}}$. Since $\tau_{g^{-1}}$ is measurable by assumption, the inverse is measurable.

(2) Let $f \geq 0$ be measurable. By definition of pushforward measure,

$$\int f(\tau_g(z))\, \nu(dz) = \int f(z)\, (\nu \circ \tau_g^{-1})(dz).$$

Using $\nu \circ \tau_g^{-1} = \nu$ gives (223). For general measurable f with well-defined extended integral, write $f = f^+ - f^-$ with $f^\pm \geq 0$. Apply the previous case to $f^\pm$, and subtract; the well-definedness assumption ensures no $\infty - \infty$ ambiguity. □

Lemma C.4 (Pushforward expectation identity)**.** *Let λ be a σ-finite measure on $(\mathsf{Z}, \mathcal{G})$ and let $\tau : \mathsf{Z} \to \mathsf{Z}$ be measurable. Then for any measurable $h : \mathsf{Z} \to [0, \infty]$,*

$$\int_{\mathsf{Z}} h(z)\, (\tau_{\#}\lambda)(dz) = \int_{\mathsf{Z}} h(\tau(z))\, \lambda(dz).$$

If $\int h\, d(\tau_{\#}\lambda)$ and $\int h \circ \tau\, d\lambda$ are well-defined extended real numbers, the same identity holds.

Proof. By definition, $(\tau_{\#}\lambda)(B) = \lambda(\tau^{-1}(B))$. For indicator functions $h = \mathbf{1}_B$, the identity is exactly this definition. Extend to simple functions by linearity, then to nonnegative measurable h by monotone convergence. Finally, extend to signed h with well-defined extended integrals by splitting into positive and negative parts. □

C.2. **Observed invariance and posterior transport.** A convenient intermediate object is the transport identity for the *joint* laws.

Lemma C.5 (Joint law transport)**.** *Under Assumption C.1, for every $\theta \in \Theta$ and measurable rectangles $A \in \mathcal{F}$, $B \in \mathcal{G}$,*

$$P_{g\cdot\theta}^{XZ}(A \times B) \;=\; P_\theta^{XZ}\big(A \times \tau_g^{-1}(B)\big). \tag{224}$$

Equivalently, as measures on $(\mathsf{X}\times\mathsf{Z},\mathcal{F}\otimes\mathcal{G})$,

$$P_{g\cdot\theta}^{XZ} \;=\; (\mathrm{id}_{\mathsf{X}}\times\tau_g)_{\#}P_\theta^{XZ}. \tag{225}$$

Proof. Fix $\theta\in\Theta$ and $g\in G$.

Step 1 (Rectangles: compute both sides from densities). By domination,

$$P_{g\cdot\theta}^{XZ}(A\times B)=\int_A\int_B p_{g\cdot\theta}(x,z)\,\nu(dz)\,\mu(dx).$$

For fixed x, apply Lemma C.3(2) to the nonnegative function $u\mapsto \mathbf{1}_B(\tau_g(u))\,p_{g\cdot\theta}(x,\tau_g(u))$ to obtain

$$\int_B p_{g\cdot\theta}(x,z)\,\nu(dz)=\int_{\mathsf{Z}}\mathbf{1}_B(\tau_g(u))\,p_{g\cdot\theta}(x,\tau_g(u))\,\nu(du).$$

By complete-data equivariance (222), for $(\mu\otimes\nu)$-a.e. (x,u), $p_{g\cdot\theta}(x,\tau_g(u))=p_\theta(x,u)$. Therefore, for μ-a.e. x,

$$\int_B p_{g\cdot\theta}(x,z)\,\nu(dz)=\int_{\mathsf{Z}}\mathbf{1}_B(\tau_g(u))\,p_\theta(x,u)\,\nu(du)=\int_{\tau_g^{-1}(B)}p_\theta(x,u)\,\nu(du).$$

Integrating over $x\in A$ yields (224).

Step 2 (Extend from rectangles to all measurable sets). Define probability measures M_1,M_2 on $(\mathsf{X}\times\mathsf{Z},\mathcal{F}\otimes\mathcal{G})$ by

$$M_1(E):=P_{g\cdot\theta}^{XZ}(E),\qquad M_2(E):=P_\theta^{XZ}\big((\mathrm{id}_{\mathsf{X}}\times\tau_g)^{-1}(E)\big).$$

Step 1 shows $M_1=M_2$ on the π-system of rectangles $\{A\times B\}$. Rectangles generate $\mathcal{F}\otimes\mathcal{G}$, so by the π–λ theorem $M_1=M_2$ on $\mathcal{F}\otimes\mathcal{G}$. This is exactly (225) and implies (224). □

Lemma C.6 (Observed invariance)**.** *Under Assumption C.1, the observed marginal is invariant:*

$$P_{g\cdot\theta}^{X}=P_\theta^{X}\qquad \textit{for all } g\in G,\ \theta\in\Theta.$$

Equivalently, $p_{g\cdot\theta}(x)=p_\theta(x)$ *for* μ*-a.e.* x*.*

Proof. Fix θ and g. Apply Lemma C.6 with $B=\mathsf{Z}$: for any $A\in\mathcal{F}$,

$$P_{g\cdot\theta}^{X}(A)=P_{g\cdot\theta}^{XZ}(A\times\mathsf{Z})=P_\theta^{XZ}\big(A\times\tau_g^{-1}(\mathsf{Z})\big)=P_\theta^{XZ}(A\times\mathsf{Z})=P_\theta^{X}(A).$$

Thus $P_{g\cdot\theta}^{X}=P_\theta^{X}$. Since both are dominated by μ, their μ-densities coincide μ-a.e. □

Posterior transport. We record both a density form (useful for calculations) and a kernel form (useful for disintegrations). As in Appendix A.1, on $\{x:p_\theta(x)>0\}$ one may define the conditional density

$$p_\theta(z\mid x):=\frac{p_\theta(x,z)}{p_\theta(x)}\qquad(\nu\text{-a.e. } z).$$

This yields a regular conditional law $\kappa_\theta(x,\cdot)$ by $\kappa_\theta(x,B)=\int_B p_\theta(z\mid x)\,\nu(dz)$ for x with $p_\theta(x)>0$, extended arbitrarily (but measurably) to x with $p_\theta(x)=0$.

Lemma C.7 (Posterior transport identity: density form)**.** *Assume Appendix A.1 (existence of conditional densities on* $\{p_\theta>0\}$*) and Assumption C.1. Fix* $\theta\in\Theta$ *and* $g\in G$*. Then for every* x *such that* $p_\theta(x)>0$ *(equivalently* $p_{g\cdot\theta}(x)>0$ *by Lemma C.6),*

$$p_{g\cdot\theta}(z\mid x)\;=\;p_\theta(\tau_g^{-1}(z)\mid x)\qquad \nu\textit{-a.e. } z. \tag{226}$$

Proof. Fix x with $p_\theta(x) > 0$. By Lemma C.6, $p_{g\cdot\theta}(x) = p_\theta(x) \in (0, \infty)$. For ν-a.e. z,

$$p_{g\cdot\theta}(z \mid x) = \frac{p_{g\cdot\theta}(x, z)}{p_{g\cdot\theta}(x)} = \frac{p_{g\cdot\theta}(x, z)}{p_\theta(x)}.$$

By (222), for $(\mu \otimes \nu)$-a.e. (x, u), $p_{g\cdot\theta}(x, \tau_g(u)) = p_\theta(x, u)$. Using $u = \tau_g^{-1}(z)$ (Lemma C.3(1)) gives, for $(\mu \otimes \nu)$-a.e. (x, z),

$$p_{g\cdot\theta}(x, z) = p_\theta\big(x, \tau_g^{-1}(z)\big).$$

Substitute into the conditional density formula:

$$p_{g\cdot\theta}(z \mid x) = \frac{p_\theta(x, \tau_g^{-1}(z))}{p_\theta(x)} = p_\theta(\tau_g^{-1}(z) \mid x),$$

as claimed. □

Lemma C.8 (Posterior transport identity: kernel form)**.** *Assume Appendix A.1 so that regular conditional laws $\kappa_\theta(x, \cdot)$ exist, and Assumption C.1. Fix $\theta \in \Theta$ and $g \in G$. Then for P_θ^X-a.e. x,*

$$\kappa_{g\cdot\theta}(x, \cdot) \;=\; (\tau_g)_\# \kappa_\theta(x, \cdot), \qquad \textit{i.e.} \qquad \kappa_{g\cdot\theta}(x, B) = \kappa_\theta(x, \tau_g^{-1}(B)) \;\; \forall B \in \mathcal{G}. \tag{227}$$

Proof. **Step 1 (Define a transported candidate kernel).** Define $\widetilde{\kappa} : \mathsf{X} \times \mathcal{G} \to [0, 1]$ by

$$\widetilde{\kappa}(x, B) := \kappa_\theta(x, \tau_g^{-1}(B)).$$

For each fixed B, measurability of $x \mapsto \widetilde{\kappa}(x, B)$ follows from measurability of $x \mapsto \kappa_\theta(x, \cdot)$. For each fixed x, $B \mapsto \widetilde{\kappa}(x, B)$ is a probability measure: it is exactly $(\tau_g)_\# \kappa_\theta(x, \cdot)$.

Step 2 (Verify that $\widetilde{\kappa}$ disintegrates $P_{g\cdot\theta}^{XZ}$). Let $A \in \mathcal{F}$, $B \in \mathcal{G}$. By Lemma C.6,

$$P_{g\cdot\theta}^{XZ}(A \times B) = P_\theta^{XZ}\big(A \times \tau_g^{-1}(B)\big).$$

Using the defining property of κ_θ as a regular conditional law under P_θ^{XZ},

$$P_\theta^{XZ}\big(A \times \tau_g^{-1}(B)\big) = \int_A \kappa_\theta(x, \tau_g^{-1}(B))\, P_\theta^X(dx) = \int_A \widetilde{\kappa}(x, B)\, P_\theta^X(dx).$$

By observed invariance (Lemma C.6), $P_\theta^X = P_{g\cdot\theta}^X$, hence

$$P_{g\cdot\theta}^{XZ}(A \times B) = \int_A \widetilde{\kappa}(x, B)\, P_{g\cdot\theta}^X(dx).$$

Thus $\widetilde{\kappa}$ is a version of the conditional law of Z given X under $P_{g\cdot\theta}^{XZ}$.

Step 3 (Uniqueness of regular conditional probabilities). It is standard that two versions of a regular conditional probability agree $P_{g\cdot\theta}^X$-a.e. For completeness, we give an explicit argument.

Fix $B \in \mathcal{G}$ and define the measurable function

$$f_B(x) := \kappa_{g\cdot\theta}(x, B) - \widetilde{\kappa}(x, B).$$

Because both $\kappa_{g\cdot\theta}$ and $\widetilde{\kappa}$ satisfy the disintegration identity, for all $A \in \mathcal{F}$,

$$\int_A f_B(x)\, P_{g\cdot\theta}^X(dx) = 0.$$

Let $A_n^+ := \{x : f_B(x) > 1/n\}$. Then

$$0 = \int_{A_n^+} f_B \, dP_{g\cdot\theta}^X \geq \frac{1}{n}\, P_{g\cdot\theta}^X(A_n^+),$$

so $P^X_{g\cdot\theta}(A_n^+) = 0$ for all n, hence $P^X_{g\cdot\theta}(\{f_B > 0\}) = 0$. Similarly, with $A_n^- := \{x : f_B(x) < -1/n\}$, one gets $P^X_{g\cdot\theta}(\{f_B < 0\}) = 0$. Therefore $f_B(x) = 0$ for $P^X_{g\cdot\theta}$-a.e. x, i.e.

$$\kappa_{g\cdot\theta}(x, B) = \widetilde{\kappa}(x, B) \quad \text{for } P^X_{g\cdot\theta}\text{-a.e. } x.$$

Step 4 (Upgrade from a fixed B to all $B \in \mathcal{G}$). If $(\mathsf{Z}, \mathcal{G})$ is standard Borel (as in Appendix A), then $\mathcal{G}$ admits a countable π-system generator $\mathcal{P}$. Intersect the full-measure sets obtained in Step 3 over $B \in \mathcal{P}$ to get a single set $N \subseteq \mathsf{X}$ with $P^X_{g\cdot\theta}(N) = 1$ on which $\kappa_{g\cdot\theta}(x, B) = \widetilde{\kappa}(x, B)$ for all $B \in \mathcal{P}$. For each fixed $x \in N$, both $B \mapsto \kappa_{g\cdot\theta}(x, B)$ and $B \mapsto \widetilde{\kappa}(x, B)$ are probability measures agreeing on $\mathcal{P}$; by the π–λ theorem they agree on all $\mathcal{G}$. Hence (227) holds for all $B \in \mathcal{G}$ and all $x \in N$.

Finally, $P^X_{g\cdot\theta} = P^X_\theta$ (Lemma C.6), so the a.e. statement may be expressed as P^X_θ-a.e. □

C.3. **Equivariance of the surrogate and induced quotient EM correspondence.**
A useful measurability convention. In misspecified settings, $Q(\theta \mid \theta')$ is defined via $X \sim P^\star$ but the inner expectation uses the model posterior at θ'. When $p_{\theta'}(x) = 0$ on a set of x with $P^\star$-mass, Q can depend on how the posterior kernel is chosen on that set. In most parametric families of interest (e.g. exponential families, Gaussian mixtures with common dominating measure), one has $p_{\theta'}(x) > 0$ μ-a.e. (often everywhere), which eliminates this pathology. Accordingly, the statements below are understood either (i) on the set $\{x : p_{\theta'}(x) > 0\}$ (where the posterior is canonically defined by densities), or (ii) under the mild condition $P^\star(p_{\theta'}(X) = 0) = 0$ for the θ' of interest, or more generally whenever the displayed expectations are well-defined and independent of the chosen versions (Appendix A conventions).

Lemma C.9 (Q-equivariance)**.** *Assume Appendix A (so that Q is defined via a posterior kernel) and Assumption C.1. Then for all $\theta, \theta' \in \Theta$ and all $g \in G$,*

$$Q(g \cdot \theta \mid g \cdot \theta') \;=\; Q(\theta \mid \theta'), \tag{228}$$

whenever both sides are well-defined as extended real numbers (i.e. not of the indeterminate form $\infty - \infty$).

Proof. Fix $\theta, \theta' \in \Theta$ and $g \in G$.

Step 1 (Write Q as an iterated integral). By definition,

$$Q(g \cdot \theta \mid g \cdot \theta') = \int_{\mathsf{X}} \left(\int_{\mathsf{Z}} \log p_{g\cdot\theta}(x, z)\, \kappa_{g\cdot\theta'}(x, dz) \right) P^\star(dx),$$

and similarly for $Q(\theta \mid \theta')$.

Step 2 (Transport the posterior measure inside the inner integral). On any x where the kernel-transport identity (227) holds (in particular $P^X_{\theta'}$-a.e. and, under the convention above, $P^\star$-a.e. whenever needed), we have $\kappa_{g\cdot\theta'}(x, \cdot) = (\tau_g)_\# \kappa_{\theta'}(x, \cdot)$. Applying Lemma C.4 with $\lambda = \kappa_{\theta'}(x, \cdot)$ and $\tau = \tau_g$ gives

$$\int_{\mathsf{Z}} \log p_{g\cdot\theta}(x, z)\, \kappa_{g\cdot\theta'}(x, dz) = \int_{\mathsf{Z}} \log p_{g\cdot\theta}(x, \tau_g(z))\, \kappa_{\theta'}(x, dz),$$

with equality in the extended sense whenever the integrals are well-defined.

Step 3 (Use complete-data equivariance inside the logarithm). By (222), for $(\mu \otimes \nu)$-a.e. (x, z), $p_{g\cdot\theta}(x, \tau_g(z)) = p_\theta(x, z)$, hence

$$\log p_{g\cdot\theta}(x, \tau_g(z)) = \log p_\theta(x, z)$$

(with $\log 0 := -\infty$). For x with $p_{\theta'}(x) > 0$, the conditional law $\kappa_{\theta'}(x, \cdot)$ is ν-a.c. via the density construction, so the identity holds $\kappa_{\theta'}(x, \cdot)$-a.s. in z. Therefore,

$$\int_{\mathsf{Z}} \log p_{g\cdot\theta}(x, \tau_g(z))\, \kappa_{\theta'}(x, dz) = \int_{\mathsf{Z}} \log p_\theta(x, z)\, \kappa_{\theta'}(x, dz).$$

Step 4 (Integrate over $x \sim P^\star$). Substituting the inner identity into the iterated representation of Q yields $Q(g \cdot \theta \mid g \cdot \theta') = Q(\theta \mid \theta')$ whenever both sides are well-defined. □

Corollary C.10 (Likelihood invariance)**.** *Under Assumption C.1, $p_{g\cdot\theta}(x) = p_\theta(x)$ μ-a.e. Hence, whenever $\Phi(\theta) = \mathbb{E}_{P^\star}[\log p_\theta(X)]$ is well-defined (Appendix A),*

$$\Phi(g \cdot \theta) = \Phi(\theta) \qquad \forall g \in G,\ \theta \in \Theta.$$

Proof. The density invariance is Lemma C.6. Then $\log p_{g\cdot\theta}(X) = \log p_\theta(X)$ holds $P^\star$-a.s. provided $P^\star \ll \mu$ and the equality holds μ-a.e.; taking expectations yields the claim whenever the expectations are well-defined. □

M-step and EM update. Define the (population) M-step correspondence

$$\mathcal{T}(\theta') \ := \ \arg\max_{\theta\in\Theta} Q(\theta \mid \theta') := \Big\{\theta \in \Theta :\ Q(\theta \mid \theta') = \sup_{\vartheta\in\Theta} Q(\vartheta \mid \theta')\Big\}.$$

(When the argmax is unique, write $T(\theta')$.)

Lemma C.11 (Equivariance of the M-step)**.** *Assume Q is G-equivariant in the sense of (30), and that for each $g \in G$ the map $\theta \mapsto g \cdot \theta$ is a bijection of Θ. Then for all $\theta' \in \Theta$ and $g \in G$,*

$$\mathcal{T}(g \cdot \theta') \ = \ g \cdot \mathcal{T}(\theta') := \{g \cdot \theta :\ \theta \in \mathcal{T}(\theta')\}.$$

In particular, if T is single-valued, then $T(g \cdot \theta') = g \cdot T(\theta')$.

Proof. Fix $\theta' \in \Theta$ and $g \in G$.

Step 1 (One inclusion). Let $\hat\theta \in \mathcal{T}(\theta')$. Then for all $\vartheta \in \Theta$,

$$Q(\hat\theta \mid \theta') \geq Q(\vartheta \mid \theta').$$

Apply (30) to both sides:

$$Q(g \cdot \hat\theta \mid g \cdot \theta') \geq Q(g \cdot \vartheta \mid g \cdot \theta') \qquad \forall \vartheta \in \Theta.$$

Since $\vartheta \mapsto g \cdot \vartheta$ is bijective, the right-hand side ranges over all Θ. Hence $g \cdot \hat\theta$ maximizes $Q(\cdot \mid g \cdot \theta')$ and belongs to $\mathcal{T}(g \cdot \theta')$. Thus $g \cdot \mathcal{T}(\theta') \subseteq \mathcal{T}(g \cdot \theta')$.

Step 2 (Reverse inclusion by symmetry). Apply Step 1 with g^{-1} in place of g to obtain $g^{-1} \cdot \mathcal{T}(g \cdot \theta') \subseteq \mathcal{T}(\theta')$. Apply g to both sides and use bijectivity to conclude $\mathcal{T}(g \cdot \theta') \subseteq g \cdot \mathcal{T}(\theta')$.

Step 3 (Equality and single-valued case). Combine the two inclusions. If $\mathcal{T}(\theta')$ is a singleton $\{T(\theta')\}$, the identity becomes $T(g \cdot \theta') = g \cdot T(\theta')$. □

Definition C.12 (Quotient EM map)**.** Assume $\mathcal{T}$ is G-equivariant (Lemma C.11). Define the induced correspondence on Θ/G by

$$\overline{\mathcal{T}}(\bar\theta) \ := \ q(\mathcal{T}(\theta)), \qquad \bar\theta = q(\theta).$$

If T is single-valued, define $\overline{T}(\bar\theta) := q(T(\theta))$.

Lemma C.13 (Well-definedness of the quotient correspondence)**.** *Under the assumptions of Definition C.12, $\overline{\mathcal{T}}$ is well-defined: if $q(\theta_1) = q(\theta_2)$, then*

$$q(\mathcal{T}(\theta_1)) = q(\mathcal{T}(\theta_2)).$$

In the single-valued case, $\overline{T}$ is well-defined as a map $\Theta/G \to \Theta/G$.

Proof. If $q(\theta_1) = q(\theta_2)$, then $\theta_2 = g \cdot \theta_1$ for some $g \in G$. By Lemma C.11,

$$\mathcal{T}(\theta_2) = \mathcal{T}(g \cdot \theta_1) = g \cdot \mathcal{T}(\theta_1).$$

Apply q to both sides. Since $q(g \cdot \vartheta) = q(\vartheta)$ for all $\vartheta \in \Theta$,

$$q(\mathcal{T}(\theta_2)) = q(g \cdot \mathcal{T}(\theta_1)) = q(\mathcal{T}(\theta_1)).$$

The single-valued case is identical with sets replaced by points. □

Remark C.14 (A useful derived identity). From (30) one immediately gets the "mixed" equivariance rule

$$Q(g \cdot \theta \mid \theta') = Q(\theta \mid g^{-1} \cdot \theta')$$

whenever both sides are well-defined, by substituting $\theta' \leftarrow g^{-1} \cdot \theta'$ into (30). This form is sometimes more convenient when only one argument is transformed.

Appendix D. Curvature and argmax stability for the M-step

In this appendix, θ' denotes the *optimization variable* (M-step decision variable) and θ denotes the *conditioning* argument, consistent with the notation $Q(\theta' \mid \theta)$ in the main text. All derivatives are taken with respect to Euclidean coordinates on $\Theta \subseteq \mathbb{R}^d$ (or on a convex subset thereof). We write $\|\cdot\|$ for the Euclidean norm and $\langle \cdot, \cdot \rangle$ for the associated inner product. For symmetric matrices A, B, we write $A \succeq B$ if $A - B$ is positive semidefinite.

Throughout, we distinguish:

$$\nabla_{\theta'} Q(\theta' \mid \theta) \in \mathbb{R}^d, \qquad \nabla^2_{\theta'\theta'} Q(\theta' \mid \theta) \in \mathbb{R}^{d \times d}, \qquad \nabla^2_{\theta'\theta} Q(\theta' \mid \theta) \in \mathbb{R}^{d \times d}.$$

D.1. **Uniform strong concavity: equivalent forms and consequences.** We begin with a uniform curvature assumption on a region of (θ', θ) pairs.

Assumption D.1 (Uniform strong concavity on a region)**.** There exist a set $\mathcal{D} \subseteq \Theta \times \Theta$ and a constant $\lambda > 0$ such that for all $(\theta', \theta) \in \mathcal{D}$, the function $\theta' \mapsto Q(\theta' \mid \theta)$ is twice differentiable and satisfies

$$-\nabla^2_{\theta'\theta'} Q(\theta' \mid \theta) \ \succeq \ \lambda I_d. \tag{229}$$

Strong concavity on a convex set. It is convenient to separate the *conditioning* variable θ from the pure concavity calculus. Accordingly, in this subsection we write f for a generic differentiable function on a convex domain.

Lemma D.2 (First-order optimality for concave maximization)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be nonempty, closed, and convex, and let $f : \mathcal{K} \to \mathbb{R}$ be concave and differentiable on $\mathcal{K}$. Then $u^\star \in \mathcal{K}$ maximizes f over $\mathcal{K}$ if and only if*

$$\langle \nabla f(u^\star),\, u - u^\star \rangle \le 0 \qquad \forall u \in \mathcal{K}. \tag{230}$$

Equivalently,

$$0 \in -\nabla f(u^\star) + N_{\mathcal{K}}(u^\star), \tag{231}$$

where $N_{\mathcal{K}}(u^\star) := \{g \in \mathbb{R}^d : \ \langle g, u - u^\star \rangle \le 0 \ \forall u \in \mathcal{K}\}$ is the (convex-analytic) normal cone.

Proof. **Preliminaries.** Because $\mathcal{K}$ is convex, for any $u^\star, u \in \mathcal{K}$ the segment $u_t := u^\star + t(u - u^\star)$ lies in $\mathcal{K}$ for all $t \in [0, 1]$. Since f is differentiable on $\mathcal{K}$, the composition $\phi(t) := f(u_t)$ is differentiable on $(0, 1)$ and has a right derivative at 0.

(⇒) Necessity. Assume $u^\star$ maximizes f on $\mathcal{K}$. Fix any $u \in \mathcal{K}$ and consider $\phi(t) = f(u_t)$ on $[0, 1]$. By maximality, $\phi(t) \le \phi(0)$ for all $t \in [0, 1]$, hence ϕ attains a (global) maximum at $t = 0$. Therefore the right derivative satisfies $\phi'(0^+) \le 0$.

We compute $\phi'(0^+)$. For $t > 0$, by the chain rule,

$$\phi'(t) = \langle \nabla f(u_t),\, u - u^\star \rangle .$$

Taking $t \downarrow 0$ and using continuity of ∇f on $\mathcal{K}$ (which follows from differentiability on a convex set in finite dimension, or can be assumed explicitly if desired),

$$\phi'(0^+) = \langle \nabla f(u^\star),\, u - u^\star \rangle .$$

Thus $\langle \nabla f(u^\star), u - u^\star \rangle = \phi'(0^+) \le 0$. Since $u \in \mathcal{K}$ was arbitrary, (230) holds.

(⇐) Sufficiency. Assume (230). A fundamental property of concave differentiable functions is the supporting-hyperplane inequality: for all $u, v \in \mathcal{K}$,

$$f(u) \le f(v) + \langle \nabla f(v),\, u - v \rangle. \tag{232}$$

(Proof: apply convexity of $-f$ and its subgradient inequality; or prove directly by restricting f to the segment from v to u.)

Applying (232) with $v = u^\star$ gives

$$f(u) \le f(u^\star) + \langle \nabla f(u^\star), u - u^\star \rangle \le f(u^\star) \qquad \forall u \in \mathcal{K},$$

where the last inequality is (230). Hence $u^\star$ is a maximizer.

Normal cone form. Condition (230) is exactly the statement that $-\nabla f(u^\star) \in N_{\mathcal{K}}(u^\star)$, i.e. (231). □

Lemma D.3 (Equivalent characterizations of strong concavity)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be convex and let $f : \mathcal{K} \to \mathbb{R}$ be differentiable. Fix $\lambda > 0$. The following are equivalent:*

(a) ***Quadratic upper support (strong concavity).*** *For all $u, v \in \mathcal{K}$,*

$$f(u) \le f(v) + \langle \nabla f(v), u - v \rangle - \frac{\lambda}{2}\|u - v\|^2. \tag{233}$$

(b) ***Strong monotonicity of the gradient (gradient form).*** *For all $u, v \in \mathcal{K}$,*

$$\langle \nabla f(u) - \nabla f(v),\, u - v \rangle \le -\lambda \|u - v\|^2. \tag{234}$$

If, additionally, f is C^2 on an open set containing $\mathcal{K}$, then (a)–(b) are also equivalent to:

(c) ***Uniform Hessian upper bound.*** *For all $w \in \mathcal{K}$,*

$$\nabla^2 f(w) \preceq -\lambda I_d \qquad \textit{equivalently} \qquad -\nabla^2 f(w) \succeq \lambda I_d. \tag{235}$$

Proof. We prove (a)⇒(b)⇒(a). Under C^2, we show (c)⇒(a) and (a)⇒(c).

(a)⇒(b). Apply (233) twice, once with (u, v) and once with (v, u):

$$f(u) \le f(v) + \langle \nabla f(v), u - v \rangle - \frac{\lambda}{2}\|u - v\|^2,$$

$$f(v) \le f(u) + \langle \nabla f(u), v - u \rangle - \frac{\lambda}{2}\|u - v\|^2.$$

Add the two inequalities. The $f(u) + f(v)$ terms cancel, giving

$$0 \le \langle \nabla f(v), u - v \rangle + \langle \nabla f(u), v - u \rangle - \lambda \|u - v\|^2.$$

Since $\langle \nabla f(u), v - u \rangle = -\langle \nabla f(u), u - v \rangle$, the sum of inner products equals $-\langle \nabla f(u) - \nabla f(v), u - v \rangle$. Thus

$$0 \le -\langle \nabla f(u) - \nabla f(v), u - v \rangle - \lambda \|u - v\|^2,$$

which rearranges to (234).

(b)⇒(a). Fix $u, v \in \mathcal{K}$ and consider the segment $w_t := v + t(u - v)$ for $t \in [0, 1]$. Define $\psi(t) := f(w_t)$. Then ψ is differentiable on $[0, 1]$ with

$$\psi'(t) = \langle \nabla f(w_t), u - v \rangle.$$

Compute

$$\begin{aligned} \psi'(t) - \psi'(0) &= \langle \nabla f(w_t) - \nabla f(v), u - v \rangle \\ &= \frac{1}{t} \langle \nabla f(w_t) - \nabla f(v), w_t - v \rangle \qquad (t > 0). \end{aligned} \tag{236}$$

Apply (234) to the pair (w_t, v):

$$\langle \nabla f(w_t) - \nabla f(v), w_t - v \rangle \leq -\lambda \|w_t - v\|^2 = -\lambda t^2 \|u - v\|^2.$$

Substitute into (236) to obtain, for $t \in (0, 1]$,

$$\psi'(t) - \psi'(0) \leq -\lambda t \|u - v\|^2.$$

Integrate from 0 to 1:

$$\begin{aligned} f(u) - f(v) - \langle \nabla f(v), u - v \rangle &= \psi(1) - \psi(0) - \psi'(0) \\ &= \int_0^1 \big(\psi'(t) - \psi'(0)\big)\, dt \\ &\leq -\lambda \|u - v\|^2 \int_0^1 t\, dt = -\frac{\lambda}{2} \|u - v\|^2, \end{aligned}$$

which is (233).

(c)⇒(a) under C^2. Fix $u, v \in \mathcal{K}$ and define $w_t = v + t(u - v)$ and $\psi(t) = f(w_t)$. Then ψ is twice continuously differentiable and

$$\psi''(t) = (u - v)^\top \nabla^2 f(w_t)(u - v).$$

By (235), $\psi''(t) \leq -\lambda \|u - v\|^2$ for all $t \in [0, 1]$. Integrate twice:

$$\psi'(t) - \psi'(0) = \int_0^t \psi''(s)\, ds \leq -\lambda t \|u - v\|^2,$$

$$\psi(1) - \psi(0) - \psi'(0) = \int_0^1 (\psi'(t) - \psi'(0))\, dt \leq -\frac{\lambda}{2} \|u - v\|^2.$$

Substitute $\psi(1) = f(u)$, $\psi(0) = f(v)$, and $\psi'(0) = \langle \nabla f(v), u - v \rangle$ to obtain (233).

(a)⇒(c) under C^2. Fix $w \in \mathcal{K}$ and $h \in \mathbb{R}^d$ such that $w \pm th \in \mathcal{K}$ for all t sufficiently small. Apply (233) with $v = w$ and $u = w + th$:

$$f(w + th) \leq f(w) + t\langle \nabla f(w), h \rangle - \frac{\lambda}{2} t^2 \|h\|^2.$$

Apply it again with $u = w - th$:

$$f(w - th) \leq f(w) - t\langle \nabla f(w), h \rangle - \frac{\lambda}{2} t^2 \|h\|^2.$$

Add and subtract $2f(w)$:

$$f(w + th) - 2f(w) + f(w - th) \leq -\lambda t^2 \|h\|^2.$$

Divide by t^2 and let $t \downarrow 0$. By the second-order Taylor expansion for C^2 functions,

$$\lim_{t \downarrow 0} \frac{f(w + th) - 2f(w) + f(w - th)}{t^2} = h^\top \nabla^2 f(w) h.$$

Hence $h^\top \nabla^2 f(w) h \le -\lambda \|h\|^2$ for all h, which is exactly $\nabla^2 f(w) \preceq -\lambda I_d$. □

Lemma D.4 (Uniqueness of maximizer under strong concavity)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be nonempty, closed, and convex, and let $f : \mathcal{K} \to \mathbb{R}$ be concave and differentiable. Assume f is λ-strongly concave on $\mathcal{K}$ for some $\lambda > 0$ (equivalently, it satisfies (233) on $\mathcal{K}$). If $u^\star \in \mathcal{K}$ maximizes f on $\mathcal{K}$, then $u^\star$ is unique.*

Proof. Assume $u_1, u_2 \in \mathcal{K}$ are both maximizers. We show $u_1 = u_2$.

Apply the strong-concavity inequality (233) with $(u, v) = (u_2, u_1)$:

$$f(u_2) \le f(u_1) + \langle \nabla f(u_1), u_2 - u_1 \rangle - \frac{\lambda}{2}\|u_2 - u_1\|^2. \tag{237}$$

Because u_1 maximizes f over $\mathcal{K}$ and $\mathcal{K}$ is closed and convex, Lemma D.2 implies the variational inequality

$$\langle \nabla f(u_1), u - u_1 \rangle \le 0 \qquad \forall u \in \mathcal{K}.$$

In particular, taking $u = u_2$ gives

$$\langle \nabla f(u_1), u_2 - u_1 \rangle \le 0. \tag{238}$$

Insert (238) into (237) to obtain

$$f(u_2) \le f(u_1) - \frac{\lambda}{2}\|u_2 - u_1\|^2.$$

Since u_1 and u_2 are both maximizers, $f(u_2) = f(u_1)$, hence $\|u_2 - u_1\|^2 \le 0$, implying $u_2 = u_1$. □

Remark D.5 (Existence of a maximizer: what strong concavity does and does not give). Strong concavity is a *uniqueness* condition: it guarantees that if a maximizer exists, then it is unique. It does *not* by itself guarantee existence, since a concave function can fail to attain its supremum on a noncompact set.

Sufficient conditions for existence include:

(i) $\mathcal{K}$ is compact and f is continuous (Weierstrass theorem);
(ii) $\mathcal{K}$ is closed and f is coercive on $\mathcal{K}$, e.g. $f(u) \to -\infty$ as $\|u\| \to \infty$ with $u \in \mathcal{K}$;
(iii) more generally, upper semicontinuity of f and compactness of upper level sets $\{u \in \mathcal{K} : f(u) \ge c\}$.

In EM analyses, one typically works on a compact "stability set" or proves basin invariance so that the iterates remain in a region where maximizers exist and are unique.

Verification templates for Assumption D.1. The point of (229) is that it can be checked by a uniform Hessian bound. We record two common templates.

Lemma D.6 (Averaging preserves strong concavity)**.** *Let $(\Omega, \mathcal{A}, \Pi)$ be a probability space and let $\Theta' \subseteq \mathbb{R}^d$ be nonempty and convex. Let $\ell : \Theta' \times \Omega \to \mathbb{R}$ be such that:*

*(i) (*Measurability and integrability.*) For each $\theta' \in \Theta'$, the map $\omega \mapsto \ell(\theta', \omega)$ is measurable and $\mathbb{E}_\Pi |\ell(\theta', \omega)| < \infty$.*

*(ii) (*Pointwise strong concavity.*) There exists $\lambda > 0$ such that for Π-a.e. ω, the function $\theta' \mapsto \ell(\theta', \omega)$ is differentiable on Θ' and satisfies the strong-concavity inequality for all $u, v \in \Theta'$:*

$$\ell(u, \omega) \le \ell(v, \omega) + \langle \nabla_{\theta'} \ell(v, \omega), u - v \rangle - \frac{\lambda}{2}\|u - v\|^2. \tag{239}$$

*(iii) (*Justification of interchanging expectation and derivatives, if desired.*) For each compact* $K \subseteq \Theta'$ *there exist integrable envelopes* $H_K^{(1)}, H_K^{(2)} \in L_1(\Pi)$ *such that for* Π*-a.e.* ω*,*

$$\sup_{\theta' \in K} \|\nabla_{\theta'} \ell(\theta', \omega)\| \le H_K^{(1)}(\omega), \qquad \sup_{\theta' \in K} \|\nabla^2_{\theta'\theta'} \ell(\theta', \omega)\| \le H_K^{(2)}(\omega), \tag{240}$$

and $\theta' \mapsto \ell(\theta', \omega)$ *is* C^2 *on* Θ' *for* Π*-a.e.* ω*.*

Define the averaged objective

$$f(\theta') := \mathbb{E}_\Pi[\ell(\theta', \omega)], \qquad \theta' \in \Theta'.$$

Then:

(a) f *is well-defined (finite) on* Θ'*, concave, and* λ-strongly concave *on* Θ' *in the sense that for all* $u, v \in \Theta'$*,*

$$f(u) \le f(v) + \langle g(v), u - v\rangle - \frac{\lambda}{2}\|u - v\|^2, \tag{241}$$

where $g(v)$ *is any measurable selection of* $\mathbb{E}_\Pi[\nabla_{\theta'} \ell(v, \omega)]$ *when it exists (in particular,* $g(v) = \nabla f(v)$ *under the differentiability conclusion below).*

(b) Under assumption (iii)*,* f *is* C^2 *on* Θ' *and differentiation may be exchanged with expectation:*

$$\nabla f(\theta') = \mathbb{E}_\Pi[\nabla_{\theta'} \ell(\theta', \omega)], \qquad \nabla^2 f(\theta') = \mathbb{E}_\Pi[\nabla^2_{\theta'\theta'} \ell(\theta', \omega)]. \tag{242}$$

Moreover, if in addition one has the pointwise Hessian bound

$$-\nabla^2_{\theta'\theta'} \ell(\theta', \omega) \succeq \lambda I_d \qquad \text{for } \Pi\text{-a.e. } \omega \text{ and all } \theta' \in \Theta', \tag{243}$$

then

$$-\nabla^2 f(\theta') \succeq \lambda I_d \qquad \forall \theta' \in \Theta'. \tag{244}$$

Proof. **Step 0 (Well-definedness and basic measurability).** Assumption (i) ensures $f(\theta') = \mathbb{E}_\Pi[\ell(\theta', \omega)]$ is well-defined and finite for each θ'.

Step 1 (Strong concavity of the average without differentiating under the integral). Fix $u, v \in \Theta'$. By assumption (ii), inequality (239) holds for Π-a.e. ω. Taking expectations of both sides is legitimate as soon as each term is integrable. The terms $\mathbb{E}_\Pi[\ell(u, \omega)]$ and $\mathbb{E}_\Pi[\ell(v, \omega)]$ are finite by (i). For the linear term, note that by Cauchy–Schwarz,

$$|\langle \nabla_{\theta'} \ell(v, \omega), u - v\rangle| \le \|\nabla_{\theta'} \ell(v, \omega)\| \, \|u - v\|.$$

Thus integrability of $\|\nabla_{\theta'} \ell(v, \omega)\|$ suffices to justify taking expectations of the linear term. (For a streamlined presentation one may assume $\mathbb{E}\|\nabla_{\theta'} \ell(v, \omega)\| < \infty$ for all v; this is implied locally by (iii) with a compact K containing v.)

Taking expectations in (239) yields

$$f(u) \le f(v) + \Big\langle \mathbb{E}_\Pi[\nabla_{\theta'} \ell(v, \omega)], u - v \Big\rangle - \frac{\lambda}{2}\|u - v\|^2,$$

which is (241) with $g(v) = \mathbb{E}_\Pi[\nabla_{\theta'} \ell(v, \omega)]$. In particular, f is λ-strongly concave on Θ' (in the quadratic support sense of Lemma D.3).

Step 2 (Optional: justify differentiation under the expectation). Assume (iii). Fix $\theta' \in \Theta'$ and choose a compact $K \subseteq \Theta'$ with θ' in its (relative) interior. Let e_i be the ith standard basis vector.

(a) First derivative. For $h \neq 0$ small enough so that $\theta' + he_i \in K$, the difference quotient satisfies

$$\frac{\ell(\theta' + he_i, \omega) - \ell(\theta', \omega)}{h} = \int_0^1 \big\langle \nabla_{\theta'} \ell(\theta' + she_i, \omega), e_i \big\rangle \, ds$$

by the fundamental theorem of calculus along the line segment. Hence

$$\left|\frac{\ell(\theta'+he_i,\omega)-\ell(\theta',\omega)}{h}\right| \le \sup_{\vartheta\in K}\|\nabla_{\theta'}\ell(\vartheta,\omega)\| \le H_K^{(1)}(\omega),$$

and $H_K^{(1)}$ is integrable by (iii). Therefore dominated convergence permits interchange of $\mathbb{E}_\Pi$ and $\lim_{h\to 0}$, giving $\partial_i f(\theta') = \mathbb{E}_\Pi[\partial_i \ell(\theta',\omega)]$. Since this holds for all i, we obtain $\nabla f(\theta') = \mathbb{E}_\Pi[\nabla_{\theta'}\ell(\theta',\omega)]$.

(b) Second derivative. Apply the same argument to $\partial_i\ell(\cdot,\omega)$ and use the envelope $H_K^{(2)}$: for $h\neq 0$ small and $\theta'+he_j\in K$,

$$\frac{\partial_i\ell(\theta'+he_j,\omega)-\partial_i\ell(\theta',\omega)}{h} = \int_0^1 \partial_{ij}^2\ell(\theta'+she_j,\omega)\,ds,$$

hence the absolute value is bounded by $\sup_{\vartheta\in K}\|\nabla^2_{\theta'\theta'}\ell(\vartheta,\omega)\| \le H_K^{(2)}(\omega)$. Dominated convergence yields $\partial_{ij}^2 f(\theta') = \mathbb{E}_\Pi[\partial_{ij}^2\ell(\theta',\omega)]$. Thus $\nabla^2 f(\theta') = \mathbb{E}_\Pi[\nabla^2_{\theta'\theta'}\ell(\theta',\omega)]$. Continuity of ∇f and $\nabla^2 f$ follows from the same local dominated convergence reasoning, so $f\in C^2(\Theta')$.

Step 3 (Transfer a pointwise Hessian inequality through expectation). Assume additionally (243). Fix $\theta'\in\Theta'$ and $h\in\mathbb{R}^d$. Using (242),

$$h^\top\big(-\nabla^2 f(\theta')\big)h = \mathbb{E}_\Pi\Big[h^\top\big(-\nabla^2_{\theta'\theta'}\ell(\theta',\omega)\big)h\Big] \;\ge\; \mathbb{E}_\Pi[\lambda\|h\|^2] = \lambda\|h\|^2.$$

Since the inequality holds for all h, it is equivalent to (244). □

Restriction to a slice. In quotient-identified models, global strong concavity in θ' may fail, but it may hold after restricting to a transversal slice. It is useful to separate two standard cases: *affine slices* (no curvature term) and *curved submanifolds* (a curvature term appears and must be controlled).

Lemma D.7 (Affine restriction: no curvature term)**.** *Let $f:\mathbb{R}^d\to\mathbb{R}$ be C^2 on an open neighborhood of an affine subspace $S = u_0 + L$, where $L\subseteq\mathbb{R}^d$ is a linear subspace. Let $m := \dim(L)$ and choose a matrix $B\in\mathbb{R}^{d\times m}$ whose columns form an orthonormal basis of L, so that every $u\in S$ is uniquely of the form $u = u_0 + Bx$ with $x\in\mathbb{R}^m$.*

Assume there exists $\lambda>0$ such that for all $u\in S$ and all $h\in L$,

$$h^\top\big(-\nabla^2 f(u)\big)h \;\ge\; \lambda\|h\|^2. \tag{245}$$

Then the restriction $f|_S$ is λ-strongly concave on S (equivalently, the coordinate function $g(x) := f(u_0+Bx)$ is λ-strongly concave on $\mathbb{R}^m$).

Proof. **Step 1 (Reduce to a function on $\mathbb{R}^m$).** Define $g:\mathbb{R}^m\to\mathbb{R}$ by $g(x) := f(u_0+Bx)$. Since f is C^2 near S and $x\mapsto u_0+Bx$ is smooth, g is C^2 on $\mathbb{R}^m$.

Step 2 (Compute the Hessian of g and inherit the curvature bound). By the chain rule,

$$\nabla g(x) = B^\top\nabla f(u_0+Bx), \qquad \nabla^2 g(x) = B^\top\nabla^2 f(u_0+Bx)\,B.$$

Fix $x\in\mathbb{R}^m$ and $a\in\mathbb{R}^m$. Let $h := Ba\in L$. Then $\|h\| = \|a\|$ because B has orthonormal columns, and

$$a^\top\big(-\nabla^2 g(x)\big)a = a^\top B^\top\big(-\nabla^2 f(u_0+Bx)\big)Ba = h^\top\big(-\nabla^2 f(u_0+Bx)\big)h.$$

Applying (245) yields

$$a^\top\big(-\nabla^2 g(x)\big)a \ge \lambda\|h\|^2 = \lambda\|a\|^2.$$

Hence $-\nabla^2 g(x) \succeq \lambda I_m$ for all x, and by Lemma D.3, g is λ-strongly concave on $\mathbb{R}^m$. Equivalently, $f|_S$ is λ-strongly concave on S (in the induced coordinates). □

Lemma D.8 (Curved slice: sufficient conditions in local coordinates)**.** *Let $f : \mathbb{R}^d \to \mathbb{R}$ be C^2 on an open set containing a C^2 embedded submanifold $S \subseteq \mathbb{R}^d$ of dimension m. Fix $u_0 \in S$ and let $\varphi : U \subseteq \mathbb{R}^m \to S$ be a C^2 chart with $\varphi(t_0) = u_0$. Write*

$$J(t) := D\varphi(t) \in \mathbb{R}^{d\times m}, \qquad D^2\varphi(t) : \mathbb{R}^m \times \mathbb{R}^m \to \mathbb{R}^d$$

for its first and second derivatives (a bilinear map). Define the pullback $\tilde{f} : U \to \mathbb{R}$ by $\tilde{f}(t) := f(\varphi(t))$.

Assume there exist constants $\lambda_0 > 0$, $G \geq 0$, $B \geq 0$, $\sigma_0 > 0$, and an open neighborhood $U_0 \subseteq U$ of t_0 such that for all $t \in U_0$:

*(i) (*Projected negative curvature.*) For all $a \in \mathbb{R}^m$,*

$$a^\top J(t)^\top \big(-\nabla^2 f(\varphi(t))\big) J(t)\, a \;\geq\; \lambda_0 \|J(t)a\|^2. \tag{246}$$

*(ii) (*Gradient bound.*) $\|\nabla f(\varphi(t))\| \leq G$.*

*(iii) (*Chart curvature bound.*) The bilinear operator norm satisfies*

$$\|D^2\varphi(t)\|_{\mathrm{bil}} := \sup_{\|a\|=\|b\|=1} \|D^2\varphi(t)[a,b]\| \;\leq\; B. \tag{247}$$

*(iv) (*Uniform conditioning of the chart.*) The smallest singular value of $J(t)$ satisfies*

$$\sigma_{\min}(J(t)) \geq \sigma_0. \tag{248}$$

If $GB \leq \frac{1}{2}\lambda_0\sigma_0^2$, then $\tilde{f}$ is λ-strongly concave on U_0 with

$$\lambda := \tfrac{1}{2}\lambda_0\sigma_0^2.$$

Consequently, f is strongly concave along S on the chart patch $\varphi(U_0)$ (in the sense that $\tilde{f} = f \circ \varphi$ is strongly concave in Euclidean coordinates on U_0).

Proof. **Step 1 (Compute the pullback Hessian, including the curvature term).** Since f and φ are C^2, $\tilde{f} = f \circ \varphi$ is C^2 and

$$\nabla \tilde{f}(t) = J(t)^\top \nabla f(\varphi(t)).$$

We now compute the quadratic form of $\nabla^2 \tilde{f}(t)$. Fix $a \in \mathbb{R}^m$ and define the curve $c(s) := \varphi(t + sa)$ for s near 0. Then $\tilde{f}(t + sa) = f(c(s))$. Differentiate twice at $s = 0$:

$$\begin{aligned} a^\top \nabla^2 \tilde{f}(t)\, a &= \frac{d^2}{ds^2}\Big|_{s=0} f(c(s)) \\ &= \underbrace{c'(0)^\top \nabla^2 f(c(0))\, c'(0)}_{\text{ambient Hessian term}} + \underbrace{\big\langle \nabla f(c(0)),\, c''(0) \big\rangle}_{\text{curvature term}}. \end{aligned} \tag{249}$$

Now $c(0) = \varphi(t)$, $c'(0) = J(t)a$, and $c''(0) = D^2\varphi(t)[a,a]$. Substituting into (249) gives the standard formula

$$a^\top \nabla^2 \tilde{f}(t)\, a = a^\top J(t)^\top \nabla^2 f(\varphi(t)) J(t)\, a + \big\langle \nabla f(\varphi(t)),\, D^2\varphi(t)[a,a] \big\rangle. \tag{250}$$

Step 2 (Lower bound the negative pullback Hessian). Multiply (250) by -1:

$$a^\top \big(-\nabla^2 \tilde{f}(t)\big) a = a^\top J(t)^\top \big(-\nabla^2 f(\varphi(t))\big) J(t)\, a - \big\langle \nabla f(\varphi(t)),\, D^2\varphi(t)[a,a] \big\rangle. \tag{251}$$

We lower bound the first term and upper bound the absolute value of the second.

(a) Projected curvature term. Assumption (246) gives

$$a^\top J(t)^\top \big(-\nabla^2 f(\varphi(t))\big) J(t)\, a \geq \lambda_0 \, \|J(t)a\|^2. \tag{252}$$

(b) Curvature term bound. By Cauchy–Schwarz and the bilinear norm definition (247),

$$\big|\big\langle \nabla f(\varphi(t)),\, D^2\varphi(t)[a,a]\big\rangle\big| \leq \|\nabla f(\varphi(t))\|\, \|D^2\varphi(t)[a,a]\| \leq G \cdot \|D^2\varphi(t)\|_{\mathrm{bil}}\, \|a\|^2 \leq GB\, \|a\|^2.$$

Therefore,

$$-\big\langle \nabla f(\varphi(t)),\, D^2\varphi(t)[a,a]\big\rangle \geq -GB\, \|a\|^2. \tag{253}$$

Combining (251), (252), and (253) yields

$$a^\top \big(-\nabla^2 \tilde f(t)\big) a \geq \lambda_0 \, \|J(t)a\|^2 - GB\, \|a\|^2. \tag{254}$$

Step 3 (Use chart conditioning to compare $\|J(t)a\|$ and $\|a\|$). By (248), $\|J(t)a\| \geq \sigma_0 \|a\|$ for all a. Substitute into (254):

$$a^\top \big(-\nabla^2 \tilde f(t)\big) a \geq (\lambda_0 \sigma_0^2 - GB)\, \|a\|^2.$$

Under the condition $GB \leq \frac{1}{2}\lambda_0\sigma_0^2$, the right-hand side is at least $\frac{1}{2}\lambda_0\sigma_0^2\|a\|^2$. Thus for all $t \in U_0$,

$$-\nabla^2 \tilde f(t) \succeq \lambda I_m, \qquad \lambda := \tfrac{1}{2}\lambda_0\sigma_0^2.$$

By Lemma D.3, $\tilde f$ is λ-strongly concave on U_0. This is exactly strong concavity of f along S on $\varphi(U_0)$ in the chosen chart. □

Remark D.9 (Geometric meaning of the curvature term). Formula (250) decomposes the second derivative of $f|_S$ in coordinates into: (i) the ambient Hessian of f projected onto tangent directions $J(t)a$, and (ii) a curvature correction involving the ambient gradient ∇f and the second derivative of the embedding φ (closely related to the second fundamental form of S). Affine slices are precisely the case $D^2\varphi \equiv 0$, hence no curvature term (Lemma D.7). Near a maximizer of $f|_S$, the tangential component of ∇f is small; on sufficiently small neighborhoods, the product $\|\nabla f\|\, \|D^2\varphi\|_{\mathrm{bil}}$ can be made subordinate to the projected negative curvature, yielding local strong concavity.

D.2. **From curvature to a well-defined and smooth EM map.** Assume (locally) that the M-step produces a unique maximizer

$$T(\theta) \in \arg\max_{\theta' \in \mathcal{K}} Q(\theta' \mid \theta),$$

where $\mathcal{K} \subseteq \Theta$ is a closed convex stability set (often compact). Strong concavity in θ' guarantees uniqueness once existence is secured (Remark D.5).

Assumption D.10 (C^2 smoothness and invertible Hessian at the maximizer)**.** There exists a neighborhood $\mathcal{N}$ of θ_0 such that:

(1) $Q(\theta' \mid \theta)$ is C^2 in (θ', θ) on a neighborhood of $(T(\theta_0), \theta_0)$;
(2) $\nabla^2_{\theta'\theta'} Q(T(\theta_0) \mid \theta_0)$ is invertible (equivalently, negative definite).

Lemma D.11 (Differentiability and Jacobian formula for an unconstrained EM map)**.** *Let $\Theta, \Theta' \subseteq \mathbb{R}^d$ be open sets and let*

$$Q : \Theta' \times \Theta \to \mathbb{R}, \qquad (\theta', \theta) \mapsto Q(\theta' \mid \theta)$$

be C^2 on a neighborhood of $(\theta_0', \theta_0) \in \Theta' \times \Theta$. Define

$$F(\theta', \theta) := \nabla_{\theta'} Q(\theta' \mid \theta) \in \mathbb{R}^d.$$

Assume:

(i) (First-order condition at the basepoint) $F(\theta_0', \theta_0) = 0$.
(ii) (Nonsingularity in θ') *the matrix*

$$D_{\theta'} F(\theta_0', \theta_0) = \nabla^2_{\theta'\theta'} Q(\theta_0' \mid \theta_0)$$

is invertible.

Then there exist neighborhoods $V \subseteq \Theta$ of θ_0 and $U \subseteq \Theta'$ of θ_0', and a unique C^1 map $T : V \to U$ such that

$$F(T(\theta), \theta) = 0 \qquad \forall \theta \in V, \tag{255}$$

and $T(\theta_0) = \theta_0'$. Moreover, T is differentiable at θ_0 with Jacobian

$$\nabla T(\theta_0) = -\Big[\nabla^2_{\theta'\theta'} Q(T(\theta_0) \mid \theta_0)\Big]^{-1} \nabla^2_{\theta'\theta} Q(T(\theta_0) \mid \theta_0). \tag{256}$$

Finally, if $A := \nabla T(\theta_0)$ satisfies $\rho(A) < 1$, then for every $\varepsilon \in \big(0, 1 - \rho(A)\big)$ there exist a norm $\|\cdot\|_\varepsilon$ on $\mathbb{R}^d$, a neighborhood $V_\varepsilon \subseteq V$ of θ_0, and a constant $c_\varepsilon \in (0,1)$ such that

$$\|T(\theta) - T(\vartheta)\|_\varepsilon \le c_\varepsilon \, \|\theta - \vartheta\|_\varepsilon \qquad \forall \theta, \vartheta \in V_\varepsilon. \tag{257}$$

Proof. **Step 1 (Set up the implicit equation).** Define $F(\theta', \theta) = \nabla_{\theta'} Q(\theta' \mid \theta)$. Since Q is C^2, F is C^1 on a neighborhood of (θ_0', θ_0). Assumption (i) states $F(\theta_0', \theta_0) = 0$.

Step 2 (Verify the IFT hypotheses and invoke the implicit function theorem). Compute the partial derivative of F with respect to θ':

$$D_{\theta'} F(\theta', \theta) = \nabla^2_{\theta'\theta'} Q(\theta' \mid \theta).$$

Assumption (ii) states that $D_{\theta'} F(\theta_0', \theta_0)$ is invertible. By the (classical) implicit function theorem in finite dimensions, there exist neighborhoods U of θ_0' and V of θ_0 and a unique C^1 map $T : V \to U$ such that $T(\theta_0) = \theta_0'$ and $F(T(\theta), \theta) = 0$ for all $\theta \in V$.

Step 3 (Differentiate the identity $F(T(\theta), \theta) = 0$ and solve). Differentiate (255) at $\theta = \theta_0$. Using the chain rule for $F : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d$,

$$0 = D\big[F(T(\theta), \theta)\big]_{\theta = \theta_0} = D_{\theta'} F(\theta_0', \theta_0)\, \nabla T(\theta_0) + D_\theta F(\theta_0', \theta_0).$$

Since $D_{\theta'} F(\theta_0', \theta_0)$ is invertible, we can solve:

$$\nabla T(\theta_0) = -\big(D_{\theta'} F(\theta_0', \theta_0)\big)^{-1} D_\theta F(\theta_0', \theta_0).$$

Finally observe that $D_\theta F(\theta', \theta) = \nabla^2_{\theta'\theta} Q(\theta' \mid \theta)$, and $\theta_0' = T(\theta_0)$, yielding (256).

Step 4 (A linear-algebra lemma: spectral radius yields an adapted norm). We use the following standard fact.

Claim D.12. *Let $A \in \mathbb{R}^{d \times d}$ and let $\varepsilon > 0$. There exists a norm $\|\cdot\|_\varepsilon$ on $\mathbb{R}^d$ with induced operator norm $\|\cdot\|_{\mathrm{op},\varepsilon}$ such that*

$$\|A\|_{\mathrm{op},\varepsilon} \le \rho(A) + \varepsilon. \tag{258}$$

Proof of Claim D.12. We give a constructive proof via Jordan form (any equivalent proof is acceptable).

Let $A = PJP^{-1}$ be a (real) Jordan decomposition over $\mathbb{C}$; operator norms below may be defined over $\mathbb{C}^d$ and then restricted to $\mathbb{R}^d$, which does not affect (258). Write J as a direct sum of Jordan blocks $J_k = \lambda_k I + N_k$, where N_k is nilpotent with ones on the superdiagonal. Let $\rho := \max_k |\lambda_k| = \rho(A)$.

Fix $\eta > 0$ (to be chosen in terms of ε). For each block size m, define the diagonal scaling matrix

$$D_\eta := \operatorname{diag}(1, \eta, \eta^2, \dots, \eta^{m-1}).$$

A direct computation shows

$$D_\eta^{-1} N_k D_\eta = \eta\, N_k,$$

hence

$$D_\eta^{-1} J_k D_\eta = \lambda_k I + \eta N_k.$$

Equip $\mathbb{C}^m$ with the ℓ_∞ norm, whose induced operator norm satisfies $\|\lambda_k I\|_{\mathrm{op},\infty} = |\lambda_k|$ and $\|N_k\|_{\mathrm{op},\infty} = 1$. Therefore

$$\left\|D_\eta^{-1} J_k D_\eta\right\|_{\mathrm{op},\infty} \le |\lambda_k| + \eta \le \rho + \eta.$$

As J is block diagonal, with the max norm on the product space we obtain

$$\|D^{-1} J D\|_{\mathrm{op},\infty} \le \rho + \eta,$$

where D is the block-diagonal matrix with the appropriate D_η on each Jordan block.

Now define a norm on $\mathbb{C}^d$ by

$$\|x\|_\varepsilon := \|D^{-1} P^{-1} x\|_\infty,$$

and let $\|\cdot\|_{\mathrm{op},\varepsilon}$ be its induced operator norm. Then

$$\|A\|_{\mathrm{op},\varepsilon} = \|D^{-1} P^{-1} A P D\|_{\mathrm{op},\infty} = \|D^{-1} J D\|_{\mathrm{op},\infty} \le \rho + \eta.$$

Choose $\eta = \varepsilon$ to obtain (258). □

Step 5 (Upgrade the linear bound to a nonlinear local contraction). Assume $\rho(A) < 1$, where $A = \nabla T(\theta_0)$. Fix $\varepsilon \in (0, 1 - \rho(A))$ and let $\|\cdot\|_\varepsilon$ be the norm from Claim D.12 with

$$\|A\|_{\mathrm{op},\varepsilon} \le \rho(A) + \varepsilon/2 < 1.$$

Since T is C^1 on V, the map $\theta \mapsto \nabla T(\theta)$ is continuous at θ_0. Hence there exists a neighborhood $V_\varepsilon \subseteq V$ of θ_0 such that

$$\|\nabla T(\theta) - A\|_{\mathrm{op},\varepsilon} \le \varepsilon/2 \qquad \forall \theta \in V_\varepsilon. \tag{259}$$

By the triangle inequality and (259),

$$\|\nabla T(\theta)\|_{\mathrm{op},\varepsilon} \le \|A\|_{\mathrm{op},\varepsilon} + \|\nabla T(\theta) - A\|_{\mathrm{op},\varepsilon} \le (\rho(A) + \varepsilon/2) + \varepsilon/2 = \rho(A) + \varepsilon =: c_\varepsilon,$$

and $c_\varepsilon \in (0, 1)$ by the choice of ε.

Now fix $\theta, \vartheta \in V_\varepsilon$. Because V_ε is open, after shrinking it if necessary we may assume it is convex (or simply replace it by a small ball), so the segment $\vartheta + t(\theta - \vartheta) \in V_\varepsilon$ for all $t \in [0, 1]$. By the fundamental theorem of calculus in $\mathbb{R}^d$ (integral form of the mean-value theorem),

$$T(\theta) - T(\vartheta) = \int_0^1 \nabla T(\vartheta + t(\theta - \vartheta))(\theta - \vartheta)\, dt.$$

Taking $\|\cdot\|_\varepsilon$ and using the uniform bound on $\|\nabla T(\cdot)\|_{\mathrm{op},\varepsilon}$ yields

$$\|T(\theta) - T(\vartheta)\|_\varepsilon \le \int_0^1 \|\nabla T(\vartheta + t(\theta - \vartheta))\|_{\mathrm{op},\varepsilon}\, dt\ \|\theta - \vartheta\|_\varepsilon \le c_\varepsilon \|\theta - \vartheta\|_\varepsilon.$$

This is (257). □

Lemma D.13 (Lipschitzness of the argmax map from curvature and mixed smoothness)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be nonempty, closed, and convex, and let $\mathcal{S} \subseteq \mathbb{R}^d$ be any index set. For each $\theta \in \mathcal{S}$, let $\theta' \mapsto Q(\theta' \mid \theta)$ be differentiable on $\mathcal{K}$ and assume:*

(i) (Uniform strong concavity in θ'.) *There exists $\lambda > 0$ such that for every $\theta \in \mathcal{S}$, the function $\theta' \mapsto Q(\theta' \mid \theta)$ is λ-strongly concave on $\mathcal{K}$ (equivalently, it satisfies (234) on $\mathcal{K}$ with the same λ).*

(ii) (Mixed gradient Lipschitzness.) *There exists $L < \infty$ such that for all $\theta, \vartheta \in \mathcal{S}$,*

$$\sup_{\theta' \in \mathcal{K}} \left\| \nabla_{\theta'} Q(\theta' \mid \theta) - \nabla_{\theta'} Q(\theta' \mid \vartheta) \right\| \le L \, \|\theta - \vartheta\|. \tag{260}$$

(iii) (Existence of maximizers.) *For each $\theta \in \mathcal{S}$, the maximizer*

$$T(\theta) \in \arg\max_{\theta' \in \mathcal{K}} Q(\theta' \mid \theta)$$

exists.

Then $T(\theta)$ is unique for each $\theta \in \mathcal{S}$, and for all $\theta, \vartheta \in \mathcal{S}$,

$$\|T(\theta) - T(\vartheta)\| \le \frac{L}{\lambda} \, \|\theta - \vartheta\|. \tag{261}$$

Proof. **Step 1 (Uniqueness).** Fix $\theta \in \mathcal{S}$. By (i), $Q(\cdot \mid \theta)$ is λ-strongly concave on the convex set $\mathcal{K}$. By Lemma D.4, any maximizer is unique. Thus T is well-defined as a single-valued map on $\mathcal{S}$.

Step 2 (Write the optimality conditions in variational form). Fix $\theta, \vartheta \in \mathcal{S}$. Since $\mathcal{K}$ is closed and convex and $Q(\cdot \mid \theta)$ is concave and differentiable, Lemma D.2 implies the first-order optimality conditions:

$$\left\langle \nabla_{\theta'} Q(T(\theta) \mid \theta), \, u - T(\theta) \right\rangle \le 0 \qquad \forall u \in \mathcal{K}, \tag{262}$$

$$\left\langle \nabla_{\theta'} Q(T(\vartheta) \mid \vartheta), \, u - T(\vartheta) \right\rangle \le 0 \qquad \forall u \in \mathcal{K}. \tag{263}$$

Choose $u = T(\vartheta)$ in (262) and $u = T(\theta)$ in (263) to obtain

$$\left\langle \nabla_{\theta'} Q(T(\theta) \mid \theta), \, T(\vartheta) - T(\theta) \right\rangle \le 0, \tag{264}$$

$$\left\langle \nabla_{\theta'} Q(T(\vartheta) \mid \vartheta), \, T(\theta) - T(\vartheta) \right\rangle \le 0. \tag{265}$$

Adding (264) and (265) yields

$$\left\langle \nabla_{\theta'} Q(T(\theta) \mid \theta) - \nabla_{\theta'} Q(T(\vartheta) \mid \vartheta), \, T(\theta) - T(\vartheta) \right\rangle \ge 0. \tag{266}$$

(Equivalently, the left-hand side is the negative of what appears in some conventions; the sign here is correct as written.)

Step 3 (Insert–subtract and isolate the strong concavity term). Decompose the gradient difference as

$$\nabla_{\theta'} Q(T(\theta) \mid \theta) - \nabla_{\theta'} Q(T(\vartheta) \mid \vartheta) = \underbrace{\left(\nabla_{\theta'} Q(T(\theta) \mid \theta) - \nabla_{\theta'} Q(T(\theta) \mid \vartheta) \right)}_{\text{(I)}} + \underbrace{\left(\nabla_{\theta'} Q(T(\theta) \mid \vartheta) - \nabla_{\theta'} Q(T(\vartheta) \mid \vartheta) \right)}_{\text{(II)}}.$$

Plug into (266):

$$\langle \text{(I)}, \, T(\theta) - T(\vartheta) \rangle + \langle \text{(II)}, \, T(\theta) - T(\vartheta) \rangle \ge 0,$$

hence

$$\langle \text{(II)}, \, T(\theta) - T(\vartheta) \rangle \ge -\langle \text{(I)}, \, T(\theta) - T(\vartheta) \rangle. \tag{267}$$

Step 4 (Lower bound the left side via strong concavity). Now view ϑ as fixed and consider the strongly concave function $u \mapsto Q(u \mid \vartheta)$ on $\mathcal{K}$. By Lemma D.3 in gradient form, strong concavity implies for all $u, v \in \mathcal{K}$,

$$\left\langle \nabla_{\theta'} Q(u \mid \vartheta) - \nabla_{\theta'} Q(v \mid \vartheta), \, u - v \right\rangle \le -\lambda \|u - v\|^2.$$

Apply this with $u = T(\theta)$ and $v = T(\vartheta)$ to obtain

$$\big\langle \nabla_{\theta'} Q(T(\theta) \mid \vartheta) - \nabla_{\theta'} Q(T(\vartheta) \mid \vartheta),\, T(\theta) - T(\vartheta) \big\rangle \le -\lambda \|T(\theta) - T(\vartheta)\|^2.$$

That is,

$$\langle (\mathrm{II}),\, T(\theta) - T(\vartheta) \rangle \le -\lambda \|T(\theta) - T(\vartheta)\|^2. \tag{268}$$

Step 5 (Upper bound the right side via mixed smoothness and conclude). By Cauchy–Schwarz,

$$-\langle (\mathrm{I}),\, T(\theta) - T(\vartheta) \rangle \le \|(\mathrm{I})\|\, \|T(\theta) - T(\vartheta)\|.$$

Combining with (267) and (268) gives

$$\lambda \|T(\theta) - T(\vartheta)\|^2 \le \|(\mathrm{I})\|\, \|T(\theta) - T(\vartheta)\|.$$

If $T(\theta) = T(\vartheta)$ we are done. Otherwise divide both sides by $\|T(\theta) - T(\vartheta)\|$:

$$\lambda \|T(\theta) - T(\vartheta)\| \le \|(\mathrm{I})\| = \big\| \nabla_{\theta'} Q(T(\theta) \mid \theta) - \nabla_{\theta'} Q(T(\theta) \mid \vartheta) \big\|.$$

Finally apply (260) (note $T(\theta) \in \mathcal{K}$) to obtain

$$\lambda \|T(\theta) - T(\vartheta)\| \le L \|\theta - \vartheta\|,$$

which is (261). □

Corollary D.14 (A simple Euclidean contraction criterion)**.** *In the setting of Lemma D.13, if $L/\lambda < 1$, then T is a contraction on $\mathcal{S}$ in Euclidean norm:*

$$\|T(\theta) - T(\vartheta)\| \le \frac{L}{\lambda} \|\theta - \vartheta\| \qquad \forall \theta, \vartheta \in \mathcal{S}.$$

Remark D.15 (Relating L to mixed Hessians via a rigorous mean-value argument). Assume Q is C^2 in (θ', θ) on $\mathcal{K} \times \mathcal{S}$ (with $\mathcal{S}$ convex for simplicity), and that

$$\sup_{\theta' \in \mathcal{K},\ \theta \in \mathcal{S}} \big\| \nabla^2_{\theta'\theta} Q(\theta' \mid \theta) \big\| \le M < \infty.$$

Then (260) holds with $L = M$.

Indeed, fix $\theta, \vartheta \in \mathcal{S}$ and $\theta' \in \mathcal{K}$. Define the path $\theta_s := \vartheta + s(\theta - \vartheta)$ for $s \in [0, 1]$ and the function $g(s) := \nabla_{\theta'} Q(\theta' \mid \theta_s)$. Then g is differentiable and, by the chain rule,

$$g'(s) = \nabla^2_{\theta'\theta} Q(\theta' \mid \theta_s)\, (\theta - \vartheta).$$

Integrating from 0 to 1 gives

$$\nabla_{\theta'} Q(\theta' \mid \theta) - \nabla_{\theta'} Q(\theta' \mid \vartheta) = \int_0^1 \nabla^2_{\theta'\theta} Q(\theta' \mid \theta_s)\, (\theta - \vartheta)\, ds,$$

so by the operator norm bound,

$$\big\| \nabla_{\theta'} Q(\theta' \mid \theta) - \nabla_{\theta'} Q(\theta' \mid \vartheta) \big\| \le \int_0^1 M\, \|\theta - \vartheta\|\, ds = M \|\theta - \vartheta\|.$$

Taking $\sup_{\theta' \in \mathcal{K}}$ yields (260) with $L = M$.

Remark D.16 (When the IFT Jacobian formula and the Lipschitz bound are used). The Jacobian identity (256) is *local* and exact, and it is the right tool for sharp local rates (spectral radius governs the rate in an adapted norm). The Euclidean Lipschitz bound (261) is *global on the set* $\mathcal{S}$ and is often simpler to verify from curvature (λ) and mixed smoothness (L), producing the classical contraction template $L/\lambda < 1$.

D.3. **Stability of maximizers under perturbations.** We now expand the argmax perturbation calculus used to compare population and sample-based M-steps.

Lemma D.17 (First-order *necessary* condition for a maximizer on a convex set)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be nonempty and convex, and let $f : \mathcal{K} \to \mathbb{R}$ be differentiable. If $u^\star \in \mathcal{K}$ satisfies*

$$f(u^\star) \;\geq\; f(u) \qquad \forall u \in \mathcal{K}$$

(i.e. $u^\star$ is a global maximizer of f over $\mathcal{K}$), then

$$\langle \nabla f(u^\star),\, u - u^\star \rangle \leq 0 \qquad \forall u \in \mathcal{K}. \tag{269}$$

Proof. Fix $u \in \mathcal{K}$ and consider the segment $u_t := u^\star + t(u - u^\star)$ for $t \in [0,1]$. Convexity of $\mathcal{K}$ implies $u_t \in \mathcal{K}$ for all $t \in [0,1]$. Define the one-dimensional function $\phi(t) := f(u_t)$. Since $u^\star$ is a global maximizer on $\mathcal{K}$, we have $\phi(t) \leq \phi(0)$ for all $t \in [0,1]$. Hence $t = 0$ is a maximizer of ϕ on $[0,1]$, so the right derivative satisfies $\phi'(0^+) \leq 0$.

Because f is differentiable at $u^\star$, ϕ is differentiable at 0 and by the chain rule,

$$\phi'(0) = \big\langle \nabla f(u^\star),\, u - u^\star \big\rangle.$$

Thus $\langle \nabla f(u^\star), u - u^\star \rangle = \phi'(0) \leq 0$. Since u was arbitrary, (269) holds. □

Lemma D.18 (First-order condition is also *sufficient* under concavity)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be nonempty, closed, and convex, and let $f : \mathcal{K} \to \mathbb{R}$ be differentiable and concave. Then $u^\star \in \mathcal{K}$ maximizes f over $\mathcal{K}$ if and only if*

$$\langle \nabla f(u^\star),\, u - u^\star \rangle \leq 0 \qquad \forall u \in \mathcal{K}. \tag{270}$$

Proof. Necessity is Lemma D.17 (no concavity needed).

For sufficiency, assume (230). Concavity of f implies the global upper-support inequality: for all $u \in \mathcal{K}$,

$$f(u) \;\leq\; f(u^\star) + \langle \nabla f(u^\star),\, u - u^\star \rangle.$$

Combine with (230) to get $f(u) \leq f(u^\star)$ for all $u \in \mathcal{K}$, so $u^\star$ is a maximizer. □

Lemma D.19 (Equivalent characterizations of strong concavity)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be convex and let $f : \mathcal{K} \to \mathbb{R}$ be differentiable. Fix $\lambda > 0$. The following are equivalent:*

*(1) (*Gradient form*) For all $u, v \in \mathcal{K}$,*

$$\langle \nabla f(u) - \nabla f(v),\, u - v \rangle \leq -\lambda \|u - v\|^2. \tag{271}$$

*(2) (*Quadratic upper support*) For all $u, v \in \mathcal{K}$,*

$$f(u) \leq f(v) + \langle \nabla f(v), u - v \rangle - \frac{\lambda}{2}\|u - v\|^2. \tag{272}$$

If additionally f is C^2 on an open set containing $\mathcal{K}$, these are also equivalent to:

*(3) (*Hessian form*) For all $w \in \mathcal{K}$,*

$$-\nabla^2 f(w) \succeq \lambda I_d. \tag{273}$$

Proof. **(3)⇒(2).** Assume $f \in C^2$ and (235). Fix $u, v \in \mathcal{K}$ and define $w_t := v + t(u - v)$ and $\psi(t) := f(w_t)$. Then $\psi'(t) = \langle \nabla f(w_t), u - v \rangle$ and $\psi''(t) = (u - v)^\top \nabla^2 f(w_t)(u - v)$. By (235), $\psi''(t) \leq -\lambda \|u - v\|^2$ for all $t \in [0,1]$. Integrate from 0 to t:

$$\psi'(t) - \psi'(0) \leq -\lambda t \|u - v\|^2.$$

Integrate again from 0 to 1:

$$\psi(1) - \psi(0) - \psi'(0) \leq -\lambda \|u - v\|^2 \int_0^1 t\, dt = -\frac{\lambda}{2}\|u - v\|^2.$$

Substitute $\psi(1) = f(u)$, $\psi(0) = f(v)$, and $\psi'(0) = \langle \nabla f(v), u - v\rangle$ to obtain (233).

(2)⇒(1). Apply (233) to (u, v) and (v, u):

$$f(u) \le f(v) + \langle \nabla f(v), u - v\rangle - \frac{\lambda}{2}\|u - v\|^2,$$

$$f(v) \le f(u) + \langle \nabla f(u), v - u\rangle - \frac{\lambda}{2}\|u - v\|^2.$$

Add and cancel $f(u) + f(v)$:

$$0 \le \langle \nabla f(v), u - v\rangle + \langle \nabla f(u), v - u\rangle - \lambda\|u - v\|^2 = -\langle \nabla f(u) - \nabla f(v), u - v\rangle - \lambda\|u - v\|^2.$$

Rearrange to obtain (234).

(1)⇒(2). Fix $u, v \in \mathcal{K}$, define $w_t := v + t(u - v)$ and $\psi(t) := f(w_t)$ as above. Then

$$\psi'(t) - \psi'(0) = \langle \nabla f(w_t) - \nabla f(v), u - v\rangle.$$

Apply (234) to the pair (w_t, v). Since $w_t - v = t(u - v)$, we have

$$\langle \nabla f(w_t) - \nabla f(v), w_t - v\rangle \le -\lambda\|w_t - v\|^2 \quad \Longrightarrow \quad \langle \nabla f(w_t) - \nabla f(v), u - v\rangle \le -\lambda t\|u - v\|^2.$$

Thus $\psi'(t) - \psi'(0) \le -\lambda t\|u - v\|^2$. Integrate from 0 to 1:

$$f(u) - f(v) - \langle \nabla f(v), u - v\rangle = \int_0^1 (\psi'(t) - \psi'(0))\, dt \le -\lambda\|u - v\|^2 \int_0^1 t\, dt = -\frac{\lambda}{2}\|u - v\|^2,$$

which is (233).

(2)⇒(3) under C^2. Assume $f \in C^2$ on a neighborhood of $\mathcal{K}$ and (233). Fix $w \in \mathcal{K}$ and $h \in \mathbb{R}^d$ such that $w \pm th \in \mathcal{K}$ for small t (possible by convexity and small t). Apply (233) with $(u, v) = (w + th, w)$:

$$f(w + th) \le f(w) + t\langle \nabla f(w), h\rangle - \frac{\lambda}{2}t^2\|h\|^2.$$

Apply it with $(u, v) = (w - th, w)$:

$$f(w - th) \le f(w) - t\langle \nabla f(w), h\rangle - \frac{\lambda}{2}t^2\|h\|^2.$$

Add and subtract $2f(w)$:

$$f(w + th) - 2f(w) + f(w - th) \le -\lambda t^2\|h\|^2.$$

Divide by t^2 and let $t \downarrow 0$. By the second-order Taylor expansion (or the definition of the Hessian),

$$\lim_{t\to 0} \frac{f(w + th) - 2f(w) + f(w - th)}{t^2} = h^\top \nabla^2 f(w) h.$$

Hence $h^\top \nabla^2 f(w) h \le -\lambda\|h\|^2$ for all h, i.e. (235). □

Lemma D.20 (Uniqueness of maximizer under strong concavity)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be convex and let $f : \mathcal{K} \to \mathbb{R}$ be λ-strongly concave for some $\lambda > 0$ (in any of the equivalent senses of Lemma D.3). If a maximizer $u^\star \in \arg\max_{\mathcal{K}} f$ exists, then it is unique.*

Proof. Let $u_1, u_2 \in \mathcal{K}$ be maximizers, so $f(u_1) = f(u_2) = \sup_{\mathcal{K}} f$. Apply the quadratic support inequality (233) (Lemma D.3) with $(u, v) = (u_2, u_1)$:

$$f(u_2) \le f(u_1) + \langle \nabla f(u_1), u_2 - u_1\rangle - \frac{\lambda}{2}\|u_2 - u_1\|^2.$$

Since u_1 is a maximizer on a convex set, Lemma D.17 gives $\langle \nabla f(u_1), u_2 - u_1 \rangle \le 0$. Therefore

$$f(u_2) \le f(u_1) - \frac{\lambda}{2}\|u_2 - u_1\|^2.$$

But $f(u_2) = f(u_1)$, so $\|u_2 - u_1\|^2 \le 0$, hence $u_2 = u_1$. □

Lemma D.21 (Quadratic drop from the maximizer)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be convex, and let $f : \mathcal{K} \to \mathbb{R}$ be differentiable and λ-strongly concave on $\mathcal{K}$. Let $u^\star = \arg\max_{\mathcal{K}} f$. Then for every $u \in \mathcal{K}$,*

$$f(u) \le f(u^\star) - \frac{\lambda}{2}\|u - u^\star\|^2. \tag{274}$$

Equivalently,

$$f(u^\star) - f(u) \ge \frac{\lambda}{2}\|u - u^\star\|^2.$$

Proof. By Lemma D.3, strong concavity implies the quadratic support bound: for all $u \in \mathcal{K}$,

$$f(u) \le f(u^\star) + \langle \nabla f(u^\star), u - u^\star \rangle - \frac{\lambda}{2}\|u - u^\star\|^2.$$

Since $u^\star$ is a maximizer and $\mathcal{K}$ is convex, Lemma D.17 yields $\langle \nabla f(u^\star), u - u^\star \rangle \le 0$. Drop this term to obtain (274). □

Lemma D.22 (Argmax stability under uniform gradient error)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be nonempty, closed, and convex. Let $f, \widehat{f} : \mathcal{K} \to \mathbb{R}$ be differentiable. Assume:*

*(i) (*Strong concavity of f in gradient form*) There exists $\lambda > 0$ such that for all $u, v \in \mathcal{K}$,*

$$\big\langle \nabla f(u) - \nabla f(v),\, u - v \big\rangle \le -\lambda \|u - v\|^2. \tag{275}$$

*(ii) (*Existence of maximizers*) There exist*

$$u^\star \in \arg\max_{u \in \mathcal{K}} f(u), \qquad \hat{u} \in \arg\max_{u \in \mathcal{K}} \widehat{f}(u).$$

*(iii) (*Uniform gradient perturbation*) There exists $\varepsilon \ge 0$ such that*

$$\sup_{u \in \mathcal{K}} \|\nabla \widehat{f}(u) - \nabla f(u)\| \le \varepsilon. \tag{276}$$

Then $u^\star$ is unique and

$$\|\hat{u} - u^\star\| \le \frac{\varepsilon}{\lambda}. \tag{277}$$

Proof. Uniqueness of $u^\star$ follows from Lemma D.4.

Fix the maximizers $u^\star$ and $\hat{u}$. We prove (277).

Step 1 (First-order necessary conditions at $u^\star$ and $\hat{u}$). Since $u^\star$ maximizes f on the convex set $\mathcal{K}$ and f is differentiable, Lemma D.17 gives

$$\langle \nabla f(u^\star),\, u - u^\star \rangle \le 0 \qquad \forall u \in \mathcal{K}. \tag{278}$$

In particular, taking $u = \hat{u}$,

$$\langle \nabla f(u^\star),\, \hat{u} - u^\star \rangle \le 0. \tag{279}$$

Likewise, since $\hat{u}$ maximizes $\widehat{f}$ on $\mathcal{K}$ and $\widehat{f}$ is differentiable,

$$\langle \nabla \widehat{f}(\hat{u}),\, u - \hat{u} \rangle \le 0 \qquad \forall u \in \mathcal{K}. \tag{280}$$

Taking $u = u^\star$ gives

$$\langle \nabla \widehat{f}(\hat{u}),\, u^\star - \hat{u} \rangle \le 0 \qquad \Longleftrightarrow \qquad \langle \nabla \widehat{f}(\hat{u}),\, \hat{u} - u^\star \rangle \ge 0. \tag{281}$$

Step 2 (Use strong concavity to control the displacement). Apply (275) with $(u, v) = (\hat{u}, u^\star)$:

$$\langle \nabla f(\hat{u}) - \nabla f(u^\star),\, \hat{u} - u^\star \rangle \le -\lambda \|\hat{u} - u^\star\|^2.$$

Rearrange:

$$\begin{aligned} \lambda \|\hat{u} - u^\star\|^2 &\le \langle \nabla f(u^\star) - \nabla f(\hat{u}),\, \hat{u} - u^\star \rangle \\ &= \langle \nabla f(u^\star), \hat{u} - u^\star \rangle - \langle \nabla f(\hat{u}), \hat{u} - u^\star \rangle. \end{aligned} \tag{282}$$

By (279), the first inner product is ≤ 0, hence

$$\lambda \|\hat{u} - u^\star\|^2 \le -\langle \nabla f(\hat{u}), \hat{u} - u^\star \rangle = \langle \nabla f(\hat{u}), u^\star - \hat{u} \rangle. \tag{283}$$

Step 3 (Insert $\nabla \widehat{f}(\hat{u})$ and exploit (281)). Write

$$\langle \nabla f(\hat{u}), u^\star - \hat{u} \rangle = \langle \nabla \widehat{f}(\hat{u}), u^\star - \hat{u} \rangle + \langle \nabla f(\hat{u}) - \nabla \widehat{f}(\hat{u}), u^\star - \hat{u} \rangle.$$

By (281), $\langle \nabla \widehat{f}(\hat{u}), u^\star - \hat{u} \rangle \le 0$, therefore

$$\langle \nabla f(\hat{u}), u^\star - \hat{u} \rangle \le \langle \nabla f(\hat{u}) - \nabla \widehat{f}(\hat{u}), u^\star - \hat{u} \rangle.$$

Combine with (283) to get

$$\lambda \|\hat{u} - u^\star\|^2 \le \langle \nabla f(\hat{u}) - \nabla \widehat{f}(\hat{u}), u^\star - \hat{u} \rangle. \tag{284}$$

Step 4 (Cauchy–Schwarz and the uniform gradient bound). By Cauchy–Schwarz,

$$\langle \nabla f(\hat{u}) - \nabla \widehat{f}(\hat{u}), u^\star - \hat{u} \rangle \le \|\nabla f(\hat{u}) - \nabla \widehat{f}(\hat{u})\|\, \|u^\star - \hat{u}\|.$$

By (276), $\|\nabla f(\hat{u}) - \nabla \widehat{f}(\hat{u})\| \le \varepsilon$. Thus

$$\lambda \|\hat{u} - u^\star\|^2 \le \varepsilon \|\hat{u} - u^\star\|.$$

If $\hat{u} = u^\star$ we are done; otherwise divide both sides by $\|\hat{u} - u^\star\|$ to obtain $\|\hat{u} - u^\star\| \le \varepsilon/\lambda$. □

Remark D.23 (Tightness of the constant ε/λ). The factor $1/\lambda$ is unimprovable in general. For instance, on $\mathcal{K} = \mathbb{R}$, take $f(u) = -\frac{\lambda}{2}(u - u^\star)^2$ and $\widehat{f}(u) = f(u) + \varepsilon u$. Then $\sup_u |\widehat{f}'(u) - f'(u)| = \varepsilon$, and the maximizer shifts by exactly ε/λ.

Lemma D.24 (Argmax stability from uniform function-value error)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be nonempty, closed, and convex. Let $f, \widehat{f} : \mathcal{K} \to \mathbb{R}$ be arbitrary functions (no differentiability assumed). Assume:*

(i) f is λ-strongly concave on $\mathcal{K}$ (e.g. in the sense of Lemma D.3 if f is differentiable, or via the midpoint/segment inequality definition).
(ii) Maximizers $u^\star \in \arg\max_{\mathcal{K}} f$ and $\hat{u} \in \arg\max_{\mathcal{K}} \widehat{f}$ exist.
(iii) $\sup_{u \in \mathcal{K}} |\widehat{f}(u) - f(u)| \le \delta$ for some $\delta \ge 0$.

Then

$$\|\hat{u} - u^\star\| \le \sqrt{\frac{4\delta}{\lambda}}.$$

Proof. **Step 1 (Relate $f(\hat{u})$ to $f(u^\star)$ using the uniform error).** Since $\hat{u}$ maximizes $\widehat{f}$,

$$\widehat{f}(\hat{u}) \ge \widehat{f}(u^\star).$$

By the uniform bound $|\widehat{f} - f| \le \delta$,

$$f(\hat{u}) - \delta \le \widehat{f}(\hat{u}) \ge \widehat{f}(u^\star) \ge f(u^\star) - \delta.$$

Hence

$$f(u^\star) - f(\hat{u}) \le 2\delta. \tag{285}$$

Step 2 (Strong concavity gives a quadratic gap lower bound). Because f is λ-strongly concave and $u^\star$ is a maximizer, the quadratic drop inequality holds:

$$f(u^\star) - f(\hat{u}) \ge \frac{\lambda}{2}\|\hat{u} - u^\star\|^2. \tag{286}$$

To prove (286) directly without differentiability assumptions: set $u = \hat{u}$ and $v = u^\star$ in the defining strong concavity inequality along the segment, namely for $t \in (0,1)$,

$$f\big(t\hat{u} + (1-t)u^\star\big) \;\ge\; tf(\hat{u}) + (1-t)f(u^\star) + \frac{\lambda}{2}t(1-t)\|\hat{u} - u^\star\|^2.$$

Since $u^\star$ is a maximizer, the left-hand side is at most $f(u^\star)$. Thus

$$f(u^\star) \ge tf(\hat{u}) + (1-t)f(u^\star) + \frac{\lambda}{2}t(1-t)\|\hat{u} - u^\star\|^2.$$

Rearrange and divide by $t > 0$:

$$f(u^\star) - f(\hat{u}) \ge \frac{\lambda}{2}(1-t)\|\hat{u} - u^\star\|^2.$$

Let $t \downarrow 0$ to obtain (286).

Step 3 (Combine upper and lower bounds). Combine (285) and (286):

$$\frac{\lambda}{2}\|\hat{u} - u^\star\|^2 \le 2\delta,$$

so $\|\hat{u} - u^\star\|^2 \le 4\delta/\lambda$, i.e. $\|\hat{u} - u^\star\| \le \sqrt{4\delta/\lambda}$. □

Lemma D.25 (Approximate stationarity implies proximity to the maximizer)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be nonempty, closed, and convex. Let $f : \mathcal{K} \to \mathbb{R}$ be differentiable and λ-strongly concave on $\mathcal{K}$, meaning it satisfies*

$$f(u) \;\le\; f(v) + \langle \nabla f(v), u - v\rangle - \frac{\lambda}{2}\|u - v\|^2 \qquad \forall u, v \in \mathcal{K}. \tag{287}$$

Assume a maximizer $u^\star \in \arg\max_{u\in\mathcal{K}} f(u)$ exists (hence is unique). Suppose $\hat{u} \in \mathcal{K}$ satisfies the approximate first-order condition

$$\sup_{u\in\mathcal{K}} \langle \nabla f(\hat{u}),\, u - \hat{u}\rangle \le \eta \tag{288}$$

for some $\eta \ge 0$. Then

$$\|\hat{u} - u^\star\| \;\le\; \sqrt{\frac{2\eta}{\lambda}}.$$

Proof. **Step 0 (Uniqueness of $u^\star$).** Since f is λ-strongly concave on the convex set $\mathcal{K}$, it has at most one maximizer. Indeed, if u_1, u_2 are maximizers, apply (287) with $(u, v) = (u_2, u_1)$:

$$f(u_2) \le f(u_1) + \langle \nabla f(u_1), u_2 - u_1\rangle - \frac{\lambda}{2}\|u_2 - u_1\|^2.$$

Because u_1 maximizes f over the convex set $\mathcal{K}$ and f is differentiable, the first-order *necessary* condition holds:

$$\langle \nabla f(u_1), u_2 - u_1\rangle \le 0.$$

Thus

$$f(u_2) \le f(u_1) - \frac{\lambda}{2}\|u_2 - u_1\|^2.$$

But $f(u_2) = f(u_1)$, so $\|u_2 - u_1\|^2 \le 0$, hence $u_2 = u_1$.

Step 1 (Quadratic support inequality centered at $\hat{u}$). Apply (287) with $v = \hat{u}$ and arbitrary $u \in \mathcal{K}$:

$$f(u) \;\le\; f(\hat{u}) + \langle \nabla f(\hat{u}), u - \hat{u}\rangle - \frac{\lambda}{2}\|u - \hat{u}\|^2. \tag{289}$$

This is valid for every $u \in \mathcal{K}$ because $\hat{u} \in \mathcal{K}$.

Step 2 (Plug in $u = u^\star$ and rearrange). Since $u^\star \in \mathcal{K}$, we may take $u = u^\star$ in (289):

$$f(u^\star) \;\le\; f(\hat{u}) + \langle \nabla f(\hat{u}), u^\star - \hat{u}\rangle - \frac{\lambda}{2}\|u^\star - \hat{u}\|^2.$$

Move terms to the other side:

$$\frac{\lambda}{2}\|u^\star - \hat{u}\|^2 \;\le\; f(\hat{u}) - f(u^\star) + \langle \nabla f(\hat{u}), u^\star - \hat{u}\rangle. \tag{290}$$

Step 3 (Control the two terms on the right). *(a) Control $f(\hat{u}) - f(u^\star)$.* Because $u^\star$ maximizes f over $\mathcal{K}$ and $\hat{u} \in \mathcal{K}$,

$$f(\hat{u}) \;\le\; f(u^\star),$$

hence

$$f(\hat{u}) - f(u^\star) \;\le\; 0. \tag{291}$$

(b) Control $\langle \nabla f(\hat{u}), u^\star - \hat{u}\rangle$. By the hypothesis (288),

$$\sup_{u \in \mathcal{K}} \langle \nabla f(\hat{u}), u - \hat{u}\rangle \le \eta.$$

Since $u^\star \in \mathcal{K}$, the supremum dominates the particular choice $u = u^\star$, giving

$$\langle \nabla f(\hat{u}), u^\star - \hat{u}\rangle \le \eta. \tag{292}$$

Step 4 (Conclude). Insert (291) and (292) into (290):

$$\frac{\lambda}{2}\|u^\star - \hat{u}\|^2 \le 0 + \eta = \eta.$$

Thus $\|u^\star - \hat{u}\|^2 \le 2\eta/\lambda$, i.e.

$$\|\hat{u} - u^\star\| \le \sqrt{\frac{2\eta}{\lambda}}.$$

□

Corollary D.26 (EM-style perturbation bound for the M-step)**.** *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be nonempty, closed, and convex. Fix θ and define*

$$f(\theta') := Q(\theta' \mid \theta), \qquad \widehat{f}(\theta') := \widehat{Q}(\theta' \mid \theta), \qquad \theta' \in \mathcal{K}.$$

Assume:

(i) $Q(\cdot \mid \theta)$ is λ-strongly concave on $\mathcal{K}$ for some $\lambda > 0$.
(ii) Maximizers exist:

$$T(\theta) \in \arg\max_{\theta' \in \mathcal{K}} Q(\theta' \mid \theta), \qquad \widehat{T}(\theta) \in \arg\max_{\theta' \in \mathcal{K}} \widehat{Q}(\theta' \mid \theta).$$

(iii) The gradients are uniformly close on $\mathcal{K}$:

$$\sup_{\theta' \in \mathcal{K}} \left\| \nabla_{\theta'} \widehat{Q}(\theta' \mid \theta) - \nabla_{\theta'} Q(\theta' \mid \theta) \right\| \le \varepsilon.$$

Then

$$\|\widehat{T}(\theta) - T(\theta)\| \le \frac{\varepsilon}{\lambda}.$$

Proof. Define f and $\widehat{f}$ as in the statement. Assumption (i) means f is λ-strongly concave on $\mathcal{K}$. Assumption (ii) provides maximizers

$$u^\star := T(\theta) \in \arg\max_{\mathcal{K}} f, \qquad \hat{u} := \widehat{T}(\theta) \in \arg\max_{\mathcal{K}} \widehat{f}.$$

Assumption (iii) is exactly the uniform gradient perturbation condition

$$\sup_{u \in \mathcal{K}} \|\nabla \widehat{f}(u) - \nabla f(u)\| \le \varepsilon.$$

Therefore Lemma D.22 applies, yielding

$$\|\widehat{T}(\theta) - T(\theta)\| = \|\hat{u} - u^\star\| \le \frac{\varepsilon}{\lambda}. \qquad \square$$

Appendix E. Empirical-process tools for operator deviation bounds

Let P be a probability measure on a measurable space $(\mathcal{X}, \mathcal{F})$ and let $X_1, \dots, X_n \overset{iid}{\sim} P$ with empirical measure $P_n := \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$. For a measurable map $f : \mathcal{X} \to \mathbb{R}^m$ (scalar $m = 1$, vector $m = d$, or matrix $m = d^2$ via vectorization) we write

$$(P_n - P)f \;:=\; \frac{1}{n} \sum_{i=1}^n f(X_i) - \mathbb{E}[f(X)].$$

We use $\|\cdot\|$ for the Euclidean norm on $\mathbb{R}^d$ and $\langle \cdot, \cdot \rangle$ for the associated inner product. For matrices $A \in \mathbb{R}^{d \times d}$, we write $\|A\|_{\mathrm{op}}$ for the spectral/operator norm induced by $\|\cdot\|$ and $\|A\|_{\mathcal{F}}$ for the Frobenius norm. We denote the unit sphere by $\S^{d-1} := \{u \in \mathbb{R}^d : \|u\| = 1\}$.
Function classes and measurability convention. For a scalar class $\mathcal{F} \subseteq \{f : \mathcal{X} \to \mathbb{R}\}$ define

$$Z(\mathcal{F}) \;:=\; \sup_{f \in \mathcal{F}} (P_n - P)f, \qquad Z_{\pm}(\mathcal{F}) \;:=\; \sup_{f \in \mathcal{F}} \big|(P_n - P)f\big|.$$

To avoid measurability pathologies of $\sup_{f \in \mathcal{F}}$ in full generality, we adopt the standard convention:

Assumption E.1 (Pointwise measurability / separability)**.** Every class $\mathcal{F}$ (and each scalarized class derived from it below) is *pointwise measurable*: there exists a countable subset $\mathcal{F}_0 \subseteq \mathcal{F}$ such that for every $f \in \mathcal{F}$ there is a sequence $f_k \in \mathcal{F}_0$ with $f_k(x) \to f(x)$ for all $x \in \mathcal{X}$.

Under Assumption E.1, the random variables $Z(\mathcal{F})$ and $Z_{\pm}(\mathcal{F})$ are measurable and one may freely apply expectation and concentration inequalities. (Otherwise, all statements below remain valid with the standard "outer expectation/probability" modifications.)

E.1. **Scalarization for vector and operator norms.** The basic mechanism behind operator deviation bounds is *scalarization*: reduce a vector- or matrix-valued supremum to finitely many scalar suprema by discretizing the unit spheres.

Lemma E.2 (Scalarization via sphere nets: vectors)**.** *Let $\eta \in (0,1)$ and let $\mathcal{N}_\eta \subseteq \S^{d-1}$ be an η-net in Euclidean norm, i.e. for every $u \in \S^{d-1}$ there exists $v \in \mathcal{N}_\eta$ with $\|u - v\| \le \eta$. Then for every (possibly random) $Y \in \mathbb{R}^d$,*

$$\|Y\| \;\le\; \frac{1}{1-\eta} \max_{v \in \mathcal{N}_\eta} \langle v, Y \rangle. \tag{293}$$

Consequently, for any class $\mathcal{H} \subseteq \{h : \mathcal{X} \to \mathbb{R}^d\}$ and any probability measures P_n, P for which $(P_n - P)h$ is defined,

$$\sup_{h \in \mathcal{H}} \|(P_n - P)h\| \;\le\; \frac{1}{1-\eta} \max_{v \in \mathcal{N}_\eta} \sup_{h \in \mathcal{H}} (P_n - P)\langle v, h \rangle. \tag{294}$$

Moreover, there exists an η-net $\mathcal{N}_\eta \subseteq \S^{d-1}$ with cardinality

$$|\mathcal{N}_\eta| \le (1 + 2/\eta)^d. \tag{295}$$

Proof. **Step 1 (Support-function representation of the Euclidean norm).** We claim that for every $Y \in \mathbb{R}^d$,

$$\|Y\| \;=\; \sup_{u \in \S^{d-1}} \langle u, Y \rangle. \tag{296}$$

Indeed, by Cauchy–Schwarz, $\langle u, Y \rangle \le \|u\| \, \|Y\| = \|Y\|$ for all $u \in \S^{d-1}$, so the supremum is $\le \|Y\|$. If $Y \neq 0$, take $u = Y/\|Y\| \in \S^{d-1}$ to attain $\langle u, Y \rangle = \|Y\|$, so the supremum is $\ge \|Y\|$. If $Y = 0$, both sides equal 0. This proves (296).

Step 2 (Discretize the sphere and derive the $(1-\eta)^{-1}$ factor). Fix $u \in \S^{d-1}$. By the η-net property, choose $v \in \mathcal{N}_\eta$ with $\|u - v\| \le \eta$. Then, for any $Y \in \mathbb{R}^d$,

$$\langle u, Y \rangle = \langle v, Y \rangle + \langle u - v, Y \rangle \le \langle v, Y \rangle + \|u - v\| \, \|Y\| \le \langle v, Y \rangle + \eta \|Y\|,$$

where we used Cauchy–Schwarz in the middle inequality. Since $v \in \mathcal{N}_\eta$, we have $\langle v, Y \rangle \le \max_{w \in \mathcal{N}_\eta} \langle w, Y \rangle$, hence

$$\langle u, Y \rangle \le \max_{w \in \mathcal{N}_\eta} \langle w, Y \rangle + \eta \|Y\|.$$

Now take $\sup_{u \in \S^{d-1}}$ of the left-hand side and use (296):

$$\|Y\| \le \max_{w \in \mathcal{N}_\eta} \langle w, Y \rangle + \eta \|Y\|.$$

Rearrange (noting $1 - \eta > 0$):

$$(1 - \eta)\|Y\| \le \max_{w \in \mathcal{N}_\eta} \langle w, Y \rangle,$$

which is exactly (293).

Step 3 (Apply to empirical-process vectors). Apply (293) with $Y = (P_n - P)h$ to obtain, for each fixed h,

$$\|(P_n - P)h\| \le \frac{1}{1-\eta} \max_{v \in \mathcal{N}_\eta} \langle v, (P_n - P)h \rangle.$$

Take $\sup_{h \in \mathcal{H}}$ of both sides:

$$\sup_{h \in \mathcal{H}} \|(P_n - P)h\| \le \frac{1}{1-\eta} \sup_{h \in \mathcal{H}} \max_{v \in \mathcal{N}_\eta} \langle v, (P_n - P)h \rangle.$$

Since $\mathcal{N}_\eta$ is finite, for any real-valued function $a(h,v)$ we have $\sup_h \max_v a(h,v) \le \max_v \sup_h a(h,v)$ (because for every h, $\max_v a(h,v) \le \max_v \sup_{h'} a(h',v)$). Thus

$$\sup_{h\in\mathcal{H}} \|(P_n - P)h\| \le \frac{1}{1-\eta} \max_{v\in\mathcal{N}_\eta} \sup_{h\in\mathcal{H}} \langle v, (P_n - P)h\rangle.$$

Finally, by linearity of $(P_n - P)$ and the definition of $\langle v, h\rangle$ as a scalar function,

$$\langle v, (P_n - P)h\rangle = (P_n - P)\langle v, h\rangle,$$

which yields (294).

Step 4 (Existence of an η-net and the volumetric cardinality bound). We show existence and (295) by a standard maximal-packing argument.

(a) Existence of a maximal η-separated subset and net property. Call a subset $S \subseteq \S^{d-1}$ *η-separated* if $\|x - y\| > \eta$ for all distinct $x, y \in S$. Because $\S^{d-1}$ is compact, there exists a finite maximal η-separated subset $\mathcal{N}_\eta$ (e.g. by a greedy construction: iteratively add a point outside the union of existing η-balls; compactness forces termination). We claim that any maximal η-separated subset is an η-net: if there existed $u \in \S^{d-1}$ with $\|u - v\| > \eta$ for all $v \in \mathcal{N}_\eta$, then $\mathcal{N}_\eta \cup \{u\}$ would still be η-separated, contradicting maximality. Hence the net property holds.

(b) Disjointness of small balls. For distinct $v, w \in \mathcal{N}_\eta$, the η-separation gives $\|v - w\| > \eta$, so the Euclidean balls $B(v, \eta/2)$ and $B(w, \eta/2)$ are disjoint (otherwise a point in the intersection would yield $\|v - w\| \le \eta$ by the triangle inequality).

(c) Containment in a slightly larger ball. For any $x \in B(v, \eta/2)$ with $v \in \mathcal{N}_\eta$, we have $\|x\| \le \|v\| + \|x - v\| \le 1 + \eta/2$. Thus

$$\bigcup_{v\in\mathcal{N}_\eta} B(v, \eta/2) \;\subseteq\; B(0, 1 + \eta/2).$$

(d) Volume comparison. Let $\mathrm{Vol}(\cdot)$ denote Lebesgue volume in $\mathbb{R}^d$. By disjointness and containment,

$$|\mathcal{N}_\eta|\, \mathrm{Vol}\big(B(0,\eta/2)\big) = \sum_{v\in\mathcal{N}_\eta} \mathrm{Vol}\big(B(v,\eta/2)\big) = \mathrm{Vol}\left(\bigcup_{v\in\mathcal{N}_\eta} B(v,\eta/2)\right) \le \mathrm{Vol}\big(B(0,1+\eta/2)\big).$$

Using $\mathrm{Vol}(B(0,r)) = r^d \,\mathrm{Vol}(B(0,1))$,

$$|\mathcal{N}_\eta| \left(\frac{\eta}{2}\right)^d \mathrm{Vol}(B(0,1)) \le (1+\eta/2)^d \,\mathrm{Vol}(B(0,1)),$$

hence

$$|\mathcal{N}_\eta| \le \left(\frac{1+\eta/2}{\eta/2}\right)^d = (1 + 2/\eta)^d,$$

which is (295). □

Lemma E.3 (Scalarization via sphere nets: operator norm)**.** *Let $\eta \in (0,1)$ and let $\mathcal{N}_\eta \subseteq \S^{d-1}$ be an η-net. Then for any matrix $A \in \mathbb{R}^{d\times d}$,*

$$\|A\|_{\mathrm{op}} \;\le\; \frac{1}{(1-\eta)^2} \max_{u\in\mathcal{N}_\eta} \max_{v\in\mathcal{N}_\eta} \langle u, Av\rangle. \tag{297}$$

If A is symmetric, then for any $\eta \in (0, 1/2)$ one may take a single net and obtain

$$\|A\|_{\mathrm{op}} \;\le\; \frac{1}{1-2\eta} \max_{v\in\mathcal{N}_\eta} \big|\langle v, Av\rangle\big|. \tag{298}$$

If additionally $A \succeq 0$, the absolute value in (298) may be dropped.

Proof. **Step 1 (Two equivalent variational characterizations of $\|A\|_{\mathrm{op}}$).** We use:

$$\|A\|_{\mathrm{op}} = \sup_{v \in \S^{d-1}} \|Av\|, \tag{299}$$

$$\|Av\| = \sup_{u \in \S^{d-1}} \langle u, Av \rangle \qquad \text{for each fixed } v \in \S^{d-1}. \tag{300}$$

Identity (299) is the definition of the induced operator norm. Identity (300) is Lemma E.2 Step 1 (support-function representation) applied to the vector Av: indeed, by Cauchy–Schwarz, $\langle u, Av \rangle \le \|Av\|$, with equality at $u = Av/\|Av\|$ if $Av \neq 0$.

Step 2 (General case: discretize both left and right singular directions). Fix $v \in \S^{d-1}$ and apply Lemma E.2 (293) to the vector $Y = Av$:

$$\|Av\| \le \frac{1}{1-\eta} \max_{u \in \mathcal{N}_\eta} \langle u, Av \rangle. \tag{301}$$

Now take $\sup_{v \in \S^{d-1}}$ of both sides and use (299):

$$\|A\|_{\mathrm{op}} = \sup_{v \in \S^{d-1}} \|Av\| \le \frac{1}{1-\eta} \sup_{v \in \S^{d-1}} \max_{u \in \mathcal{N}_\eta} \langle u, Av \rangle.$$

As in Lemma E.2 Step 3, since $\mathcal{N}_\eta$ is finite we may exchange sup and max:

$$\sup_{v \in \S^{d-1}} \max_{u \in \mathcal{N}_\eta} \langle u, Av \rangle \le \max_{u \in \mathcal{N}_\eta} \sup_{v \in \S^{d-1}} \langle u, Av \rangle.$$

For each fixed u, $\langle u, Av \rangle = \langle A^\top u, v \rangle$, hence by (300) applied to $A^\top u$,

$$\sup_{v \in \S^{d-1}} \langle u, Av \rangle = \sup_{v \in \S^{d-1}} \langle A^\top u, v \rangle = \|A^\top u\|.$$

Therefore,

$$\|A\|_{\mathrm{op}} \le \frac{1}{1-\eta} \max_{u \in \mathcal{N}_\eta} \|A^\top u\|.$$

Apply Lemma E.2 again to each vector $Y = A^\top u$:

$$\|A^\top u\| \le \frac{1}{1-\eta} \max_{v \in \mathcal{N}_\eta} \langle v, A^\top u \rangle = \frac{1}{1-\eta} \max_{v \in \mathcal{N}_\eta} \langle u, Av \rangle.$$

Combining the last two displays yields (297).

Step 3 (Symmetric case: reduction to a single net via Rayleigh quotients). Assume $A = A^\top$ and fix $\eta \in (0, 1/2)$.

(a) Rayleigh-quotient characterization. We claim

$$\|A\|_{\mathrm{op}} = \sup_{x \in \S^{d-1}} |\langle x, Ax \rangle|. \tag{302}$$

Indeed, since A is symmetric it admits an orthonormal eigenbasis with real eigenvalues $\{\lambda_i\}$. Writing $x = \sum_i \alpha_i e_i$ with $\sum_i \alpha_i^2 = 1$,

$$\langle x, Ax \rangle = \sum_i \lambda_i \alpha_i^2,$$

so the supremum of $\langle x, Ax \rangle$ over $\|x\| = 1$ equals $\lambda_{\max}(A)$, attained at a top eigenvector, and the infimum equals $\lambda_{\min}(A)$. Therefore $\sup_{\|x\|=1} |\langle x, Ax \rangle| = \max\{\lambda_{\max}(A), -\lambda_{\min}(A)\} = \|A\|_{\mathrm{op}}$, proving (302).

(b) Net approximation and rearrangement. Let $x^\star \in \S^{d-1}$ achieve the supremum in (302) (existence holds since $\S^{d-1}$ is compact and $x \mapsto |\langle x, Ax\rangle|$ is continuous). Choose $v \in \mathcal{N}_\eta$ with $\|x^\star - v\| \le \eta$. Then

$$\begin{aligned}|\langle x^\star, Ax^\star\rangle| &\le |\langle v, Av\rangle| + |\langle x^\star, Ax^\star\rangle - \langle v, Av\rangle| \\ &= |\langle v, Av\rangle| + |\langle x^\star - v, Ax^\star\rangle + \langle v, A(x^\star - v)\rangle| \\ &\le |\langle v, Av\rangle| + |\langle x^\star - v, Ax^\star\rangle| + |\langle v, A(x^\star - v)\rangle|.\end{aligned}$$

Bound each error term by Cauchy–Schwarz and the operator norm:

$$|\langle x^\star - v, Ax^\star\rangle| \le \|x^\star - v\| \cdot \|Ax^\star\| \le \eta \, \|A\|_{\mathrm{op}},$$

and similarly (using $\|v\| = 1$),

$$|\langle v, A(x^\star - v)\rangle| \le \|v\| \cdot \|A(x^\star - v)\| \le \|A\|_{\mathrm{op}} \, \|x^\star - v\| \le \eta \, \|A\|_{\mathrm{op}}.$$

Thus

$$\|A\|_{\mathrm{op}} = |\langle x^\star, Ax^\star\rangle| \le |\langle v, Av\rangle| + 2\eta \, \|A\|_{\mathrm{op}}.$$

Rearrange using $1 - 2\eta > 0$:

$$(1 - 2\eta)\|A\|_{\mathrm{op}} \le |\langle v, Av\rangle| \le \max_{w \in \mathcal{N}_\eta} |\langle w, Aw\rangle|,$$

which yields (298).

(c) Positive semidefinite case. If $A \succeq 0$, then $\langle v, Av\rangle \ge 0$ for all v, so $|\langle v, Av\rangle| = \langle v, Av\rangle$ and the absolute value may be dropped. □

Corollary E.4 (Matrix-valued classes: reduction to scalar classes)**.** *Let $\mathcal{M} \subseteq \{M : \mathcal{X} \to \mathbb{R}^{d\times d}\}$ and let $\mathcal{N}_\eta \subseteq \S^{d-1}$ be an η-net with $\eta \in (0, 1)$. Then*

$$\sup_{M \in \mathcal{M}} \big\|(P_n - P)M\big\|_{\mathrm{op}} \;\le\; \frac{1}{(1-\eta)^2} \max_{u \in \mathcal{N}_\eta} \max_{v \in \mathcal{N}_\eta} \sup_{M \in \mathcal{M}} (P_n - P)\, \langle u, Mv\rangle.$$

If each $M(x)$ is symmetric, then for any $\eta \in (0, 1/2)$ one also has the single-net bound

$$\sup_{M \in \mathcal{M}} \big\|(P_n - P)M\big\|_{\mathrm{op}} \;\le\; \frac{1}{1-2\eta} \max_{v \in \mathcal{N}_\eta} \sup_{M \in \mathcal{M}} (P_n - P)\, \big|\langle v, Mv\rangle\big|.$$

Proof. Apply Lemma E.3 to the deterministic matrix $A = (P_n - P)M$ for each fixed $M \in \mathcal{M}$:

$$\|(P_n - P)M\|_{\mathrm{op}} \le \frac{1}{(1-\eta)^2} \max_{u \in \mathcal{N}_\eta} \max_{v \in \mathcal{N}_\eta} \langle u, (P_n - P)Mv\rangle.$$

Take $\sup_{M \in \mathcal{M}}$ of both sides. Since $(P_n - P)$ is linear,

$$\langle u, (P_n - P)Mv\rangle = (P_n - P)\langle u, Mv\rangle,$$

and the first bound follows. The symmetric-case bound follows similarly from (298). □

Remark E.5 (Why nets are useful for operator deviation bounds). Corollary E.4 reduces operator norms to finitely many scalar empirical processes indexed by the classes $\{\langle u, Mv\rangle : M \in \mathcal{M}\}$ (or $\{\langle v, Mv\rangle\}$ in the symmetric case). One then applies scalar tools (symmetrization, contraction, entropy bounds, concentration) and finally union-bounds over $|\mathcal{N}_\eta|^2 \le (1 + 2/\eta)^{2d}$ directions (or $|\mathcal{N}_\eta|$ in the symmetric case).

E.2. **Symmetrization, contraction, and Rademacher complexity.** The workhorse bounds on $\mathbb{E}[Z(\mathcal{F})]$ proceed by symmetrization to Rademacher processes, followed by a complexity estimate.

Lemma E.6 (Symmetrization)**.** *Let $(X_i)_{i=1}^n$ be i.i.d. with law P on $(\mathcal{X}, \mathcal{A})$ and let $P_n := \frac{1}{n}\sum_{i=1}^n \delta_{X_i}$. Let $\mathcal{F} \subseteq \{f : \mathcal{X} \to \mathbb{R}\}$ satisfy Assumption E.1 so that the random variable $\sup_{f\in\mathcal{F}}(P_n - P)f$ is measurable.*[1] *Let $\epsilon_1, \dots, \epsilon_n$ be i.i.d. Rademacher signs independent of $(X_i)_{i=1}^n$. Then*

$$\mathbb{E}\Big[\sup_{f\in\mathcal{F}}(P_n - P)f\Big] \;\le\; 2\,\mathbb{E}\Big[\sup_{f\in\mathcal{F}}\frac{1}{n}\sum_{i=1}^n \epsilon_i f(X_i)\Big]. \tag{303}$$

Moreover, letting $\mathcal{F}^{\pm} := \mathcal{F} \cup (-\mathcal{F})$,

$$\mathbb{E}\Big[\sup_{f\in\mathcal{F}}|(P_n - P)f|\Big] = \mathbb{E}\Big[\sup_{g\in\mathcal{F}^{\pm}}(P_n - P)g\Big] \;\le\; 2\,\mathbb{E}\Big[\sup_{g\in\mathcal{F}^{\pm}}\frac{1}{n}\sum_{i=1}^n \epsilon_i g(X_i)\Big]. \tag{304}$$

Proof. **Step 0 (Ghost sample and notation).** Let $X_1', \dots, X_n' \overset{iid}{\sim} P$ be an independent copy ("ghost sample"), independent of X_1^n. Write $P_n' := \frac{1}{n}\sum_{i=1}^n \delta_{X_i'}$.

Step 1 (Replace P by the conditional expectation of P_n'). Fix $f \in \mathcal{F}$. Since $X_1'^n$ is i.i.d. with law P and independent of X_1^n,

$$\mathbb{E}\big[P_n' f \,\big|\, X_1^n\big] = \mathbb{E}[P_n' f] = Pf.$$

Hence, for each fixed f,

$$(P_n - P)f = P_n f - \mathbb{E}\big[P_n' f \,\big|\, X_1^n\big] = \mathbb{E}\big[(P_n - P_n')f \,\big|\, X_1^n\big].$$

Taking $\sup_{f\in\mathcal{F}}$ and using the elementary inequality $\sup_f \mathbb{E}[Z_f \mid X] \le \mathbb{E}[\sup_f Z_f \mid X]$ (a direct consequence of Jensen since sup is convex, or simply because $\sup_f Z_f \ge Z_{f_0}$ for each f_0 and then take conditional expectations), we obtain

$$\sup_{f\in\mathcal{F}}(P_n - P)f = \sup_{f\in\mathcal{F}}\mathbb{E}\big[(P_n - P_n')f \,\big|\, X_1^n\big] \le \mathbb{E}\left[\sup_{f\in\mathcal{F}}(P_n - P_n')f \,\middle|\, X_1^n\right].$$

Taking expectations yields

$$\mathbb{E}\Big[\sup_{f\in\mathcal{F}}(P_n - P)f\Big] \le \mathbb{E}\Big[\sup_{f\in\mathcal{F}}(P_n - P_n')f\Big]. \tag{305}$$

Step 2 (Random sign-swap identity). Let $\epsilon_1, \dots, \epsilon_n$ be i.i.d. Rademachers independent of $(X_i, X_i')_{i=1}^n$. Define the swapped pairs

$$(\widetilde{X}_i, \widetilde{X}_i') := \begin{cases} (X_i, X_i') & \text{if } \epsilon_i = +1, \\ (X_i', X_i) & \text{if } \epsilon_i = -1. \end{cases}$$

[1] If one does not wish to impose measurability, the same proof goes through with outer expectations throughout.

Because (X_i, X_i') is exchangeable and ϵ_i is symmetric, the joint law of $(\widetilde{X}_i, \widetilde{X}_i')_{i=1}^n$ is the same as that of $(X_i, X_i')_{i=1}^n$. Therefore,

$$\begin{aligned}\mathbb{E}\Big[\sup_{f\in\mathcal{F}}(P_n - P_n')f\Big] &= \mathbb{E}\Big[\sup_{f\in\mathcal{F}}\frac{1}{n}\sum_{i=1}^n\big(f(X_i) - f(X_i')\big)\Big] \\ &= \mathbb{E}\Big[\sup_{f\in\mathcal{F}}\frac{1}{n}\sum_{i=1}^n\big(f(\widetilde{X}_i) - f(\widetilde{X}_i')\big)\Big]. \end{aligned} \tag{306}$$

But by the definition of $(\widetilde{X}_i, \widetilde{X}_i')$,

$$f(\widetilde{X}_i) - f(\widetilde{X}_i') = \begin{cases} f(X_i) - f(X_i') & \epsilon_i = +1, \\ f(X_i') - f(X_i) & \epsilon_i = -1, \end{cases} = \epsilon_i\big(f(X_i) - f(X_i')\big).$$

Plugging this into (306) gives the exact identity

$$\mathbb{E}\Big[\sup_{f\in\mathcal{F}}(P_n - P_n')f\Big] = \mathbb{E}\Big[\sup_{f\in\mathcal{F}}\frac{1}{n}\sum_{i=1}^n\epsilon_i\big(f(X_i) - f(X_i')\big)\Big]. \tag{307}$$

Step 3 (Split and use symmetry). Using $\sup(A+B) \le \sup A + \sup B$,

$$\sup_{f\in\mathcal{F}}\frac{1}{n}\sum_{i=1}^n\epsilon_i\big(f(X_i) - f(X_i')\big) \le \sup_{f\in\mathcal{F}}\frac{1}{n}\sum_{i=1}^n\epsilon_i f(X_i) + \sup_{f\in\mathcal{F}}\frac{1}{n}\sum_{i=1}^n(-\epsilon_i)f(X_i').$$

Take expectations. Since $(X_i')_{i=1}^n$ is an independent copy of $(X_i)_{i=1}^n$ and $(-\epsilon_i)$ has the same law as (ϵ_i), the two terms have the same expectation. Hence from (307),

$$\mathbb{E}\Big[\sup_{f\in\mathcal{F}}(P_n - P_n')f\Big] \le 2\,\mathbb{E}\Big[\sup_{f\in\mathcal{F}}\frac{1}{n}\sum_{i=1}^n\epsilon_i f(X_i)\Big].$$

Combine with (305) to obtain (303).

Step 4 (Absolute-value form). For any real collection $(a_f)_{f\in\mathcal{F}}$,

$$\sup_{f\in\mathcal{F}}|a_f| = \max\Big\{\sup_{f\in\mathcal{F}}a_f,\ \sup_{f\in\mathcal{F}}(-a_f)\Big\} = \sup_{g\in\mathcal{F}\cup(-\mathcal{F})}a_g = \sup_{g\in\mathcal{F}^{\pm}}a_g.$$

Apply this with $a_f = (P_n - P)f$ to get the equality in (304), and then apply (303) to the class $\mathcal{F}^{\pm}$ to get the inequality. □

Rademacher complexity. Define the expected empirical Rademacher complexity

$$\mathfrak{R}_n(\mathcal{F}) \ :=\ \mathbb{E}\Big[\sup_{f\in\mathcal{F}}\frac{1}{n}\sum_{i=1}^n\epsilon_i f(X_i)\Big].$$

Lemma E.6 implies $\mathbb{E}[\sup_f(P_n - P)f] \le 2\mathfrak{R}_n(\mathcal{F})$ (and similarly for $\mathcal{F}^{\pm}$).

Lemma E.7 (Contraction inequality (scalar Lipschitz maps; correct constants))**.** *Let $\phi_i : \mathbb{R}\to\mathbb{R}$ be L-Lipschitz and satisfy $\phi_i(0) = 0$ for $i = 1,\dots,n$. Let $\mathcal{F} \subseteq \{f : \mathcal{X}\to\mathbb{R}\}$ satisfy Assumption E.1. Then*

$$\mathbb{E}\Big[\sup_{f\in\mathcal{F}}\frac{1}{n}\sum_{i=1}^n\epsilon_i\,\phi_i\big(f(X_i)\big)\Big] \ \le\ 2L\,\mathbb{E}\Big[\sup_{f\in\mathcal{F}}\frac{1}{n}\sum_{i=1}^n\epsilon_i\,f(X_i)\Big]. \tag{308}$$

If additionally $\mathcal{F}$ is symmetric *(i.e. $f \in \mathcal{F} \Rightarrow -f \in \mathcal{F}$), then the factor 2 may be removed:*

$$\mathbb{E}\Big[\sup_{f\in\mathcal{F}} \frac{1}{n}\sum_{i=1}^{n} \epsilon_i\, \phi_i\big(f(X_i)\big)\Big] \;\le\; L\,\mathbb{E}\Big[\sup_{f\in\mathcal{F}} \frac{1}{n}\sum_{i=1}^{n} \epsilon_i\, f(X_i)\Big].$$

Proof. This is the classical Ledoux–Talagrand contraction principle. A fully detailed proof (including the measurability reduction under Assumption E.1) appears, for example, as Theorem 4.12 in Ledoux–Talagrand (Probability in Banach Spaces, 1991) and in many modern empirical-process texts. We record the key reductions to make explicit what is used.

Step 1 (Reduction to a finite set of coordinate vectors). Condition on X_1^n. Consider the (random) subset of $\mathbb{R}^n$

$$T := \big\{(f(X_1),\dots,f(X_n)) :\; f \in \mathcal{F}\big\} \subseteq \mathbb{R}^n.$$

By separability/pointwise-measurability (Assumption E.1), there exists a countable $\mathcal{F}_0 \subseteq \mathcal{F}$ such that the suprema over $\mathcal{F}$ equal those over $\mathcal{F}_0$ almost surely; then approximate $\mathcal{F}_0$ by finite truncations and use monotone convergence to reduce to a finite T (this is standard and ensures measurability at each step).

Step 2 (Apply the contraction principle on $\mathbb{R}^n$). For a finite $T \subseteq \mathbb{R}^n$, define

$$Z_\phi(T) := \mathbb{E}_\epsilon\Big[\sup_{t\in T}\sum_{i=1}^{n} \epsilon_i\phi_i(t_i)\Big], \qquad Z_{\mathrm{id}}(T) := \mathbb{E}_\epsilon\Big[\sup_{t\in T}\sum_{i=1}^{n} \epsilon_i t_i\Big].$$

The Ledoux–Talagrand contraction principle states that if each ϕ_i is L-Lipschitz and $\phi_i(0) = 0$, then

$$Z_\phi(T) \;\le\; 2L\, Z_{\mathrm{id}}(T),$$

and if T is symmetric (i.e. $T = -T$, equivalent to symmetry of $\mathcal{F}$), then $Z_\phi(T) \le L\, Z_{\mathrm{id}}(T)$. Divide by n and uncondition to obtain (308) (and its symmetric refinement).

Step 3 (Why $\phi_i(0) = 0$ matters). If $\phi_i(0) \neq 0$, then $\sum_i \epsilon_i\phi_i(f(X_i))$ contains an uncontrolled term $\sum_i \epsilon_i\phi_i(0)$ independent of f, which has typical size $\asymp \sqrt{n}$ and is not bounded by the RHS. Centering ϕ_i at 0 removes this additive constant. □

Remark E.8 (Vector and matrix contraction). In many EM applications one first scalarizes vector/matrix deviations using sphere nets (Appendix E.1) and then applies Lemma E.7 to the resulting scalar classes. This typically suffices. There are also genuinely vector-valued contraction inequalities (with Lipschitz maps $\Phi_i : \mathbb{R}^d \to \mathbb{R}$), but scalarization usually avoids needing them.

E.3. **Entropy integrals and chaining bounds.**

Setup. Let $(\mathcal{F}, d)$ be a pseudo-metric space. For $\varepsilon > 0$, let $N(\mathcal{F}, d, \varepsilon)$ denote the covering number: the minimum cardinality of an ε-net in $(\mathcal{F}, d)$, and write $\mathrm{diam}(\mathcal{F}) := \sup_{f,g} d(f, g)$.

Lemma E.9 (Finite sub-Gaussian maximum)**.** *Let $Z_1,\dots,Z_m$ be centered random variables such that for some $\sigma > 0$,*

$$\mathbb{E}\big[e^{\lambda Z_j}\big] \;\le\; \exp\Big(\frac{\lambda^2\sigma^2}{2}\Big) \qquad \forall\lambda \in \mathbb{R},\ \forall j \le m.$$

Then

$$\mathbb{E}\Big[\max_{1\le j\le m} Z_j\Big] \;\le\; \sigma\sqrt{2\log m}.$$

Proof. For any $\lambda > 0$, by Jensen and $\max_j a_j \le \log \sum_j e^{a_j}$,

$$\mathbb{E}\Big[\max_j Z_j\Big] = \frac{1}{\lambda}\mathbb{E}\Big[\log e^{\lambda \max_j Z_j}\Big] \le \frac{1}{\lambda}\log \mathbb{E}\Big[e^{\lambda \max_j Z_j}\Big] \le \frac{1}{\lambda}\log \mathbb{E}\Big[\sum_{j=1}^{m} e^{\lambda Z_j}\Big] \le \frac{1}{\lambda}\log\Big(m e^{\lambda^2\sigma^2/2}\Big) = \frac{\log m}{\lambda} + \frac{\lambda\sigma^2}{2}$$

Optimize over $\lambda > 0$ at $\lambda = \sqrt{2\log m}/\sigma$. □

Lemma E.10 (Dudley entropy integral for Gaussian processes)**.** *Let $\{G_f\}_{f\in\mathcal{F}}$ be a centered Gaussian process such that*

$$\mathbb{E}\big[(G_f - G_g)^2\big] = d(f,g)^2 \qquad \forall f, g \in \mathcal{F}.$$

Then there exists a universal constant $C > 0$ such that

$$\mathbb{E}\Big[\sup_{f\in\mathcal{F}} G_f\Big] \;\le\; C\int_0^{\operatorname{diam}(\mathcal{F})} \sqrt{\log N(\mathcal{F}, d, \varepsilon)}\, d\varepsilon. \tag{309}$$

Proof. **Step 0 (Basepoint and compactness reduction).** Fix $f_0 \in \mathcal{F}$ and note

$$\sup_{f\in\mathcal{F}} G_f = G_{f_0} + \sup_{f\in\mathcal{F}} (G_f - G_{f_0}).$$

Since $\mathbb{E}[G_{f_0}] = 0$, it suffices to bound $\mathbb{E}[\sup_f (G_f - G_{f_0})]$. Replace G_f by $G_f - G_{f_0}$ so that $G_{f_0} = 0$.

If $N(\mathcal{F}, d, \varepsilon) = \infty$ for some ε, then the RHS is $+\infty$ and the bound is trivial. Hence assume finite covering numbers at all scales of interest (e.g. d-totally bounded).

Step 1 (Dyadic scales and nets). Let $\Delta := \operatorname{diam}(\mathcal{F})$ and define $\varepsilon_k := 2^{-k}\Delta$ for $k \ge 0$. For each $k \ge 0$, let $\mathcal{F}_k$ be a minimal ε_k-net, so $|\mathcal{F}_k| = N(\mathcal{F}, d, \varepsilon_k)$. For each $f \in \mathcal{F}$, choose a selector $\pi_k(f) \in \mathcal{F}_k$ with $d(f, \pi_k(f)) \le \varepsilon_k$. Under Assumption E.1 (or outer-expectation conventions), one can choose these selectors measurably on a countable dense subset and extend by approximation; we treat this standard measurability point as handled by the assumption.

Step 2 (Chaining decomposition). For any $K \ge 1$,

$$G_f = G_{\pi_0(f)} + \sum_{k=1}^{K} \big(G_{\pi_k(f)} - G_{\pi_{k-1}(f)}\big) + \big(G_f - G_{\pi_K(f)}\big).$$

Taking $\sup_f$ and using $\sup(A + B) \le \sup A + \sup B$ gives

$$\sup_{f\in\mathcal{F}} G_f \le \sup_f G_{\pi_0(f)} + \sum_{k=1}^{K} \sup_f \big(G_{\pi_k(f)} - G_{\pi_{k-1}(f)}\big) + \sup_f \big(G_f - G_{\pi_K(f)}\big). \tag{310}$$

Step 3 (Remainder term). For each fixed f, $G_f - G_{\pi_K(f)}$ is centered Gaussian with variance $d(f, \pi_K(f))^2 \le \varepsilon_K^2$. Thus it is ε_K-sub-Gaussian: $\mathbb{E}[\exp(\lambda(G_f - G_{\pi_K(f)}))] \le \exp(\lambda^2\varepsilon_K^2/2)$. By Lemma E.9 applied on a finite ε_K-net approximation and then letting the net refine (or by standard Gaussian continuity arguments), one obtains

$$\mathbb{E}\Big[\sup_{f\in\mathcal{F}} (G_f - G_{\pi_K(f)})\Big] \;\le\; C_0\,\varepsilon_K$$

for a universal C_0. In particular, $\varepsilon_K \to 0$ implies this remainder vanishes as $K \to \infty$.

Step 4 (Control each increment level by a finite maximum). Fix $k \ge 1$. For any f,

$$\operatorname{Var}\big(G_{\pi_k(f)} - G_{\pi_{k-1}(f)}\big) = d(\pi_k(f), \pi_{k-1}(f))^2 \le \big(d(\pi_k(f), f) + d(f, \pi_{k-1}(f))\big)^2 \le (\varepsilon_k + \varepsilon_{k-1})^2 \le (2\varepsilon_{k-1})^2.$$

Also, the pair $(\pi_k(f), \pi_{k-1}(f))$ ranges over $\mathcal{F}_k \times \mathcal{F}_{k-1}$, hence

$$\sup_{f\in\mathcal{F}} \big(G_{\pi_k(f)} - G_{\pi_{k-1}(f)}\big) \le \max_{(a,b)\in\mathcal{F}_k\times\mathcal{F}_{k-1}} (G_a - G_b).$$

The collection $\{G_a - G_b\}_{(a,b)}$ is centered Gaussian and each element has variance $\le (2\varepsilon_{k-1})^2$. Apply Lemma E.9 (Gaussian is sub-Gaussian with the same parameter) with $m \le |\mathcal{F}_k|\,|\mathcal{F}_{k-1}|$ and $\sigma = 2\varepsilon_{k-1}$:

$$\mathbb{E}\Big[\sup_f (G_{\pi_k(f)} - G_{\pi_{k-1}(f)})\Big] \le 2\varepsilon_{k-1}\sqrt{2\log(|\mathcal{F}_k|\,|\mathcal{F}_{k-1}|)} \le C_1\,\varepsilon_{k-1}\Big(\sqrt{\log|\mathcal{F}_k|} + \sqrt{\log|\mathcal{F}_{k-1}|}\Big)$$

for a universal constant C_1.

Step 5 (Sum and compare to an integral). Take expectations in (310), let $K \to \infty$ using Step 3, and use $|\mathcal{F}_k| \le N(\mathcal{F}, d, \varepsilon_k)$:

$$\mathbb{E}\Big[\sup_{f\in\mathcal{F}} G_f\Big] \le C_2 \sum_{k\ge1} \varepsilon_{k-1}\sqrt{\log N(\mathcal{F}, d, \varepsilon_k)} \le C_3 \sum_{k\ge1} \varepsilon_k \sqrt{\log N(\mathcal{F}, d, \varepsilon_k)}.$$

Since $\varepsilon \mapsto N(\mathcal{F}, d, \varepsilon)$ is nonincreasing, for $\varepsilon \in (\varepsilon_{k+1}, \varepsilon_k]$, $\sqrt{\log N(\mathcal{F}, d, \varepsilon)} \ge \sqrt{\log N(\mathcal{F}, d, \varepsilon_k)}$. Therefore

$$\int_{\varepsilon_{k+1}}^{\varepsilon_k} \sqrt{\log N(\mathcal{F}, d, \varepsilon)}\,d\varepsilon \ge (\varepsilon_k - \varepsilon_{k+1})\sqrt{\log N(\mathcal{F}, d, \varepsilon_k)} = \varepsilon_{k+1}\sqrt{\log N(\mathcal{F}, d, \varepsilon_k)} \asymp \varepsilon_k \sqrt{\log N(\mathcal{F}, d, \varepsilon_k)}.$$

Summing over k yields (309) up to a universal constant. □

Lemma E.11 (Dudley entropy integral for sub-Gaussian increment processes)**.** *Let $\{Z_f\}_{f\in\mathcal{F}}$ be a centered process such that for some pseudo-metric d,*

$$\mathbb{E}\Big[\exp\big(\lambda(Z_f - Z_g)\big)\Big] \;\le\; \exp\Big(\frac{\lambda^2 d(f,g)^2}{2}\Big) \qquad \forall\lambda\in\mathbb{R},\ \forall f,g\in\mathcal{F}.$$

Then there exists a universal constant $C > 0$ such that

$$\mathbb{E}\Big[\sup_{f\in\mathcal{F}} Z_f\Big] \;\le\; C\int_0^{\operatorname{diam}(\mathcal{F})} \sqrt{\log N(\mathcal{F}, d, \varepsilon)}\,d\varepsilon.$$

Proof. Repeat the proof of Lemma E.10, replacing Gaussianity by the assumed sub-Gaussian increment bound, and replacing the Gaussian maximum step by Lemma E.9. The increment variance calculation in Step 4 is replaced by the corresponding sub-Gaussian parameter bound $d(\pi_k(f), \pi_{k-1}(f)) \le 2\varepsilon_{k-1}$, which follows identically from the triangle inequality in d. □

Corollary E.12 (Conditional Dudley bound for Rademacher averages)**.** *Fix a sample X_1^n and define the empirical L_2 pseudo-metric*

$$d_n(f,g) \;:=\; \Big(\frac{1}{n}\sum_{i=1}^n (f(X_i) - g(X_i))^2\Big)^{1/2}.$$

Let

$$\mathcal{R}_n(f) := \frac{1}{n}\sum_{i=1}^n \epsilon_i f(X_i).$$

Then conditional on X_1^n, the process $\{\mathcal{R}_n(f)\}_{f\in\mathcal{F}}$ has sub-Gaussian increments w.r.t. the metric $\tilde{d}_n(f,g) := d_n(f,g)/\sqrt{n}$, and hence

$$\mathbb{E}_\epsilon\Big[\sup_{f\in\mathcal{F}} \mathcal{R}_n(f) \;\Big|\; X_1^n\Big] \;\le\; \frac{C}{\sqrt{n}}\int_0^{\operatorname{diam}(\mathcal{F}, d_n)} \sqrt{\log N(\mathcal{F}, d_n, \varepsilon)}\,d\varepsilon \tag{311}$$

for a universal constant $C > 0$.

Proof. **Step 1 (Sub-Gaussian increments).** Fix $f, g \in \mathcal{F}$ and write $a_i := f(X_i) - g(X_i)$. Then, conditional on X_1^n,

$$\mathcal{R}_n(f) - \mathcal{R}_n(g) = \frac{1}{n}\sum_{i=1}^n \epsilon_i a_i.$$

Using independence and $\mathbb{E}[e^{\lambda \epsilon a}] = \cosh(\lambda a) \le \exp(\lambda^2 a^2/2)$,

$$\mathbb{E}_\epsilon\Big[\exp\Big(\lambda(\mathcal{R}_n(f)-\mathcal{R}_n(g))\Big)\,\Big|\,X_1^n\Big] = \prod_{i=1}^n \mathbb{E}_\epsilon\Big[\exp\Big(\frac{\lambda}{n}\epsilon_i a_i\Big)\Big] \le \prod_{i=1}^n \exp\Big(\frac{\lambda^2 a_i^2}{2n^2}\Big) = \exp\Big(\frac{\lambda^2}{2n^2}\sum_{i=1}^n a_i^2\Big) = \exp\Big(\frac{\lambda^2}{2}\cdot\frac{d}{}$$

Thus the increment is sub-Gaussian with parameter $\tilde{d}_n(f,g) = d_n(f,g)/\sqrt{n}$.

Step 2 (Apply Lemma E.11 conditionally). Apply Lemma E.11 (conditionally on X_1^n) with metric $\tilde{d}_n$. Since $N(\mathcal{F}, \tilde{d}_n, \varepsilon) = N(\mathcal{F}, d_n, \varepsilon\sqrt{n})$, one obtains (311) after the change of variables $\varepsilon \mapsto \varepsilon/\sqrt{n}$. □

Remark E.13 (From Dudley to empirical-process bounds). Combining Corollary E.12 with Lemma E.6 yields

$$\mathbb{E}\Big[\sup_{f\in\mathcal{F}}(P_n - P)f\Big] \le \frac{2C}{\sqrt{n}}\mathbb{E}\left[\int_0^{\operatorname{diam}(\mathcal{F},d_n)} \sqrt{\log N(\mathcal{F}, d_n, \varepsilon)}\, d\varepsilon\right].$$

To obtain deterministic entropies (e.g. in $L_2(P)$ rather than $L_2(P_n)$), one typically upper bounds $N(\mathcal{F}, d_n, \varepsilon)$ by $N(\mathcal{F}, L_2(P), c\varepsilon)$ on a high-probability event, or uses bracketing entropies.

E.4. **Concentration inequalities for suprema.** We now record concentration inequalities that upgrade an expectation bound into a high-probability deviation bound.

Assumption E.14 (Bounded envelope)**.** Let $\mathcal{F} \subseteq \{f : \mathcal{X} \to \mathbb{R}\}$ satisfy $\sup_{f\in\mathcal{F}} \|f\|_\infty \le b$ and define the variance proxy

$$v \;:=\; \sup_{f\in\mathcal{F}} \operatorname{Var}(f(X)) \;\le\; \sup_{f\in\mathcal{F}} \mathbb{E}[f(X)^2].$$

Lemma E.15 (Bousquet/Talagrand inequality for suprema)**.** *Let $X_1, \dots, X_n$ be independent with common law P on $(\mathcal{X}, \mathcal{A})$ and let $P_n := \frac{1}{n}\sum_{i=1}^n \delta_{X_i}$. Let $\mathcal{F} \subseteq \{f : \mathcal{X} \to \mathbb{R}\}$ satisfy Assumption E.1 so that the supremum below is measurable. Define the supremum functional*

$$Z \;:=\; \sup_{f\in\mathcal{F}}(P_n - P)f.$$

Assume that there exist finite constants $b \ge 0$ and $v \ge 0$ such that the following two properties hold:

(A) Bounded differences (one-step Lipschitz). *For each $i \in \{1, \dots, n\}$ and every replacement $X_i \mapsto X_i'$ (with X_i' independent of all else and distributed as X_i), the resulting value $Z^{(i)}$ satisfies*

$$Z - Z^{(i)} \;\le\; \frac{b}{n} \qquad \text{a.s.} \tag{312}$$

(B) Predictable quadratic variation. *Let $\mathcal{F}_i := \sigma(X_1, \dots, X_i)$ and set the Doob martingale $M_i := \mathbb{E}[Z \mid \mathcal{F}_i]$ with differences $D_i := M_i - M_{i-1}$ ($M_0 := \mathbb{E}Z$). Assume the predictable quadratic variation satisfies*

$$\sum_{i=1}^n \mathbb{E}\big[D_i^2 \mid \mathcal{F}_{i-1}\big] \;\le\; \frac{v}{n} \qquad \text{a.s.} \tag{313}$$

Then for all $t > 0$, with probability at least $1 - e^{-t}$,

$$Z \;\le\; \mathbb{E}[Z] + \sqrt{\frac{2vt}{n}} + \frac{bt}{3n}. \tag{314}$$

Moreover, with probability at least $1 - e^{-t}$,

$$\sup_{f\in\mathcal{F}} \big|(P_n - P)f\big| \;\le\; \mathbb{E}\Big[\sup_{f\in\mathcal{F}} \big|(P_n - P)f\big|\Big] + \sqrt{\frac{2vt}{n}} + \frac{bt}{3n}. \tag{315}$$

Proof. We prove (314); the absolute-value form (315) is then immediate.

Step 0 (Reduce to a martingale tail bound). By construction, $(M_i)_{i=0}^n$ is a martingale with respect to $(\mathcal{F}_i)$ and

$$Z - \mathbb{E}Z \;=\; M_n - M_0 \;=\; \sum_{i=1}^{n} D_i.$$

Thus it suffices to prove a Bernstein/Freedman-type tail bound for $\sum_i D_i$.

Step 1 (Derive a uniform bound on each martingale increment). We claim that, under (312),

$$D_i \;\le\; \frac{b}{n} \qquad \text{a.s. for each } i. \tag{316}$$

To see this, fix i and condition on $\mathcal{F}_{i-1}$. Consider two possible values $x, x' \in \mathcal{X}$ for X_i, keeping X_{-i} fixed. Let $Z(x)$ denote the value of the supremum functional when X_i is set to x (and all other coordinates fixed). Then the conditional random variable Z given $\mathcal{F}_{i-1}$ can be viewed as $Z(X_i)$. The bounded-differences assumption (312) says that for any such pair (x, x'),

$$Z(x) - Z(x') \;\le\; \frac{b}{n}.$$

Taking conditional expectations with respect to X_i given $\mathcal{F}_{i-1}$ and using that $M_i = \mathbb{E}[Z \mid \mathcal{F}_i] = Z(X_i)$ and $M_{i-1} = \mathbb{E}[Z \mid \mathcal{F}_{i-1}] = \mathbb{E}[Z(X_i) \mid \mathcal{F}_{i-1}]$, we obtain the essential supremum bound $M_i - M_{i-1} \le b/n$, i.e. (316). (Equivalently: a random variable supported in an interval of length b/n differs from its conditional mean by at most b/n above.)

Step 2 (A sharp one-step Laplace bound for bounded increments). We prove the following deterministic inequality.

Claim E.16 (Bennett-type upper envelope)**.** *Fix $c > 0$. For every $\lambda \ge 0$ and every real $x \le c$,*

$$e^{\lambda x} \;\le\; 1 + \lambda x + \frac{e^{\lambda c} - \lambda c - 1}{c^2}\, x^2. \tag{317}$$

Proof of Claim E.16. Fix $\lambda \ge 0$ and define the function

$$\psi(x) \;:=\; e^{\lambda x} - 1 - \lambda x - \alpha x^2, \qquad \text{where} \qquad \alpha := \frac{e^{\lambda c} - \lambda c - 1}{c^2}.$$

Then $\psi(0) = 0$ and also $\psi(c) = 0$ by the definition of α. Moreover,

$$\psi''(x) = \lambda^2 e^{\lambda x} - 2\alpha.$$

Since $x \mapsto e^{\lambda x}$ is increasing, ψ'' is increasing in x. We also have $\psi''(c) = \lambda^2 e^{\lambda c} - 2\alpha \ge 0$ because $\alpha \le \frac{1}{2}\lambda^2 e^{\lambda c}$ (Taylor remainder bound). Hence ψ is convex on $(-\infty, c]$ once x is sufficiently close to c, and, with $\psi(0) = \psi(c) = 0$, this implies $\psi(x) \le 0$ for all $x \in [0, c]$; for $x < 0$ the inequality is even easier since $e^{\lambda x} \le 1 + \lambda x$ and the quadratic term is nonnegative. Thus $\psi(x) \le 0$ for all $x \le c$, which is (317). □

Step 3 (Conditional mgf control for the martingale differences). Fix $\lambda \in [0, \infty)$ and set $c := b/n$. Apply Claim E.16 to $x = D_i$ and $c = b/n$ and then take $\mathbb{E}[\cdot \mid \mathcal{F}_{i-1}]$. Using $\mathbb{E}[D_i \mid \mathcal{F}_{i-1}] = 0$ (martingale difference property) and the bound $D_i \le c$ from (316), we get

$$\begin{aligned}\mathbb{E}\Big[e^{\lambda D_i} \mid \mathcal{F}_{i-1}\Big] &\le 1 + \lambda \mathbb{E}[D_i \mid \mathcal{F}_{i-1}] + \frac{e^{\lambda c} - \lambda c - 1}{c^2}\, \mathbb{E}\big[D_i^2 \mid \mathcal{F}_{i-1}\big] \\ &= 1 + \frac{e^{\lambda c} - \lambda c - 1}{c^2}\, \mathbb{E}\big[D_i^2 \mid \mathcal{F}_{i-1}\big] \\ &\le \exp\left(\frac{e^{\lambda c} - \lambda c - 1}{c^2}\, \mathbb{E}\big[D_i^2 \mid \mathcal{F}_{i-1}\big]\right), \qquad (318)\end{aligned}$$

where the last step uses $1 + u \le e^u$ for all u.

Step 4 (Iterate the conditional mgf bound). Iterating (318) and using the tower property yields

$$\mathbb{E}\left[\exp\left(\lambda \sum_{i=1}^n D_i\right)\right] \;\le\; \mathbb{E}\left[\exp\left(\frac{e^{\lambda c} - \lambda c - 1}{c^2} \sum_{i=1}^n \mathbb{E}[D_i^2 \mid \mathcal{F}_{i-1}]\right)\right]. \qquad (319)$$

Now invoke the predictable-variation assumption (313):

$$\sum_{i=1}^n \mathbb{E}[D_i^2 \mid \mathcal{F}_{i-1}] \;\le\; \frac{v}{n} \quad \text{a.s.}$$

Therefore the exponent on the right-hand side of (319) is bounded above by the deterministic constant

$$\frac{e^{\lambda c} - \lambda c - 1}{c^2} \cdot \frac{v}{n},$$

so

$$\log \mathbb{E}[\exp(\lambda(Z - \mathbb{E}Z))] = \log \mathbb{E}\left[\exp\left(\lambda \sum_{i=1}^n D_i\right)\right] \;\le\; \frac{e^{\lambda c} - \lambda c - 1}{c^2} \cdot \frac{v}{n}, \qquad c = \frac{b}{n}. \qquad (320)$$

Step 5 (Convert Bennett form to sub-gamma form). We use the standard inequality (proved below)

$$e^u - u - 1 \;\le\; \frac{u^2}{2(1 - u/3)} \qquad \forall u \in [0, 3). \qquad (321)$$

Applying (321) with $u = \lambda c$ in (320) gives, for all $\lambda \in [0, 3/c)$,

$$\log \mathbb{E}\Big[e^{\lambda(Z - \mathbb{E}Z)}\Big] \;\le\; \frac{\lambda^2}{2\big(1 - \lambda c/3\big)} \cdot \frac{v}{n}, \qquad c = \frac{b}{n}. \qquad (322)$$

Claim E.17 (Proof of (321))**.** *For all $u \in [0, 3)$, $e^u - u - 1 \le \frac{u^2}{2(1-u/3)}$.*

Proof of Claim E.17. For $u \in [0, 3)$, expand $e^u - u - 1 = \sum_{k \ge 2} u^k/k!$ and note that for all $k \ge 3$,

$$\frac{1}{k!} \;\le\; \frac{1}{2 \cdot 3^{k-2}},$$

since $k! = 2 \cdot 3 \cdot 4 \cdots k \;\ge\; 2 \cdot 3^{k-2}$. Hence

$$e^u - u - 1 = \frac{u^2}{2} + \sum_{k \ge 3} \frac{u^k}{k!} \le \frac{u^2}{2} + \sum_{k \ge 3} \frac{u^k}{2 \cdot 3^{k-2}} = \frac{u^2}{2}\left(1 + \sum_{m \ge 1} \left(\frac{u}{3}\right)^m\right) = \frac{u^2}{2} \cdot \frac{1}{1 - u/3},$$

which is (321). □

Step 6 (Chernoff bound and explicit optimization). Fix $t > 0$ and $s > 0$. By Chernoff's method and (322), for any $\lambda \in (0, 3/c)$,

$$\Pr(Z - \mathbb{E}Z \geq s) \leq \exp\left(-\lambda s + \frac{\lambda^2}{2(1-\lambda c/3)} \cdot \frac{v}{n}\right).$$

A standard sub-gamma optimization (and it can be checked by direct substitution) shows that taking

$$s := \sqrt{\frac{2vt}{n}} + \frac{ct}{3} \qquad \left(\text{recall } c = \frac{b}{n}\right)$$

yields $\Pr(Z - \mathbb{E}Z \geq s) \leq e^{-t}$. Equivalently, with probability at least $1 - e^{-t}$,

$$Z \leq \mathbb{E}Z + \sqrt{\frac{2vt}{n}} + \frac{bt}{3n},$$

which is (314).

Step 7 (Absolute-value version). Let $\mathcal{F}^{\pm} := \mathcal{F} \cup (-\mathcal{F})$. Then

$$\sup_{f \in \mathcal{F}} |(P_n - P)f| = \sup_{g \in \mathcal{F}^{\pm}} (P_n - P)g.$$

Assumptions (312)–(313) are inherited by $\mathcal{F}^{\pm}$ with the same constants b, v (since replacing f by $-f$ does not change ranges or the martingale increment bounds). Applying (314) to $\mathcal{F}^{\pm}$ yields (315). □

Remark E.18 (How to use Lemma E.15 in empirical-process proofs). Lemma E.15 is deliberately modular: to apply it, you verify (A) and (B) for your specific supremum functional.

- In most empirical-process applications, (A) follows from a uniform envelope bound: if $\sup_{f \in \mathcal{F}} \operatorname{osc}(f) := \sup_x f(x) - \inf_x f(x) \leq b$, then changing one sample point changes every empirical mean by at most b/n, hence (312) holds.
- Property (B) is a *predictable variance* bound for the Doob martingale of Z. There are two common ways to ensure it: (i) directly bound $\sum_i \mathbb{E}[D_i^2 \mid \mathcal{F}_{i-1}]$ using structural properties of the class (this is what Talagrand/Bousquet do for empirical-process suprema); or (ii) in high-level arguments, *assume* (313) as part of Assumption E.14, since it is exactly the variance input needed for a Bernstein-type concentration.

Once (A) and (B) are in place, Lemma E.15 gives a sharp Bernstein tail around $\mathbb{E}Z$.

E.5. **Matrix concentration for operator-norm deviations.** For operator deviation bounds involving sums of random matrices (e.g. empirical Hessians or covariance-type objects), it is often preferable to use matrix Bernstein inequalities, which provide dimension-dependent tails without explicit discretization.

Setup. Let $Y_1, \ldots, Y_n$ be independent random self-adjoint matrices in $\mathbb{R}^{d \times d}$ with $\mathbb{E}[Y_i] = 0$. Define the matrix variance proxy

$$V := \left\| \sum_{i=1}^{n} \mathbb{E}\left[Y_i^2\right] \right\|_{\mathrm{op}}.$$

Lemma E.19 (Matrix Bernstein inequality (self-adjoint case))**.** *Assume* $\|Y_i\|_{\mathrm{op}} \le R$ *almost surely for all* i*. Then for all* $t > 0$*,*

$$\Pr\Big(\Big\|\sum_{i=1}^{n} Y_i\Big\|_{\mathrm{op}} \ \ge\ t\Big) \ \le\ 2d\cdot\exp\Big(-\frac{t^2/2}{V + Rt/3}\Big). \tag{323}$$

Consequently, with probability at least $1 - e^{-s}$*,*

$$\Big\|\sum_{i=1}^{n} Y_i\Big\|_{\mathrm{op}} \ \le\ \sqrt{2Vs} + \frac{2Rs}{3} + \sqrt{2Vs}\,\sqrt{\log(2d)} + \frac{2R}{3}\log(2d), \tag{324}$$

and in particular (absorbing $\log(2d)$ *into* s*) one obtains the familiar form*

$$\Big\|\sum_{i=1}^{n} Y_i\Big\|_{\mathrm{op}} \ \lesssim\ \sqrt{V(\log d + s)} + R(\log d + s).$$

Proof. A full proof uses the matrix Laplace transform method: (i) apply Markov's inequality to $\operatorname{Tr}\exp(\lambda\sum_i Y_i)$, (ii) use a matrix MGF bound $\mathbb{E}[\exp(\lambda Y_i)] \preceq \exp\big(\frac{\lambda^2}{2(1-\lambda R/3)}\mathbb{E}[Y_i^2]\big)$ valid for $\lambda \in (0, 3/R)$, and (iii) control $\operatorname{Tr}\exp(\cdot)$ by $d\cdot\exp(\lambda_{\max}(\cdot))$. We provide the critical steps explicitly.

Step 1 (Laplace transform and trace domination). For $\lambda > 0$,

$$\Pr\Big(\lambda_{\max}\Big(\sum_i Y_i\Big) \ge t\Big) = \Pr\Big(\exp\Big(\lambda\lambda_{\max}\Big(\sum_i Y_i\Big)\Big) \ge e^{\lambda t}\Big) \le e^{-\lambda t}\,\mathbb{E}\Big[\exp\Big(\lambda\lambda_{\max}\Big(\sum_i Y_i\Big)\Big)\Big].$$

Since $\exp(\lambda\lambda_{\max}(A)) \le \operatorname{Tr}(e^{\lambda A})$ for self-adjoint A (because Tr sums eigenvalues and dominates the maximum),

$$\Pr\Big(\lambda_{\max}\Big(\sum_i Y_i\Big) \ge t\Big) \le e^{-\lambda t}\,\mathbb{E}\operatorname{Tr}\exp\Big(\lambda\sum_i Y_i\Big).$$

Step 2 (Iterative conditioning and MGF bound). Let $S_k := \sum_{i=1}^{k} Y_i$. By iterated expectation and conditional Jensen together with the (noncommutative) Golden–Thompson/Lieb machinery, one can show

$$\mathbb{E}\operatorname{Tr} e^{\lambda S_n} \le \operatorname{Tr}\exp\Big(\sum_{i=1}^{n}\log\mathbb{E}e^{\lambda Y_i}\Big). \tag{325}$$

This is the standard reduction that replaces the MGF of a sum by a product of MGFs inside a trace exponential.

Step 3 (Bernstein-type MGF estimate for each summand). Under $\|Y_i\|_{\mathrm{op}} \le R$ and $\mathbb{E}[Y_i] = 0$, a matrix-valued Bernstein MGF bound holds: for $\lambda \in (0, 3/R)$,

$$\log\mathbb{E}e^{\lambda Y_i} \ \preceq\ \frac{\lambda^2}{2(1-\lambda R/3)}\,\mathbb{E}[Y_i^2]. \tag{326}$$

This follows by comparing the power series of $e^{\lambda Y_i}$ to $I + \lambda Y_i + \frac{\lambda^2}{2(1-\lambda R/3)}Y_i^2$ using the scalar inequality $e^y \le 1 + y + \frac{y^2}{2(1-y/3)}$ for $|y| < 3$ applied to the eigenvalues, and then taking expectations and the operator log.

Step 4 (Assemble the bound). Insert (326) into (325):

$$\mathbb{E}\operatorname{Tr} e^{\lambda S_n} \le \operatorname{Tr}\exp\Big(\frac{\lambda^2}{2(1-\lambda R/3)}\sum_{i=1}^{n}\mathbb{E}[Y_i^2]\Big).$$

Let $W := \sum_i \mathbb{E}[Y_i^2] \succeq 0$. Then $\operatorname{Tr} e^{\alpha W} \le d \cdot e^{\alpha \|W\|_{\mathrm{op}}} = d \cdot e^{\alpha V}$, so

$$\mathbb{E} \operatorname{Tr} e^{\lambda S_n} \le d \cdot \exp\Big(\frac{\lambda^2}{2(1-\lambda R/3)} V\Big).$$

Hence

$$\Pr\big(\lambda_{\max}(S_n) \ge t\big) \le d \cdot \exp\Big(-\lambda t + \frac{\lambda^2}{2(1-\lambda R/3)} V\Big).$$

Optimize over $\lambda \in (0, 3/R)$ to obtain

$$\Pr\big(\lambda_{\max}(S_n) \ge t\big) \le d \exp\Big(-\frac{t^2/2}{V + Rt/3}\Big).$$

Applying the same bound to $-S_n$ controls $\lambda_{\min}(S_n)$ and yields the two-sided operator norm bound (323). The inverted high-probability forms follow by solving the inequality for t and using $\log(2d)$. □

Remark E.20 (How to use matrix Bernstein for empirical Hessians (clean plug-in form)). Let $H : \mathcal{X} \to \mathbb{S}^d$ be a measurable self-adjoint (symmetric) matrix-valued function. Given i.i.d. $X_1, \dots, X_n \sim P$, define the centered summands

$$Y_i \;:=\; H(X_i) - \mathbb{E}[H(X)] \;\in\; \mathbb{S}^d, \qquad i = 1, \dots, n,$$

so that $\mathbb{E}[Y_i] = 0$ and

$$(P_n - P)H \;=\; \frac{1}{n} \sum_{i=1}^n Y_i.$$

Assume there exists $R < \infty$ such that

$$\|Y_i\|_{\mathrm{op}} \;\le\; R \qquad \text{a.s. for all } i, \tag{327}$$

and define the matrix-variance proxy

$$V \;:=\; \Big\| \sum_{i=1}^n \mathbb{E}[Y_i^2] \Big\|_{\mathrm{op}} \;=\; n \, \big\| \mathbb{E}[Y_1^2] \big\|_{\mathrm{op}}, \tag{328}$$

where $Y_i^2 := Y_i Y_i$ (well-defined since Y_i is self-adjoint).

Then the (self-adjoint) matrix Bernstein inequality (e.g. Tropp) yields for all $t \ge 0$,

$$\Pr\Big(\Big\| \sum_{i=1}^n Y_i \Big\|_{\mathrm{op}} \;\ge\; t \Big) \;\le\; 2d \cdot \exp\left(-\frac{t^2}{2V + \frac{2}{3} R t} \right). \tag{329}$$

Equivalently, letting $\delta \in (0,1)$ and solving (329) for t gives the explicit high-probability bound: with probability at least $1 - \delta$,

$$\Big\| \sum_{i=1}^n Y_i \Big\|_{\mathrm{op}} \;\le\; \sqrt{2V \log\Big(\frac{2d}{\delta}\Big)} \;+\; \frac{2R}{3} \log\Big(\frac{2d}{\delta}\Big). \tag{330}$$

Dividing by n yields the empirical-process form: with probability at least $1 - \delta$,

$$\|(P_n - P)H\|_{\mathrm{op}} \;\le\; \sqrt{\frac{2 \| \mathbb{E}[Y_1^2] \|_{\mathrm{op}}}{n} \log\Big(\frac{2d}{\delta}\Big)} \;+\; \frac{2R}{3n} \log\Big(\frac{2d}{\delta}\Big). \tag{331}$$

Interpretation for EM Hessians. In EM analyses, $H(X)$ is typically an observed-information or complete-data Hessian contribution (or a mixed Hessian block). Equation (331) gives a *dimension-aware* operator-norm deviation without any explicit sphere-net/union bound. This is usually preferable to net-based scalarization when d is moderate-to-large.

Two technical notes.

(a) If $H(X)$ is not exactly symmetric, apply the bound to its symmetrization $\frac{1}{2}(H+H^\top)$, since $\|(P_n-P)H\|_{\mathrm{op}} = \|(P_n-P)\frac{1}{2}(H+H^\top)\|_{\mathrm{op}}$.

(b) In applications, R and $\|\mathbb{E}[Y_1^2]\|_{\mathrm{op}}$ can be bounded using (i) almost-sure operator bounds on $H(X)$, and (ii) second-moment bounds; for instance, $\mathbb{E}[Y_1^2] \preceq \mathbb{E}[H(X)^2]$ and $\|Y_1\|_{\mathrm{op}} \le 2\|H(X)\|_{\mathrm{op}}$ a.s.

E.6. **Parameter-indexed classes: covering numbers from Lipschitzness.**

Lemma E.21 (Covering numbers via parameter Lipschitzness (fully explicit))**.** *Let $(\mathcal{F}, \|\cdot\|)$ be a seminormed linear space of real-valued measurable functions on $\mathcal{X}$ (e.g. $\|g\| = \|g\|_{L_2(P)}$). Let $\Theta_0 \subseteq \mathbb{R}^p$ and consider the parameter-indexed class*

$$\mathcal{F}_0 \;:=\; \{f_\theta : \theta \in \Theta_0\} \;\subseteq\; \mathcal{F}.$$

Assume there exists $L > 0$ such that for all $\theta, \vartheta \in \Theta_0$,

$$\|f_\theta - f_\vartheta\| \;\le\; L\|\theta - \vartheta\|_2. \tag{332}$$

Then for every $\varepsilon > 0$,

$$N\big(\mathcal{F}_0, \|\cdot\|, \varepsilon\big) \;\le\; N\big(\Theta_0, \|\cdot\|_2, \varepsilon/L\big). \tag{333}$$

In particular, if Θ_0 is contained in the Euclidean ball $B_2(0,R) \subseteq \mathbb{R}^p$, then for every $\varepsilon > 0$,

$$N\big(\mathcal{F}_0, \|\cdot\|, \varepsilon\big) \;\le\; \Big(1 + \frac{2LR}{\varepsilon}\Big)^p. \tag{334}$$

Proof. **Step 1 (Push forward a parameter net to a function net).** Fix $\varepsilon > 0$ and let $\{\theta_1, \dots, \theta_m\} \subseteq \Theta_0$ be an (ε/L)-net of Θ_0 in $\|\cdot\|_2$, so that for every $\theta \in \Theta_0$ there exists j with

$$\|\theta - \theta_j\|_2 \;\le\; \varepsilon/L.$$

By the Lipschitz assumption (332),

$$\|f_\theta - f_{\theta_j}\| \;\le\; L\|\theta - \theta_j\|_2 \;\le\; \varepsilon.$$

Thus $\{f_{\theta_j}\}_{j=1}^m$ is an ε-net of $\mathcal{F}_0$ in the seminorm $\|\cdot\|$.

Step 2 (Conclude the covering-number inequality). By definition of covering number, the existence of an ε-net of $\mathcal{F}_0$ of size m implies $N(\mathcal{F}_0, \|\cdot\|, \varepsilon) \le m$. Minimizing over all (ε/L)-nets of Θ_0 yields (333).

Step 3 (Volumetric specialization on a Euclidean ball). If $\Theta_0 \subseteq B_2(0,R)$, then $N(\Theta_0, \|\cdot\|_2, \varepsilon/L) \le N(B_2(0,R), \|\cdot\|_2, \varepsilon/L)$. Apply Lemma E.22 below with $\rho = \varepsilon/L$ to get $N(\Theta_0, \|\cdot\|_2, \varepsilon/L) \le (1+2R/(\varepsilon/L))^p = (1+2LR/\varepsilon)^p$, and combine with (333) to obtain (334). □

Lemma E.22 (Volumetric covering bound for Euclidean balls)**.** *For $p \in \mathbb{N}$, $R > 0$, and $\rho \in (0, 2R]$,*

$$N\big(B_2(0,R), \|\cdot\|_2, \rho\big) \;\le\; \Big(1 + \frac{2R}{\rho}\Big)^p. \tag{335}$$

Proof. Let $\mathcal{N} \subseteq B_2(0,R)$ be a maximal ρ-separated set: $\|x-y\|_2 > \rho$ for distinct $x, y \in \mathcal{N}$, and no further point can be added while preserving separation. Maximality implies $\mathcal{N}$ is a ρ-net of $B_2(0,R)$ (otherwise a point at distance $> \rho$ from all net points could be added).

Now consider Euclidean balls $\{B_2(x, \rho/2) : x \in \mathcal{N}\}$. They are pairwise disjoint by ρ-separation, and all are contained in $B_2(0, R+\rho/2)$ because $\|x\| \le R$ implies $B_2(x,\rho/2) \subseteq B_2(0, R+\rho/2)$. By comparing volumes,

$$|\mathcal{N}| \cdot \mathrm{Vol}\big(B_2(0,\rho/2)\big) \;\le\; \mathrm{Vol}\big(B_2(0, R+\rho/2)\big).$$

Using $\mathrm{Vol}(B_2(0,r)) = c_p r^p$ for a constant c_p depending only on p,

$$|\mathcal{N}| \;\le\; \left(\frac{R+\rho/2}{\rho/2}\right)^p \;=\; \left(1+\frac{2R}{\rho}\right)^p.$$

Since $\mathcal{N}$ is a ρ-net, $N(B_2(0,R), \|\cdot\|_2, \rho) \le |\mathcal{N}|$, proving (335). □

Remark E.23 (Combining everything: a typical EM deviation pipeline (more explicit)). A common, fully explicit route to uniform deviation bounds in EM proofs is:

(1) **Choose the object and its norm.** Typical targets are vector gradients $\sup_{\theta\in\Theta_0}\|(P_n - P)g_\theta\|$ or matrix Hessians $\sup_{\theta\in\Theta_0}\|(P_n - P)H_\theta\|_{\mathrm{op}}$.
(2) **If needed, scalarize.** For vector/matrix *suprema* over classes, scalarize via sphere nets as in Appendix E.1 and union-bound over net points. This reduces the problem to finitely many scalar empirical processes.
(3) **Control the scalar supremum in expectation.** Use symmetrization (Lemma E.6) to reduce $\mathbb{E}[\sup_f (P_n - P)f]$ to a Rademacher average. Then bound the Rademacher average via entropy integrals / chaining: combine parameter-to-function covering reductions (Lemma E.21) with Dudley/generic chaining (Remark E.13).
(4) **Upgrade to high probability.** Apply a Bernstein-type concentration inequality for suprema (Lemma E.15) and, if scalarization was used, add the union bound over net points.
(5) **Optional shortcut for matrix sums.** If the object is already a sum of self-adjoint matrices (e.g. empirical Hessians), bypass explicit nets and scalarization using matrix Bernstein directly (Remark E.20).

This appendix is organized so each step is a direct plug-in lemma with explicit constants and no hidden reductions.

Appendix F. Canonical instantiations under nonidentifiability and misspecification

This appendix instantiates the symmetry and quotient-measurability templates used in the main text for several canonical model classes. For each class we verify: (i) a latent-transport representation (Assumption C.1); (ii) the induced observed invariance $P^X_{g\cdot\theta} = P^X_\theta$; (iii) posterior transport and Q-equivariance (Appendix C); and (iv) where needed, explicit measurable sections of the quotient map.

Standing measurable setup and notation. Let $(\mathcal{X}, \mathcal{F})$ be the observed space and fix a σ-finite dominating measure μ on $(\mathcal{X}, \mathcal{F})$. Assume $P^\star \ll \mu$ throughout Appendix F. Whenever latent variables are used, let $(\mathcal{Z}, G)$ be standard Borel with a σ-finite measure ν, and assume Appendix A (joint domination and existence of regular conditional laws).

Given $\theta \in \Theta$ with complete-data density $p_\theta(x,z)$ w.r.t. $\mu\otimes\nu$, write the marginal density

$$p_\theta(x) := \int_{\mathcal{Z}} p_\theta(x,z)\,\nu(dz) \in [0,\infty],$$

and let $\kappa_\theta(x,\cdot)$ denote a regular conditional distribution of $Z \mid X = x$ under P_θ; when $p_\theta(x) > 0$ one may take

$$\kappa_\theta(x,B) = \frac{\int_B p_\theta(x,z)\,\nu(dz)}{p_\theta(x)} \qquad (B \in G),$$

and define $\kappa_\theta(x,\cdot)$ arbitrarily on $\{x : p_\theta(x) = 0\}$.

Let G be a group acting on Θ (Appendix B) and write $q : \Theta \to \Theta/G$ for the orbit map.

F.1. **A general quotient-measurability principle.** The following lemma is repeatedly used (often implicitly) to pass from invariance on Θ to well-defined objects on Θ/G.

Lemma F.1 (Invariant functions factor through the quotient)**.** *Let $(\Theta, \mathcal{B}(\Theta))$ be a measurable space and let G be a group acting on Θ. Assume the action is* measurable in the usual sense *that for each fixed $g \in G$, the map $T_g : \Theta \to \Theta$, $T_g(\theta) = g \cdot \theta$, is $\mathcal{B}(\Theta)$–measurable. Let $q : \Theta \to \Theta/G$ be the orbit map, and equip Θ/G with the* quotient σ-field

$$\mathcal{B}(\Theta/G) := \big\{A \subseteq \Theta/G : \ q^{-1}(A) \in \mathcal{B}(\Theta)\big\}.$$

Let $\overline{\mathbb{R}} := \mathbb{R} \cup \{\pm\infty\}$ with its Borel σ-field $\mathcal{B}(\overline{\mathbb{R}})$. If $f : \Theta \to \overline{\mathbb{R}}$ is $\mathcal{B}(\Theta)$–measurable and G-invariant, i.e.

$$f(g \cdot \theta) = f(\theta) \qquad \forall g \in G, \ \forall \theta \in \Theta,$$

then there exists a unique *$\mathcal{B}(\Theta/G)$–measurable function $\bar{f} : \Theta/G \to \overline{\mathbb{R}}$ such that*

$$f = \bar{f} \circ q.$$

Proof. **Step 0 (Two basic facts about the quotient σ-field).**

(a) By definition of $\mathcal{B}(\Theta/G)$, the orbit map $q : (\Theta, \mathcal{B}(\Theta)) \to (\Theta/G, \mathcal{B}(\Theta/G))$ is measurable: indeed, for any $A \in \mathcal{B}(\Theta/G)$ we have $q^{-1}(A) \in \mathcal{B}(\Theta)$ by definition.
(b) The map q is surjective: every orbit class $\bar{\theta} \in \Theta/G$ is of the form $q(\theta)$ for some $\theta \in \Theta$.

Step 1 (Define $\bar{f}$ and verify it is well-defined). For $\bar{\theta} \in \Theta/G$, pick any $\theta \in \Theta$ such that $q(\theta) = \bar{\theta}$ and define

$$\bar{f}(\bar{\theta}) \ := \ f(\theta). \tag{336}$$

We must check that (336) does not depend on the choice of representative. If $\theta_1, \theta_2 \in \Theta$ satisfy $q(\theta_1) = q(\theta_2)$, then θ_1 and θ_2 lie in the same orbit, so there exists $g \in G$ with $\theta_2 = g \cdot \theta_1$. By G-invariance of f,

$$f(\theta_2) = f(g \cdot \theta_1) = f(\theta_1).$$

Thus $\bar{f}$ is well-defined.

Step 2 (Factorization identity $f = \bar{f} \circ q$). For any $\theta \in \Theta$, we have $\bar{f}(q(\theta)) = f(\theta)$ by the definition (336) applied to the orbit $q(\theta)$. Hence $f = \bar{f} \circ q$ pointwise on Θ.

Step 3 (Uniqueness). Suppose $\tilde{f} : \Theta/G \to \overline{\mathbb{R}}$ also satisfies $f = \tilde{f} \circ q$. Fix $\bar{\theta} \in \Theta/G$ and choose $\theta \in \Theta$ with $q(\theta) = \bar{\theta}$ (surjectivity). Then

$$\tilde{f}(\bar{\theta}) = \tilde{f}(q(\theta)) = f(\theta) = \bar{f}(q(\theta)) = \bar{f}(\bar{\theta}).$$

Thus $\tilde{f} = \bar{f}$ everywhere on Θ/G, proving uniqueness.

Step 4 (Measurability of $\bar{f}$). Let $B \in \mathcal{B}(\overline{\mathbb{R}})$ be an arbitrary Borel set. Consider $\bar{f}^{-1}(B) \subseteq \Theta/G$. Using $f = \bar{f} \circ q$, we compute its pullback:

$$q^{-1}\big(\bar{f}^{-1}(B)\big) = \{\theta \in \Theta : \ \bar{f}(q(\theta)) \in B\} = \{\theta \in \Theta : \ f(\theta) \in B\} = f^{-1}(B).$$

Since f is $\mathcal{B}(\Theta)$–measurable, $f^{-1}(B) \in \mathcal{B}(\Theta)$. Therefore $q^{-1}(\bar{f}^{-1}(B)) \in \mathcal{B}(\Theta)$, and by the definition of the quotient σ-field, $\bar{f}^{-1}(B) \in \mathcal{B}(\Theta/G)$. As this holds for every $B \in \mathcal{B}(\overline{\mathbb{R}})$, the function $\bar{f}$ is measurable. □

Remark F.2 (Quotient objective under misspecification (fully explicit)). Assume *observed invariance*:

$$P^X_{g\cdot\theta} = P^X_\theta \qquad \forall g \in G,\ \forall \theta \in \Theta.$$

Assume furthermore that each P^X_θ admits a density p_θ w.r.t. a common dominating measure μ, and that the invariance holds at the density level μ-a.e. (equivalently, $p_{g\cdot\theta} = p_\theta$ μ-a.e.). Define the population log-likelihood under misspecification,

$$\Phi(\theta) \ := \ \mathbb{E}_{P^\star}\big[\log p_\theta(X)\big] \quad \text{(whenever the expectation is well-defined in } (-\infty,\infty]).$$

Then Φ is G-invariant: for any $g \in G$ and $\theta \in \Theta$,

$$\Phi(g\cdot\theta) = \mathbb{E}_{P^\star}\big[\log p_{g\cdot\theta}(X)\big] = \mathbb{E}_{P^\star}\big[\log p_\theta(X)\big] = \Phi(\theta),$$

because $p_{g\cdot\theta}(x) = p_\theta(x)$ μ-a.e. and hence $P^\star$-a.s. on the support where the log is evaluated. Therefore, by Lemma F.1, Φ factors through the quotient: there exists a unique measurable $\bar\Phi : \Theta/G \to \overline{\mathbb{R}}$ such that $\Phi = \bar\Phi \circ q$. Consequently, the orbit-level maximizer set

$$\arg\max_{\bar\theta \in \Theta/G} \bar\Phi(\bar\theta)$$

is the natural identifiable target in the presence of nonidentifiability induced by G-symmetry.

Lemma F.3 (Mixture-model transport, invariance, posterior transport, and EM equivariance)**.** *Consider the finite-mixture complete-data density*

$$p_\theta(x,z) := \pi_z\, f(x;\vartheta_z), \qquad (x,z) \in \mathcal{X}\times[k],$$

with parameter space $\Theta = \Delta_{k-1}\times\Xi^k$ *and the* S_k*-action*

$$(g\cdot\theta) := (\pi_{g^{-1}(1)},\dots,\pi_{g^{-1}(k)},\ \vartheta_{g^{-1}(1)},\dots,\vartheta_{g^{-1}(k)}), \qquad g \in S_k.$$

Let ν *be counting measure on* $[k]$*, and define* $\tau_g : [k] \to [k]$ *by* $\tau_g(z) := g(z)$*. Then Assumption C.1 holds (with* τ_g *and* ν*), and consequently:*

(1) Observed invariance: $P^X_{g\cdot\theta} = P^X_\theta$ *for all* $g \in S_k$ *and* $\theta \in \Theta$*. Equivalently, the observed density* $p_\theta(x) = \sum_{z=1}^k \pi_z f(x;\vartheta_z)$ *is invariant:* $p_{g\cdot\theta}(x) = p_\theta(x)$ *for all* x*.*

(2) Posterior transport: *for all* $B \subseteq [k]$ *and for all* x *such that* $p_\theta(x) > 0$*,*

$$\kappa_{g\cdot\theta}(x,B) = \kappa_\theta(x, g^{-1}(B)).$$

(3) Q-equivariance and M-step equivariance: *for all* $\theta,\theta' \in \Theta$ *and* $g \in S_k$*,*

$$Q(g\cdot\theta \mid g\cdot\theta') = Q(\theta \mid \theta')$$

whenever both sides are well-defined. Consequently, for the M-step correspondence $\mathcal{T}(\theta') \in \arg\max_{\theta\in\Theta} Q(\theta \mid \theta')$*,*

$$\mathcal{T}(g\cdot\theta') \ = \ g\cdot\mathcal{T}(\theta') \qquad \textit{(equivariance at the level of argmax sets).}$$

Proof. **Step 1 (Verify the transport identity for the complete-data density).** Fix $g \in S_k$, $\theta \in \Theta$, and $(x,z) \in \mathcal{X}\times[k]$. By definition of the action and $\tau_g(z) = g(z)$,

$$p_{g\cdot\theta}(x,\tau_g(z)) = p_{g\cdot\theta}(x,g(z)) = \pi'_{g(z)}\, f\big(x;\vartheta'_{g(z)}\big),$$

where (π',ϑ') denote the parameters after permutation, i.e. $\pi'_j = \pi_{g^{-1}(j)}$ and $\vartheta'_j = \vartheta_{g^{-1}(j)}$. Hence

$$p_{g\cdot\theta}(x,g(z)) = \pi_{g^{-1}(g(z))}\, f\big(x;\vartheta_{g^{-1}(g(z))}\big) = \pi_z\, f(x;\vartheta_z) = p_\theta(x,z).$$

Thus the exact transport relation holds:

$$p_{g\cdot\theta}(x,\tau_g(z)) \ = \ p_\theta(x,z) \qquad \forall (x,z). \tag{337}$$

Step 2 (Check the group and measure-preservation properties of τ_g). Because τ_g is a permutation of the finite set $[k]$,

$$\tau_{gh} = \tau_g \circ \tau_h, \qquad \tau_e = \mathrm{id},$$

and counting measure ν is preserved: $\nu(\tau_g^{-1}(B)) = \nu(B)$ for all $B \subseteq [k]$. These are exactly the structural requirements typically imposed in latent-transport assumptions.

Step 3 (Observed invariance). The observed density is obtained by marginalization over z (a finite sum):

$$p_{g\cdot\theta}(x) = \sum_{z=1}^{k} p_{g\cdot\theta}(x, z).$$

Using (337) with $z' = \tau_g(z) = g(z)$ (a bijection of $[k]$),

$$\sum_{z=1}^{k} p_{g\cdot\theta}(x, z) = \sum_{z=1}^{k} p_{g\cdot\theta}(x, g(z)) = \sum_{z=1}^{k} p_\theta(x, z) = p_\theta(x).$$

Therefore $p_{g\cdot\theta}(x) = p_\theta(x)$ for all x, i.e. $P^X_{g\cdot\theta} = P^X_\theta$.

Step 4 (Posterior transport). For x with $p_\theta(x) > 0$ (equivalently $p_{g\cdot\theta}(x) > 0$ by Step 3), Bayes' rule yields the posterior probabilities

$$\kappa_\theta(x, \{z\}) = \Pr_\theta(Z = z \mid X = x) = \frac{p_\theta(x, z)}{p_\theta(x)}, \qquad \kappa_{g\cdot\theta}(x, \{z\}) = \frac{p_{g\cdot\theta}(x, z)}{p_{g\cdot\theta}(x)}.$$

Using Step 3 for the denominator and (337) for the numerator (rewritten as $p_{g\cdot\theta}(x, z) = p_\theta(x, g^{-1}(z))$),

$$\kappa_{g\cdot\theta}(x, \{z\}) = \frac{p_\theta(x, g^{-1}(z))}{p_\theta(x)} = \kappa_\theta(x, \{g^{-1}(z)\}).$$

Extend from singletons to arbitrary $B \subseteq [k]$ by additivity (finite sum):

$$\kappa_{g\cdot\theta}(x, B) = \sum_{z\in B} \kappa_{g\cdot\theta}(x, \{z\}) = \sum_{z\in B} \kappa_\theta(x, \{g^{-1}(z)\}) = \sum_{z'\in g^{-1}(B)} \kappa_\theta(x, \{z'\}) = \kappa_\theta(x, g^{-1}(B)).$$

Step 5 (Q-equivariance). Write $Q(\theta \mid \theta')$ in the concrete finite-mixture form (whenever integrable):

$$Q(\theta \mid \theta') = \mathbb{E}_{P^\star}\left[\sum_{z=1}^{k} \kappa_{\theta'}(X, \{z\}) \log p_\theta(X, z)\right].$$

Then, using Step 4 (posterior transport) and Step 1 (complete-data transport),

$$\begin{aligned}
Q(g\cdot\theta \mid g\cdot\theta') &= \mathbb{E}_{P^\star}\left[\sum_{z=1}^{k} \kappa_{g\cdot\theta'}(X, \{z\}) \log p_{g\cdot\theta}(X, z)\right] \\
&= \mathbb{E}_{P^\star}\left[\sum_{z=1}^{k} \kappa_{\theta'}(X, \{g^{-1}(z)\}) \log p_\theta(X, g^{-1}(z))\right] \qquad \text{(Steps 4 and 1)} \\
&= \mathbb{E}_{P^\star}\left[\sum_{z'=1}^{k} \kappa_{\theta'}(X, \{z'\}) \log p_\theta(X, z')\right] \qquad \text{(change variable } z' = g^{-1}(z)\text{)} \\
&= Q(\theta \mid \theta').
\end{aligned}$$

Step 6 (M-step equivariance). Let $\mathcal{T}(\theta') \in \arg\max_{\theta\in\Theta} Q(\theta \mid \theta')$ denote the argmax correspondence. By Step 5, for any θ,

$$Q(\theta \mid \theta') = Q(g\cdot\theta \mid g\cdot\theta').$$

Therefore θ maximizes $Q(\cdot \mid \theta')$ if and only if $g\cdot\theta$ maximizes $Q(\cdot \mid g\cdot\theta')$, which is precisely the equivariance identity $\mathcal{T}(g\cdot\theta') = g\cdot\mathcal{T}(\theta')$. □

A global measurable section under a finite action. We record a single reusable selector for *any* finite group action on a Borel subset of Euclidean space.

Assumption F.4 (Borel embedding and Borel action)**.** Θ is identified with a Borel subset of $\mathbb{R}^D$ for some D, and for each $g \in G$ the map $\theta \mapsto g\cdot\theta$ is Borel.

Fix a Borel linear order $\preceq$ on $\mathbb{R}^D$ (e.g. lexicographic order).

Lemma F.5 (Borel section for finite actions via orbit minimization)**.** *Let $(\Theta, \mathcal{B}(\Theta))$ be a* standard Borel space *and let G be a finite group acting on Θ. Assume the action is* Borel *in the sense that for each $g \in G$ the map*

$$T_g : \Theta \to \Theta, \qquad T_g(\theta) := g\cdot\theta$$

is $\mathcal{B}(\Theta)$–measurable. Fix a Borel linear order $\preceq$ *on Θ, i.e. a total order such that its graph*

$$\Gamma_{\preceq} := \{(\theta,\vartheta) \in \Theta\times\Theta : \ \theta \preceq \vartheta\}$$

belongs to $\mathcal{B}(\Theta)\otimes\mathcal{B}(\Theta)$.[2] *Define the orbit-minimizing selector*

$$\tilde{s}(\theta) \ := \ \min_{\preceq}\{g\cdot\theta : \ g \in G\}. \tag{338}$$

Then:

(a) $\tilde{s} : \Theta \to \Theta$ is Borel measurable;
(b) $\tilde{s}$ is constant on orbits: $\tilde{s}(g\cdot\theta) = \tilde{s}(\theta)$ for all $g \in G$ and $\theta \in \Theta$;
(c) letting $q : \Theta \to \Theta/G$ be the orbit map and equipping Θ/G with the quotient σ-field

$$\mathcal{B}(\Theta/G) := \{A \subseteq \Theta/G : \ q^{-1}(A) \in \mathcal{B}(\Theta)\},$$

there exists a unique measurable map $s : \Theta/G \to \Theta$ such that $\tilde{s} = s\circ q$;
(d) the map s is a measurable section*: for all $\bar{\theta} \in \Theta/G$, one has $q(s(\bar{\theta})) = \bar{\theta}$.*

Proof. Write $G = \{g_1, \dots, g_m\}$ with $m := |G|$.

Step 0 (Measurable comparison sets). For each pair $(i,j) \in \{1,\dots,m\}^2$, define

$$A_{ij} \ := \ \{\theta \in \Theta : \ g_i\cdot\theta \ \preceq \ g_j\cdot\theta\}.$$

We claim $A_{ij} \in \mathcal{B}(\Theta)$. Indeed, consider the Borel map

$$\Theta \ni \theta \ \longmapsto \ (g_i\cdot\theta, \ g_j\cdot\theta) \in \Theta\times\Theta,$$

which is measurable because $\theta \mapsto g_i\cdot\theta$ and $\theta \mapsto g_j\cdot\theta$ are measurable and products of measurable maps are measurable. Then

$$A_{ij} = \{\,\theta : \ (g_i\cdot\theta, \ g_j\cdot\theta) \in \Gamma_{\preceq}\,\} = \big(\theta \mapsto (g_i\cdot\theta, g_j\cdot\theta)\big)^{-1}(\Gamma_{\preceq}),$$

and since $\Gamma_{\preceq} \in \mathcal{B}(\Theta)\otimes\mathcal{B}(\Theta)$, we conclude $A_{ij} \in \mathcal{B}(\Theta)$.

[2] Every standard Borel space admits such an order: identify Θ Borel-isomorphically with a Borel subset of $[0,1]$ and pull back the usual order.

Step 1 (A Borel partition selecting an argmin index). Define recursively Borel sets $B_1, \dots, B_m \subseteq \Theta$ by

$$B_1 \;:=\; \bigcap_{j=1}^{m} A_{1j}, \qquad B_i \;:=\; \Big(\bigcap_{j=1}^{m} A_{ij}\Big) \;\setminus\; \bigcup_{r=1}^{i-1} B_r, \quad i = 2, \dots, m.$$

Each B_i is Borel because it is obtained from finitely many Borel sets using finite intersections, unions, and complements.

Claim 1: $\{B_i\}_{i=1}^m$ is a partition of Θ.

- They are disjoint by construction (B_i removes all previous B_r).
- They cover Θ because for each θ, the finite set $\{g \cdot \theta : \; g \in G\}$ has a (unique) $\preceq$-minimum point, so the set of indices achieving the minimum is nonempty; taking the *smallest* such index i ensures $\theta \in B_i$. Formally: define

  $$M(\theta) := \min_{\preceq}\{g_j \cdot \theta : \; j = 1, \dots, m\}.$$

  Let $I(\theta) := \{i : \; g_i \cdot \theta = M(\theta)\}$, which is nonempty. Let $i^\star := \min I(\theta)$. Then $g_{i^\star} \cdot \theta \preceq g_j \cdot \theta$ for all j, so $\theta \in \cap_j A_{i^\star j}$, and by minimality of $i^\star$, $\theta \notin \cup_{r<i^\star} B_r$, hence $\theta \in B_{i^\star}$.

Step 2 (Borel measurability of $\tilde{s}$). On B_i the element $g_i \cdot \theta$ is $\preceq$-minimal among $\{g \cdot \theta : \; g \in G\}$, hence equals $\tilde{s}(\theta)$. Therefore we have the pointwise representation

$$\tilde{s}(\theta) = \sum_{i=1}^{m} \mathbf{1}_{B_i}(\theta)\,(g_i \cdot \theta), \tag{339}$$

where the sum is "piecewise" (exactly one indicator is 1 at each θ). To prove measurability, it is enough to show that for any $C \in \mathcal{B}(\Theta)$, the preimage $\tilde{s}^{-1}(C)$ is Borel. Using (339),

$$\tilde{s}^{-1}(C) = \bigcup_{i=1}^{m} \Big(B_i \;\cap\; \{\theta : \; g_i \cdot \theta \in C\}\Big).$$

Each $\{\theta : \; g_i \cdot \theta \in C\} = T_{g_i}^{-1}(C)$ is Borel since T_{g_i} is measurable. Since B_i is Borel and $m < \infty$, the union is Borel. Hence $\tilde{s}$ is Borel measurable.

Step 3 (Orbit-constancy of $\tilde{s}$). Fix $h \in G$ and $\theta \in \Theta$. Because G is a group, left multiplication by h permutes G bijectively, so

$$\{g \cdot (h \cdot \theta) : \; g \in G\} = \{(gh) \cdot \theta : \; g \in G\} = \{g' \cdot \theta : \; g' \in G\}.$$

Taking $\preceq$-minima of equal sets yields

$$\tilde{s}(h \cdot \theta) = \min_{\preceq}\{g \cdot (h \cdot \theta) : g \in G\} = \min_{\preceq}\{g' \cdot \theta : g' \in G\} = \tilde{s}(\theta).$$

Thus $\tilde{s}$ is constant on orbits.

Step 4 (Define s on the quotient and prove it is well-defined). Let $q : \Theta \to \Theta/G$ be the orbit map $q(\theta) = [\theta]$. Define $s : \Theta/G \to \Theta$ by

$$s(\bar{\theta}) \;:=\; \tilde{s}(\theta) \quad \text{for any } \theta \in \Theta \text{ with } q(\theta) = \bar{\theta}. \tag{340}$$

This is well-defined: if $q(\theta_1) = q(\theta_2)$, then $\theta_2 = g \cdot \theta_1$ for some $g \in G$, and Step 3 gives $\tilde{s}(\theta_2) = \tilde{s}(g \cdot \theta_1) = \tilde{s}(\theta_1)$.

Step 5 (Measurability of s and factorization $\tilde{s} = s \circ q$). By construction, $\tilde{s}(\theta) = s(q(\theta))$, i.e. $\tilde{s} = s \circ q$. To prove s is measurable, let $C \in \mathcal{B}(\Theta)$. We show $s^{-1}(C) \in \mathcal{B}(\Theta/G)$ by verifying its pullback is Borel:

$$q^{-1}\big(s^{-1}(C)\big) = \{\theta \in \Theta : \ s(q(\theta)) \in C\} = \{\theta \in \Theta : \ \tilde{s}(\theta) \in C\} = \tilde{s}^{-1}(C) \ \in \ \mathcal{B}(\Theta),$$

since $\tilde{s}$ is measurable (Step 2). By definition of the quotient σ-field, this implies $s^{-1}(C) \in \mathcal{B}(\Theta/G)$. Thus s is measurable. Uniqueness of s follows from surjectivity of q: if $s_1 \circ q = s_2 \circ q$, then $s_1 = s_2$.

Step 6 (s is a section: $q \circ s = \mathrm{id}$). Fix $\bar{\theta} \in \Theta/G$ and choose θ with $q(\theta) = \bar{\theta}$. By definition of $\tilde{s}(\theta)$ as the $\preceq$-minimum in $\{g \cdot \theta : \ g \in G\}$, there exists some $g \in G$ such that $\tilde{s}(\theta) = g \cdot \theta$. Therefore

$$q(s(\bar{\theta})) = q(\tilde{s}(\theta)) = q(g \cdot \theta) = q(\theta) = \bar{\theta},$$

where we used that q is constant on orbits: $q(g \cdot \theta) = q(\theta)$. Hence $q \circ s = \mathrm{id}_{\Theta/G}$, i.e. s is a section. □

Remark F.6 (Canonical relabeling for mixtures and other finite-label models). Lemma F.5 yields a *canonical* orbit representative whenever the label group G is finite and acts Borelly: one fixes a Borel order $\preceq$ on Θ and selects the $\preceq$-minimum element of each finite orbit. This avoids ad hoc coordinate ordering rules and handles ties automatically (ties can only arise from stabilizers producing the same parameter point, in which case the minimum point is still well-defined).

F.2. **Over-specified mixtures and misspecified projection geometry.**

Lemma F.7 (Invariance of Φ and orbit structure of $\bar{\Theta}$). *Let $(\mathcal{X}, \mathcal{F})$ be the observation space and μ a σ-finite measure on $(\mathcal{X}, \mathcal{F})$. Assume $P^\star \ll \mu$. For each $\theta \in \Theta$, let $P_\theta^X \ll \mu$ and let $p_\theta := \frac{dP_\theta^X}{d\mu}$ be a chosen μ-density (i.e. a Radon–Nikodym derivative, defined μ-a.e.). Define the extended-real objective*

$$\Phi(\theta) \ := \ \mathbb{E}_{P^\star}\big[\log p_\theta(X)\big] \ \in [-\infty, \infty]$$

(where $\log 0 := -\infty$), and the maximizer set

$$\bar{\Theta} \ := \ \arg\max_{\theta \in \Theta} \Phi(\theta) \ = \ \{\theta \in \Theta : \ \Phi(\theta) = \sup_{\vartheta \in \Theta} \Phi(\vartheta)\}.$$

Let G act measurably on Θ and assume observed invariance *holds:*

$$P_{g\cdot\theta}^X \ = \ P_\theta^X \qquad \forall\, \theta \in \Theta, \ \forall\, g \in G.$$

Then:

(i) $\Phi(g \cdot \theta) = \Phi(\theta)$ for all $\theta \in \Theta$ and $g \in G$ (as extended real numbers);
(ii) $\bar{\Theta}$ is G-invariant: if $\theta \in \bar{\Theta}$ then $g \cdot \theta \in \bar{\Theta}$ for all $g \in G$;
(iii) the orbit-level set $q(\bar{\Theta}) \subseteq \Theta/G$ is the identifiable projection target; equivalently, Φ factors through the quotient and its maximizers in Θ/G are precisely $q(\bar{\Theta})$.

Proof. Fix $\theta \in \Theta$ and $g \in G$.

Step 1 (Equality of laws implies equality of densities μ-a.e.). By assumption, $P_{g\cdot\theta}^X = P_\theta^X$ as measures on $(\mathcal{X}, \mathcal{F})$. Since both measures are absolutely continuous w.r.t. μ, their Radon–Nikodym derivatives are unique μ-a.e. Hence

$$p_{g\cdot\theta}(x) \ = \ p_\theta(x) \qquad \text{for } \mu\text{-a.e. } x \in \mathcal{X}. \tag{341}$$

Step 2 (Transfer μ-a.e. equality to $P^\star$-a.s. equality). Because $P^\star \ll \mu$, any μ-null set is $P^\star$-null. Therefore (341) implies

$$p_{g\cdot\theta}(X) \;=\; p_\theta(X) \qquad P^\star\text{-a.s.} \tag{342}$$

Applying $\log(\cdot)$ (with $\log 0 := -\infty$) preserves equality pointwise on the event in (342), hence

$$\log p_{g\cdot\theta}(X) \;=\; \log p_\theta(X) \qquad P^\star\text{-a.s.} \tag{343}$$

Step 3 (Almost-sure equality implies equality of expectations, including the extended-real case). The random variables $\log p_{g\cdot\theta}(X)$ and $\log p_\theta(X)$ are measurable as maps into $[-\infty,\infty)$. Since they are equal $P^\star$-a.s. by (343), their integrals w.r.t. $P^\star$ coincide in the extended sense:

$$\mathbb{E}_{P^\star}\big[\log p_{g\cdot\theta}(X)\big] \;=\; \mathbb{E}_{P^\star}\big[\log p_\theta(X)\big].$$

This proves (i), i.e. $\Phi(g\cdot\theta) = \Phi(\theta)$.

Step 4 (G-invariance of the maximizer set). Let $\theta \in \bar\Theta$. By definition,

$$\Phi(\theta) \;=\; \sup_{\vartheta\in\Theta} \Phi(\vartheta).$$

Using (i),

$$\Phi(g\cdot\theta) \;=\; \Phi(\theta) \;=\; \sup_{\vartheta\in\Theta} \Phi(\vartheta),$$

so $g\cdot\theta \in \bar\Theta$. This proves (ii).

Step 5 (Orbit-level target and factorization through the quotient). Let $q : \Theta \to \Theta/G$ be the orbit map and endow Θ/G with the quotient σ-field $\mathcal{B}(\Theta/G) = \{A :\; q^{-1}(A) \in \mathcal{B}(\Theta)\}$. By (i), Φ is G-invariant, so by Lemma F.1 there exists a unique measurable $\bar\Phi : \Theta/G \to [-\infty,\infty)$ such that

$$\Phi \;=\; \bar\Phi \circ q.$$

We now identify the maximizers of $\bar\Phi$ with $q(\bar\Theta)$.

Claim 5a: $q(\bar\Theta) \subseteq \arg\max_{\bar\theta\in\Theta/G} \bar\Phi(\bar\theta)$. If $\theta \in \bar\Theta$, then for any $\bar\vartheta \in \Theta/G$ choose $\vartheta \in \Theta$ with $q(\vartheta) = \bar\vartheta$. Then

$$\bar\Phi(q(\theta)) = \Phi(\theta) \;\geq\; \Phi(\vartheta) = \bar\Phi(q(\vartheta)) = \bar\Phi(\bar\vartheta),$$

so $q(\theta)$ maximizes $\bar\Phi$.

Claim 5b: $\arg\max_{\bar\theta\in\Theta/G} \bar\Phi(\bar\theta) \subseteq q(\bar\Theta)$. Let $\bar\theta^\star$ maximize $\bar\Phi$, and pick $\theta^\star \in \Theta$ with $q(\theta^\star) = \bar\theta^\star$. For any $\vartheta \in \Theta$,

$$\Phi(\theta^\star) = \bar\Phi(\bar\theta^\star) \;\geq\; \bar\Phi(q(\vartheta)) = \Phi(\vartheta),$$

so $\theta^\star \in \bar\Theta$ and hence $\bar\theta^\star = q(\theta^\star) \in q(\bar\Theta)$.

Combining Claims 5a–5b yields

$$q(\bar\Theta) \;=\; \arg\max_{\bar\theta\in\Theta/G} \bar\Phi(\bar\theta).$$

This is precisely the statement that the identifiable target is the orbit-level maximizer set, proving (iii). □

Remark F.8 (Quotient projection set and additional nonidentifiabilities). In over-specified mixtures, $\bar\Theta$ may contain multiple *distinct* orbits and may have positive dimension even after quotienting. Lemma F.7 isolates the unavoidable symmetry via the quotient target $q(\bar\Theta)$. Any remaining nonidentifiabilities (beyond G) manifest as geometry of $q(\bar\Theta)$ itself (e.g. continua of orbit families), which motivates orbit-level metrics and transverse curvature conditions in the main text.

F.3. **Finite-label models beyond i.i.d. mixtures.** Many latent-variable models share the same finite-label symmetry as mixtures. The key point is: if the complete-data density is built from *label-indexed* parameter blocks, then simultaneous relabeling of parameters and latent labels induces latent transport.

Hidden Markov models (HMMs) with state-label switching. Fix $k \geq 2$ and horizon $T \geq 1$. Let $\mathcal{Z} = [k]^T$ with product σ-field and counting measure ν. Let $X_1, \ldots, X_T$ be observed in $\mathcal{X}^T$ with dominating measure $\mu^{\otimes T}$.
Parameterization. Let $\theta = (\pi, A, \vartheta_1, \ldots, \vartheta_k)$ where $\pi \in \Delta_{k-1}$ is the initial distribution, $A \in [0,1]^{k\times k}$ is row-stochastic, and $\vartheta_j \in \Xi$ index emission densities $f(\cdot;\vartheta_j)$ on $(\mathcal{X}, \mathcal{F}, \mu)$, as in Appendix F.1. For $z_{1:T} \in [k]^T$ and $x_{1:T} \in \mathcal{X}^T$, define the complete-data density

$$p_\theta(x_{1:T}, z_{1:T}) := \pi_{z_1} \prod_{t=2}^{T} A_{z_{t-1},z_t} \prod_{t=1}^{T} f(x_t;\vartheta_{z_t}) \qquad \text{w.r.t. } \mu^{\otimes T} \otimes \nu, \tag{344}$$

where ν is counting measure on $[k]^T$.
Permutation action. Let $G = S_k$ act by simultaneous relabeling:

$$(g\cdot\pi)_j := \pi_{g^{-1}(j)}, \qquad (g\cdot A)_{ij} := A_{g^{-1}(i),g^{-1}(j)}, \qquad (g\cdot\vartheta)_j := \vartheta_{g^{-1}(j)}.$$

Define $\tau_g : [k]^T \to [k]^T$ by $(\tau_g(z))_t := g(z_t)$.

Lemma F.9 (HMM transport and EM equivariance)**.** *For the HMM* (344) *with the S_k action above, Assumption C.1 holds with latent space $\mathcal{Z} = [k]^T$, base measure ν counting measure, and τ_g acting componentwise. Consequently:*

*(i) (*Observed invariance*) $P_{g\cdot\theta}^{X_{1:T}} = P_\theta^{X_{1:T}}$ for all $g \in S_k$ and $\theta \in \Theta$.*

*(ii) (*Posterior transport*) for all $B \subseteq [k]^T$ and $P_\theta^{X_{1:T}}$-a.e. $x_{1:T}$,*

$$\kappa_{g\cdot\theta}(x_{1:T}, B) \;=\; \kappa_\theta\big(x_{1:T}, \tau_g^{-1}(B)\big).$$

(iii) (Q-equivariance and M-step equivariance*) whenever the Q-functional is well-defined,*

$$Q(g\cdot\theta \mid g\cdot\theta') \;=\; Q(\theta \mid \theta') \qquad \forall\, \theta, \theta' \in \Theta,\ \forall\, g \in S_k,$$

and therefore the M-step correspondence $\mathcal{T}(\theta') \in \arg\max_\theta Q(\theta \mid \theta')$ satisfies $\mathcal{T}(g\cdot\theta') = g\cdot\mathcal{T}(\theta')$ (equivariance in the correspondence sense).

Proof. We verify the items in Assumption C.1 explicitly, and then derive (i)–(iii).

Step 1 (Measurability and group structure of τ_g). The latent space $\mathcal{Z} = [k]^T$ is finite, hence equipped with the discrete σ-field $2^{\mathcal{Z}}$. Therefore every map $\tau_g : \mathcal{Z} \to \mathcal{Z}$ is measurable. Moreover, τ_g is a bijection with inverse $\tau_{g^{-1}}$. Finally, for $g, h \in S_k$ and $z \in \mathcal{Z}$,

$$(\tau_{gh}(z))_t = (gh)(z_t) = g(h(z_t)) = (\tau_g(\tau_h(z)))_t,$$

so $\tau_{gh} = \tau_g \circ \tau_h$.

Step 2 (Base-measure preservation). Since $\mathcal{Z}$ is finite and ν is counting measure, any bijection τ_g preserves ν: for every $B \subseteq \mathcal{Z}$,

$$\nu(\tau_g(B)) = |\,\tau_g(B)\,| = |\,B\,| = \nu(B).$$

Equivalently, $\nu \circ \tau_g^{-1} = \nu$ for all g.

Step 3 (Equivariance of the complete-data density). Fix $\theta = (\pi, A, \vartheta)$, $g \in S_k$, and $(x_{1:T}, z_{1:T}) \in \mathcal{X}^T \times \mathcal{Z}$. Write $z' := \tau_g(z)$ so $z'_t = g(z_t)$.

(a) Initial factor. By definition of the action on π,

$$\pi^{(g\cdot\theta)}_{z'_1} = (g\cdot\pi)_{g(z_1)} = \pi_{g^{-1}(g(z_1))} = \pi_{z_1}.$$

(b) Transition factors. For each $t \geq 2$, using the action on A,

$$A^{(g\cdot\theta)}_{z'_{t-1},z'_t} = (g\cdot A)_{g(z_{t-1}),g(z_t)} = A_{g^{-1}(g(z_{t-1})),\, g^{-1}(g(z_t))} = A_{z_{t-1},z_t}.$$

(c) Emission factors. For each t, using the action on ϑ,

$$\vartheta^{(g\cdot\theta)}_{z'_t} = (g\cdot\vartheta)_{g(z_t)} = \vartheta_{g^{-1}(g(z_t))} = \vartheta_{z_t}.$$

Hence $f(x_t;\vartheta^{(g\cdot\theta)}_{z'_t}) = f(x_t;\vartheta_{z_t})$.

Multiplying (a)–(c) over the HMM factorization (344) yields

$$p_{g\cdot\theta}(x_{1:T},\tau_g(z_{1:T})) \;=\; p_\theta(x_{1:T},z_{1:T}) \qquad \text{for all } (x_{1:T},z_{1:T}). \tag{345}$$

Together with Steps 1–2, this is exactly the latent-transport property in Assumption C.1.

Step 4 (Observed invariance by marginalization and change of variables). Let $P^{X,Z}_\theta$ denote the joint law on $\mathcal{X}^T\times\mathcal{Z}$ with density (344) w.r.t. $\mu^{\otimes T}\otimes\nu$, and let P^X_θ be its marginal on $\mathcal{X}^T$. For any measurable $A\in\mathcal{F}^{\otimes T}$,

$$\begin{aligned}
P^X_{g\cdot\theta}(A) &= \sum_{z\in\mathcal{Z}}\int_{\mathcal{X}^T} \mathbf{1}_A(x)\, p_{g\cdot\theta}(x,z)\,\mu^{\otimes T}(dx)\\
&= \sum_{z\in\mathcal{Z}}\int_{\mathcal{X}^T} \mathbf{1}_A(x)\, p_{g\cdot\theta}\big(x,\tau_g(z)\big)\,\mu^{\otimes T}(dx) \qquad \text{(reindex the finite sum since } \tau_g \text{ is a bijection)}\\
&= \sum_{z\in\mathcal{Z}}\int_{\mathcal{X}^T} \mathbf{1}_A(x)\, p_\theta(x,z)\,\mu^{\otimes T}(dx) \qquad \text{by (345)}\\
&= P^X_\theta(A).
\end{aligned}$$

Thus $P^{X_{1:T}}_{g\cdot\theta} = P^{X_{1:T}}_\theta$, proving (i).

Step 5 (Posterior transport). Fix $B\subseteq\mathcal{Z}$. Because $\mathcal{Z}$ is finite, the posterior $\kappa_\theta(x,\cdot)$ can be written explicitly by Bayes' rule whenever the marginal density $p^X_\theta(x) := \sum_{z\in\mathcal{Z}} p_\theta(x,z)$ is positive:

$$\kappa_\theta(x,B) \;=\; \frac{\sum_{z\in B} p_\theta(x,z)}{\sum_{z\in\mathcal{Z}} p_\theta(x,z)}.$$

We show the claimed identity on the set $\{x:\ p^X_\theta(x)>0\}$; outside it, the posterior can be defined arbitrarily.

Using (345) and the bijection τ_g,

$$\sum_{z\in B} p_{g\cdot\theta}(x,z) = \sum_{z\in B} p_{g\cdot\theta}\big(x,\tau_g(\tau_g^{-1}(z))\big) = \sum_{z'\in\tau_g^{-1}(B)} p_{g\cdot\theta}\big(x,\tau_g(z')\big) = \sum_{z'\in\tau_g^{-1}(B)} p_\theta(x,z').$$

Similarly,

$$\sum_{z\in\mathcal{Z}} p_{g\cdot\theta}(x,z) = \sum_{z'\in\mathcal{Z}} p_\theta(x,z') = p^X_\theta(x),$$

where we used $\tau_g^{-1}(\mathcal{Z}) = \mathcal{Z}$. Therefore, whenever $p_\theta^X(x) > 0$ (equivalently, $p_{g\cdot\theta}^X(x) > 0$ by Step 4),

$$\kappa_{g\cdot\theta}(x, B) = \frac{\sum_{z\in B} p_{g\cdot\theta}(x, z)}{\sum_{z\in\mathcal{Z}} p_{g\cdot\theta}(x, z)} = \frac{\sum_{z'\in\tau_g^{-1}(B)} p_\theta(x, z')}{\sum_{z'\in\mathcal{Z}} p_\theta(x, z')} = \kappa_\theta\big(x, \tau_g^{-1}(B)\big).$$

This proves (ii) (for P_θ^X-a.e. x; the exceptional set is $\{p_\theta^X = 0\}$).

Step 6 (Q-equivariance and M-step equivariance). Assume $Q(\cdot \mid \cdot)$ is defined (i.e. the relevant integrability conditions hold). Fix $\theta, \theta' \in \Theta$ and $g \in S_k$.

By definition of the EM Q-functional (population version),

$$Q(\theta \mid \theta') = \mathbb{E}_{P^\star}\Big[\, \mathbb{E}_{Z\sim\kappa_{\theta'}(X,\cdot)}\big[\log p_\theta(X, Z)\big]\,\Big].$$

Apply this with (θ, θ') replaced by $(g\cdot\theta, g\cdot\theta')$:

$$Q(g\cdot\theta \mid g\cdot\theta') = \mathbb{E}_{P^\star}\Big[\, \mathbb{E}_{Z\sim\kappa_{g\cdot\theta'}(X,\cdot)}\big[\log p_{g\cdot\theta}(X, Z)\big]\,\Big].$$

Now use posterior transport (Step 5): if $Z \sim \kappa_{g\cdot\theta'}(X, \cdot)$, then $Z' := \tau_g^{-1}(Z)$ satisfies $Z' \sim \kappa_{\theta'}(X, \cdot)$. Moreover, by (345),

$$\log p_{g\cdot\theta}(X, Z) = \log p_{g\cdot\theta}\big(X, \tau_g(Z')\big) = \log p_\theta(X, Z').$$

Therefore, by a change of variables in the finite sum defining the conditional expectation,

$$\mathbb{E}_{Z\sim\kappa_{g\cdot\theta'}(X,\cdot)}\big[\log p_{g\cdot\theta}(X, Z)\big] = \mathbb{E}_{Z'\sim\kappa_{\theta'}(X,\cdot)}\big[\log p_\theta(X, Z')\big].$$

Taking $\mathbb{E}_{P^\star}$ of both sides yields $Q(g\cdot\theta \mid g\cdot\theta') = Q(\theta \mid \theta')$, proving the first part of (iii).

For the M-step correspondence, let $\mathcal{T}(\theta') \in \arg\max_\theta Q(\theta \mid \theta')$. Using $Q(g\cdot\theta \mid g\cdot\theta') = Q(\theta \mid \theta')$, we have for any θ,

$$Q(g\cdot\theta \mid g\cdot\theta') = Q(\theta \mid \theta') \le Q(\mathcal{T}(\theta') \mid \theta') = Q(g\cdot\mathcal{T}(\theta') \mid g\cdot\theta'),$$

so $g\cdot\mathcal{T}(\theta')$ is a maximizer of $\theta \mapsto Q(\theta \mid g\cdot\theta')$. This is exactly the equivariance statement $\mathcal{T}(g\cdot\theta') = g\cdot\mathcal{T}(\theta')$ in correspondence form. □

Measurable section. If Θ is a standard Borel space (e.g. a Borel subset of Euclidean space) and the S_k-action is Borel, then Lemma F.5 yields a global Borel selector $s : \Theta/S_k \to \Theta$ that chooses a canonical representative in each orbit (e.g. the $\preceq$-minimum element of the finite orbit).

Stochastic block models (SBMs) and latent-class network models. Let $n \ge 2$ and $k \ge 2$. Let $X = (X_{ij})_{1\le i<j\le n}$ be an observed adjacency matrix, viewed as an element of a finite measurable space $(\mathcal{X}, \mathcal{F})$ (e.g. $\mathcal{X} = \{0,1\}^{\binom{n}{2}}$ with $\mathcal{F} = 2^{\mathcal{X}}$) dominated by counting measure μ. Let latent labels be $Z = (Z_1, \dots, Z_n) \in [k]^n$ with counting measure ν on $\mathcal{Z} := [k]^n$.

Parameterization. Let $\theta = (\pi, B)$ where $\pi \in \Delta_{k-1}$ are class proportions and $B \in [0,1]^{k\times k}$ is symmetric. Given $Z = z$, edges are conditionally independent with

$$\Pr_\theta(X_{ij} = 1 \mid Z = z) = B_{z_i, z_j} \qquad (1 \le i < j \le n),$$

and $Z_i \overset{iid}{\sim} \pi$. The complete-data density (w.r.t. $\mu \otimes \nu$) is

$$p_\theta(x, z) = \Big(\prod_{i=1}^n \pi_{z_i}\Big) \cdot \Big(\prod_{1\le i<j\le n} B_{z_i,z_j}^{x_{ij}} (1 - B_{z_i,z_j})^{1-x_{ij}}\Big), \qquad (x, z) \in \mathcal{X}\times\mathcal{Z}. \tag{346}$$

Permutation action and transport. Let $G = S_k$ act on parameters by simultaneous relabeling:

$$(g \cdot \pi)_a := \pi_{g^{-1}(a)}, \qquad (g \cdot B)_{ab} := B_{g^{-1}(a),g^{-1}(b)}.$$

Let $\tau_g : \mathcal{Z} \to \mathcal{Z}$ act componentwise: $(\tau_g(z))_i := g(z_i)$.

Lemma F.10 (SBM transport and invariance)**.** *For the SBM (346) with the S_k action above, Assumption C.1 holds with latent space $\mathcal{Z} = [k]^n$, base measure ν counting measure, and τ_g acting componentwise. Consequently:*

*(i) (*Observed invariance*) $P^X_{g\cdot\theta} = P^X_\theta$ for all $g \in S_k$ and $\theta \in \Theta$.*

*(ii) (*Posterior transport*) for all $B \subseteq [k]^n$ and P^X_θ-a.e. x,*

$$\kappa_{g\cdot\theta}(x, B) \;=\; \kappa_\theta\big(x, \tau_g^{-1}(B)\big).$$

(iii) (Q-equivariance and M-step equivariance*) whenever $Q(\cdot \mid \cdot)$ is well-defined,*

$$Q(g \cdot \theta \mid g \cdot \theta') \;=\; Q(\theta \mid \theta') \qquad \forall\, \theta, \theta' \in \Theta,\ \forall\, g \in S_k,$$

and hence the M-step correspondence is S_k-equivariant.

Proof. We verify the latent-transport conditions and then derive (i)–(iii).

Step 1 (Group properties and measurability of the latent relabeling). The latent space $\mathcal{Z} = [k]^n$ is finite, hence carries the discrete σ-field $2^{\mathcal{Z}}$; thus every map τ_g is measurable. For $g, h \in S_k$ and $z \in \mathcal{Z}$,

$$(\tau_{gh}(z))_i = (gh)(z_i) = g(h(z_i)) = (\tau_g(\tau_h(z)))_i,$$

so $\tau_{gh} = \tau_g \circ \tau_h$. Also τ_g is a bijection with inverse $\tau_{g^{-1}}$.

Step 2 (Preservation of counting measure). Let ν be counting measure on $\mathcal{Z}$. Since τ_g is a bijection on a finite set, it preserves cardinality: for every $A \subseteq \mathcal{Z}$,

$$\nu(\tau_g(A)) = |\tau_g(A)| = |A| = \nu(A).$$

Equivalently, $\nu \circ \tau_g^{-1} = \nu$.

Step 3 (Equivariance of the complete-data density). Fix $\theta = (\pi, B)$, $g \in S_k$, and $(x, z) \in \mathcal{X} \times \mathcal{Z}$. Write $z' := \tau_g(z)$ so $z'_i = g(z_i)$.

(a) Class-proportion factor. Using the action on π,

$$\prod_{i=1}^n (g \cdot \pi)_{z'_i} = \prod_{i=1}^n \pi_{g^{-1}(z'_i)} = \prod_{i=1}^n \pi_{g^{-1}(g(z_i))} = \prod_{i=1}^n \pi_{z_i}.$$

(b) Edge-likelihood factor. For each $1 \le i < j \le n$,

$$(g \cdot B)_{z'_i,z'_j} = (g \cdot B)_{g(z_i),g(z_j)} = B_{g^{-1}(g(z_i)),\, g^{-1}(g(z_j))} = B_{z_i,z_j}.$$

Therefore

$$\prod_{i<j} (g \cdot B)^{x_{ij}}_{z'_i,z'_j} \big(1 - (g \cdot B)_{z'_i,z'_j}\big)^{1-x_{ij}} = \prod_{i<j} B^{x_{ij}}_{z_i,z_j} (1 - B_{z_i,z_j})^{1-x_{ij}}.$$

Multiplying (a) and (b) gives the transport identity

$$p_{g\cdot\theta}\big(x, \tau_g(z)\big) \;=\; p_\theta(x, z) \qquad \forall (x, z) \in \mathcal{X} \times \mathcal{Z}. \tag{347}$$

Together with Steps 1–2, this is precisely Assumption C.1.

Step 4 (Observed invariance by marginalization and reindexing). Let $P_\theta^{X,Z}$ be the joint law with density p_θ w.r.t. $\mu \otimes \nu$ and P_θ^X its marginal. For any measurable $A \subseteq \mathcal{X}$,

$$\begin{aligned} P_{g\cdot\theta}^X(A) &= \sum_{z\in\mathcal{Z}} \sum_{x\in A} p_{g\cdot\theta}(x,z) = \sum_{z\in\mathcal{Z}} \sum_{x\in A} p_{g\cdot\theta}\big(x, \tau_g(z)\big) \qquad \text{(reindex using the bijection } \tau_g\text{)} \\ &= \sum_{z\in\mathcal{Z}} \sum_{x\in A} p_\theta(x,z) = P_\theta^X(A), \end{aligned}$$

where we used (347). Hence $P_{g\cdot\theta}^X = P_\theta^X$, proving (i).

Step 5 (Posterior transport). Fix $B \subseteq \mathcal{Z}$. Since $\mathcal{Z}$ is finite, Bayes' rule yields, on the set where the marginal likelihood $p_\theta^X(x) := \sum_z p_\theta(x,z)$ is positive,

$$\kappa_\theta(x,B) = \frac{\sum_{z\in B} p_\theta(x,z)}{\sum_{z\in\mathcal{Z}} p_\theta(x,z)}.$$

Using (347) and the bijection τ_g,

$$\sum_{z\in B} p_{g\cdot\theta}(x,z) = \sum_{z\in B} p_{g\cdot\theta}\big(x, \tau_g(\tau_g^{-1}(z))\big) = \sum_{z'\in\tau_g^{-1}(B)} p_{g\cdot\theta}\big(x,\tau_g(z')\big) = \sum_{z'\in\tau_g^{-1}(B)} p_\theta(x,z').$$

Similarly, $\sum_{z\in\mathcal{Z}} p_{g\cdot\theta}(x,z) = \sum_{z\in\mathcal{Z}} p_\theta(x,z)$. Thus, whenever the denominators are nonzero,

$$\kappa_{g\cdot\theta}(x,B) = \kappa_\theta\big(x, \tau_g^{-1}(B)\big).$$

Since $P_{g\cdot\theta}^X = P_\theta^X$, this holds P_θ^X-a.e., proving (ii).

Step 6 (Q-equivariance and M-step equivariance). Assume $Q(\cdot \mid \cdot)$ is well-defined (integrability). By definition,

$$Q(\theta \mid \theta') = \mathbb{E}_{P^\star}\Big[\mathbb{E}_{Z\sim\kappa_{\theta'}(X,\cdot)}\big[\log p_\theta(X,Z)\big]\Big].$$

Apply this with (θ,θ') replaced by $(g\cdot\theta, g\cdot\theta')$ and use Step 5: if $Z \sim \kappa_{g\cdot\theta'}(X,\cdot)$ then $Z' := \tau_g^{-1}(Z) \sim \kappa_{\theta'}(X,\cdot)$. Then (347) gives $\log p_{g\cdot\theta}(X,Z) = \log p_\theta(X,Z')$. Hence, by a change of variables over the finite latent space,

$$\mathbb{E}_{Z\sim\kappa_{g\cdot\theta'}(X,\cdot)}[\log p_{g\cdot\theta}(X,Z)] = \mathbb{E}_{Z'\sim\kappa_{\theta'}(X,\cdot)}[\log p_\theta(X,Z')],$$

and taking $\mathbb{E}_{P^\star}$ yields $Q(g\cdot\theta \mid g\cdot\theta') = Q(\theta \mid \theta')$. The usual maximizer-transfer argument then implies equivariance of the M-step correspondence, proving (iii). □

F.4. **Signed-permutation symmetries (hyperoctahedral groups).** Finite sign and permutation symmetries arise whenever the latent prior is invariant under coordinate relabeling and sign flips. This includes many ICA/dictionary-learning style latent models (where identifiability is only up to signed permutation).

Signed permutation group. Let $r \geq 1$ and define the finite group

$$\mathsf{SP}(r) := \Big\{ R \in \mathbb{R}^{r\times r} : \ R \text{ has exactly one nonzero entry in each row/column, equal to } \pm 1 \Big\}.$$

Equivalently,

$$\mathsf{SP}(r) = \{ D\Pi : \ D \in \{\pm 1\}^{r\times r} \text{ diagonal}, \ \Pi \in S_r \text{ permutation matrix}\}.$$

Every $R \in \mathsf{SP}(r)$ is orthogonal: $R^\top R = I_r$ and $R^{-1} = R^\top$.

ICA / linear latent-factor template with i.i.d. symmetric priors. Let $\mathcal{Z} = \mathbb{R}^r$ with its Borel σ-field and let $\nu = \nu_0^{\otimes r}$, where ν_0 is a probability measure on $\mathbb{R}$ that is symmetric about 0, i.e. $\nu_0(B) = \nu_0(-B)$ for all Borel $B \subseteq \mathbb{R}$ (equivalently, $U \sim \nu_0 \Rightarrow -U \overset{d}{=} U$). Then ν is invariant under signed permutations in the following strong (measure-pushforward) sense.

Lemma F.11 (Invariance of the product prior under $\mathsf{SP}(r)$)**.** *Let $\nu = \nu_0^{\otimes r}$ with ν_0 symmetric. For each $R \in \mathsf{SP}(r)$ define $\tau_R : \mathbb{R}^r \to \mathbb{R}^r$ by $\tau_R(z) = Rz$. Then τ_R is a Borel bijection and*

$$\nu \circ \tau_R^{-1} = \nu \qquad \text{for all } R \in \mathsf{SP}(r).$$

Proof. **Step 1 (Borel bijection).** Since τ_R is linear and invertible, it is continuous (hence Borel measurable) and bijective with inverse $\tau_{R^{-1}}$.

Step 2 (Reduction to generators). It suffices to verify $\nu(\tau_R^{-1}(C)) = \nu(C)$ for all cylinder rectangles $C = \prod_{j=1}^r C_j$ with Borel $C_j \subseteq \mathbb{R}$, because such rectangles form a π-system generating the Borel σ-field on $\mathbb{R}^r$, and both ν and $\nu \circ \tau_R^{-1}$ are probability measures.

Step 3 (Explicit form of τ_R). Write $R = D\Pi$ with Π a permutation matrix and $D = \mathrm{diag}(\sigma_1, \ldots, \sigma_r)$ with $\sigma_j \in \{\pm 1\}$. Then $(\tau_R(z))_j = \sigma_j z_{\pi^{-1}(j)}$, i.e. τ_R permutes coordinates and flips signs of selected coordinates.

Step 4 (Compute $\nu(\tau_R^{-1}(C))$). We have

$$\tau_R^{-1}(C) = \Big\{z : \ \sigma_j z_{\pi^{-1}(j)} \in C_j \ \forall j\Big\} = \Big\{z : \ z_i \in \sigma_{\pi(i)}^{-1} C_{\pi(i)} \ \forall i\Big\},$$

where $\sigma^{-1}C := \{\sigma^{-1}x : x \in C\}$ equals C if $\sigma = 1$ and $-C$ if $\sigma = -1$. By product structure of ν,

$$\nu(\tau_R^{-1}(C)) = \prod_{i=1}^r \nu_0\big(\sigma_{\pi(i)}^{-1} C_{\pi(i)}\big).$$

By symmetry of ν_0, $\nu_0(-B) = \nu_0(B)$ for all B, hence $\nu_0(\sigma^{-1}B) = \nu_0(B)$ for $\sigma \in \{\pm 1\}$. Therefore

$$\nu(\tau_R^{-1}(C)) = \prod_{i=1}^r \nu_0(C_{\pi(i)}) = \prod_{j=1}^r \nu_0(C_j) = \nu(C),$$

since π is a permutation. This proves $\nu \circ \tau_R^{-1} = \nu$. □

Model. Let Θ contain a mixing matrix $A \in \mathbb{R}^{d \times r}$ and possibly nuisance parameters η (e.g. noise scale/covariance). Assume the conditional law of X given $Z = z$ depends on A only through Az; concretely,

$$X \mid Z = z \sim p_{\mathrm{noise}}(\,\cdot\,; Az, \eta),$$

where for each (m, η) the map $x \mapsto p_{\mathrm{noise}}(x; m, \eta)$ is a μ-density on $(\mathcal{X}, \mathcal{F})$, and $(x, z, A, \eta) \mapsto p_{\mathrm{noise}}(x; Az, \eta)$ is jointly measurable. With ν as above, define the complete-data density w.r.t. $\mu \otimes \nu$ by

$$p_{(A,\eta)}(x, z) = p_{\mathrm{noise}}(x; Az, \eta).$$

Action and transport. Let $G = \mathsf{SP}(r)$ act on Θ by *right* transformation on A:

$$R \cdot (A, \eta) := (AR^{-1}, \eta), \qquad R \in \mathsf{SP}(r),$$

and define latent transport $\tau_R(z) := Rz$ on $\mathcal{Z} = \mathbb{R}^r$.

Lemma F.12 (Signed-permutation transport and EM equivariance)**.** *In the ICA-style template above, Assumption C.1 holds with $G = \mathsf{SP}(r)$, $\tau_R(z) = Rz$, and prior $\nu = \nu_0^{\otimes r}$. Consequently:*

*(i) (*Observed invariance*) $P^X_{R\cdot(A,\eta)} = P^X_{(A,\eta)}$ for all $R \in \mathsf{SP}(r)$;*

*(ii) (*Posterior transport*) for all Borel* $B \subseteq \mathbb{R}^r$ *and* $P^X_{(A,\eta)}$*-a.e.* x*,*

$$\kappa_{R\cdot(A,\eta)}(x, B) = \kappa_{(A,\eta)}\big(x, \tau_R^{-1}(B)\big);$$

(iii) (Q-equivariance and M-step equivariance*) whenever* Q *is well-defined,*

$$Q(R\cdot\theta \mid R\cdot\theta') = Q(\theta \mid \theta') \qquad \forall\theta, \theta' \in \Theta,\ \forall R \in \mathsf{SP}(r),$$

and therefore the M-step correspondence $\mathcal{T}$ *satisfies* $\mathcal{T}(R\cdot\theta') = R\cdot\mathcal{T}(\theta')$*.*

Proof. Write $\theta = (A, \eta)$.

Step 1 (Group action and measurability). $\mathsf{SP}(r)$ is a finite group under matrix multiplication. The action $\theta \mapsto R\cdot\theta = (AR^{-1}, \eta)$ is continuous in A for fixed R, hence Borel measurable on any Borel parameter space $\Theta \subseteq \mathbb{R}^{d\times r} \times \mathsf{H}$ (with H Euclidean for η). The latent transport $\tau_R(z) = Rz$ is continuous and bijective, hence Borel measurable with measurable inverse.

Step 2 (Base-measure invariance on the latent space). By Lemma F.11, $\nu \circ \tau_R^{-1} = \nu$ for all $R \in \mathsf{SP}(r)$.

Step 3 (Complete-data equivariance identity). For every $(x, z) \in \mathcal{X} \times \mathbb{R}^r$,

$$p_{R\cdot\theta}\big(x, \tau_R(z)\big) = p_{(AR^{-1},\eta)}(x, Rz) = p_{\text{noise}}(x; AR^{-1}(Rz), \eta) = p_{\text{noise}}(x; Az, \eta) = p_\theta(x, z).$$

This is exactly the density-transport identity required in Assumption C.1.

Step 4 (Observed invariance by marginalization and change of variables). Let p^X_θ be the μ-density of P^X_θ, given by integrating out z against ν:

$$p^X_\theta(x) = \int_{\mathbb{R}^r} p_\theta(x, z)\,\nu(dz) = \int_{\mathbb{R}^r} p_{\text{noise}}(x; Az, \eta)\,\nu(dz).$$

Then, using Step 3 and Step 2,

$$\begin{aligned} p^X_{R\cdot\theta}(x) &= \int_{\mathbb{R}^r} p_{R\cdot\theta}(x, z)\,\nu(dz) = \int_{\mathbb{R}^r} p_{R\cdot\theta}\big(x, \tau_R(z')\big)\,\nu(d(\tau_R(z'))) \\ &= \int_{\mathbb{R}^r} p_{R\cdot\theta}\big(x, \tau_R(z')\big)\,\nu(dz') \qquad (\text{since } \nu\circ\tau_R^{-1} = \nu) \\ &= \int_{\mathbb{R}^r} p_\theta(x, z')\,\nu(dz') = p^X_\theta(x), \end{aligned}$$

so $P^X_{R\cdot\theta} = P^X_\theta$, proving (i).

Step 5 (Posterior transport). Whenever $p^X_\theta(x) > 0$, Bayes' formula yields the posterior as a probability kernel

$$\kappa_\theta(x, B) = \frac{\int_B p_\theta(x, z)\,\nu(dz)}{\int_{\mathbb{R}^r} p_\theta(x, z)\,\nu(dz)}.$$

Using Step 3 and the change of variables $z = \tau_R(z')$ plus Step 2,

$$\int_B p_{R\cdot\theta}(x, z)\,\nu(dz) = \int_{\tau_R^{-1}(B)} p_{R\cdot\theta}\big(x, \tau_R(z')\big)\,\nu(dz') = \int_{\tau_R^{-1}(B)} p_\theta(x, z')\,\nu(dz'),$$

and similarly $\int p_{R\cdot\theta}(x, z)\nu(dz) = \int p_\theta(x, z)\nu(dz)$. Therefore $\kappa_{R\cdot\theta}(x, B) = \kappa_\theta(x, \tau_R^{-1}(B))$ for all B and for all x with $p^X_\theta(x) > 0$. Since $P^X_{R\cdot\theta} = P^X_\theta$ (Step 4), this holds P^X_θ-a.e., proving (ii).

Step 6 (Q-equivariance and M-step equivariance). Assume $Q(\cdot \mid \cdot)$ is well-defined (integrability of $\log p_\theta(X,Z)$ under the relevant posteriors). By definition,

$$Q(\theta \mid \theta') = \mathbb{E}_{P^\star}\Big[\int_{\mathbb{R}^r} \log p_\theta(X,z)\,\kappa_{\theta'}(X,dz)\Big].$$

Apply this with (θ,θ') replaced by $(R\cdot\theta, R\cdot\theta')$ and use Steps 3 and 5 with the change of variables $z = \tau_R(z')$:

$$\begin{aligned}\int \log p_{R\cdot\theta}(X,z)\,\kappa_{R\cdot\theta'}(X,dz) &= \int \log p_{R\cdot\theta}\big(X,\tau_R(z')\big)\,\kappa_{\theta'}(X,dz') \\ &= \int \log p_\theta(X,z')\,\kappa_{\theta'}(X,dz').\end{aligned}$$

Taking $\mathbb{E}_{P^\star}$ yields $Q(R\cdot\theta \mid R\cdot\theta') = Q(\theta \mid \theta')$. Equivariance of the M-step correspondence $\mathcal{T}(\theta') \in \arg\max_\theta Q(\theta \mid \theta')$ follows immediately: if $\theta \in \mathcal{T}(\theta')$, then for all ϑ,

$$Q(R\cdot\theta \mid R\cdot\theta') = Q(\theta \mid \theta') \geq Q(\vartheta \mid \theta') = Q(R\cdot\vartheta \mid R\cdot\theta'),$$

so $R\cdot\theta \in \mathcal{T}(R\cdot\theta')$. □

Measurable section. Since $\mathsf{SP}(r)$ is finite, Lemma F.5 yields a global Borel section on $\Theta/\mathsf{SP}(r)$ whenever Assumption F.4 holds.

F.5. **Rotational symmetries in linear-Gaussian factor models ($G = O(r)$).** This is the continuous-group analogue of label switching. It is canonical in factor analysis and probabilistic PCA, where A is only identified up to right orthogonal transformations.
Model and domination. Let $\mathcal{Z} = \mathbb{R}^r$ with its Borel σ-field and let $\nu = \mathcal{N}(0, I_r)$ be the standard Gaussian *measure* on $\mathbb{R}^r$. Let μ be Lebesgue measure on $\mathbb{R}^d$, fix $d \geq r$, and let $\Psi \in \mathbb{R}^{d\times d}$ be positive definite. Under P_A, sample $Z \sim \nu$ and then

$$X \mid Z = z \sim \mathcal{N}(Az, \Psi).$$

Write $\varphi_d(\cdot; m, \Psi)$ for the μ-density of $\mathcal{N}(m,\Psi)$. Since we take $\mu\otimes\nu$ as the dominating measure for (X,Z), the joint density is simply the conditional density:

$$p_A(x,z) \;:=\; \frac{dP_A^{X,Z}}{d(\mu\otimes\nu)}(x,z) \;=\; \varphi_d(x; Az, \Psi).$$

Action and transport. Let $G = O(r)$ act by right multiplication $R\cdot A := AR$ and define $\tau_R(z) := R^\top z$.

Lemma F.13 (Latent transport and observed invariance in the factor model)**.** *In the factor model above, Assumption C.1 holds with $G = O(r)$, $\tau_R(z) = R^\top z$, and $\nu = \mathcal{N}(0, I_r)$. Consequently, $P_{AR}^X = P_A^X$ for all $R \in O(r)$; posteriors transport by τ_R; and Q (hence the M-step) is $O(r)$-equivariant whenever defined.*

Proof. **Step 1 (Group action, transport map, and measurability).** The map $(A,R) \mapsto AR$ is continuous, hence the action $A \mapsto AR$ is Borel measurable on $\mathbb{R}^{d\times r}$ for each fixed R. For each $R \in O(r)$, $\tau_R(z) = R^\top z$ is linear and invertible, hence a Borel bijection with inverse $\tau_{R^\top}$.

Step 2 (Gaussian base-measure invariance). Let $Z \sim \mathcal{N}(0, I_r)$. For any $R \in O(r)$, $R^\top Z \overset{d}{=} Z$. Equivalently, $\nu \circ \tau_R^{-1} = \nu$. A quick verification: the characteristic function satisfies

$$\mathbb{E}[e^{i\langle t, R^\top Z\rangle}] = \mathbb{E}[e^{i\langle Rt, Z\rangle}] = \exp\Big(-\tfrac{1}{2}\|Rt\|^2\Big) = \exp\Big(-\tfrac{1}{2}\|t\|^2\Big),$$

using R orthogonal.

Step 3 (Complete-data equivariance identity). For all $(x, z) \in \mathbb{R}^d \times \mathbb{R}^r$ and $R \in O(r)$,

$$p_{AR}\big(x, \tau_R(z)\big) = \varphi_d\big(x; AR(R^\top z), \Psi\big) = \varphi_d(x; Az, \Psi) = p_A(x, z).$$

Thus the density-transport identity holds.

Step 4 (Observed invariance by marginalization). The observed density is

$$p_A^X(x) = \int_{\mathbb{R}^r} \varphi_d(x; Az, \Psi)\, \nu(dz).$$

Using Step 3 with $z = \tau_R(z')$ and Step 2,

$$p_{AR}^X(x) = \int \varphi_d(x; ARz, \Psi)\, \nu(dz) = \int \varphi_d\big(x; AR(\tau_{R^\top}(z')), \Psi\big)\, \nu(dz') = \int \varphi_d(x; Az', \Psi)\, \nu(dz') = p_A^X(x),$$

so $P_{AR}^X = P_A^X$.

Step 5 (Posterior transport). For $p_A^X(x) > 0$,

$$\kappa_A(x, B) = \frac{\int_B p_A(x, z)\, \nu(dz)}{\int p_A(x, z)\, \nu(dz)}.$$

Using Step 3 and change of variables $z = \tau_R(z')$ with Step 2,

$$\int_B p_{AR}(x, z)\, \nu(dz) = \int_{\tau_R^{-1}(B)} p_{AR}\big(x, \tau_R(z')\big)\, \nu(dz') = \int_{\tau_R^{-1}(B)} p_A(x, z')\, \nu(dz').$$

The denominator is invariant by the same argument, so $\kappa_{AR}(x, B) = \kappa_A(x, \tau_R^{-1}(B))$.

Step 6 (Q-equivariance). As in Lemma F.12, combine Step 3 (transport for $\log p$) and Step 5 (posterior transport) inside the defining conditional expectation for Q. This yields $Q(AR_1 \mid AR_2) = Q(A \mid A)$ with the appropriate pairing (more generally $Q(R \cdot A \mid R \cdot A') = Q(A \mid A')$), whenever both sides are finite. Equivariance of the M-step correspondence follows. □

Local measurable sections (minor charts and polar normalization). Global measurable sections for $A \mapsto AR$ generally do not exist on all of $\mathbb{R}^{d\times r}$, but local sections on an open cover of the full-column-rank set suffice for slice-based arguments.
Minor charts. Fix $I \subseteq \{1, \dots, d\}$ with $|I| = r$ and let A_I be the $r \times r$ submatrix formed by rows indexed by I. Define the open set

$$\Omega_I := \{A \in \mathbb{R}^{d\times r} : \ \det(A_I) \neq 0\}.$$

Every full-column-rank A lies in at least one Ω_I.
Polar factors and normalization. For invertible $B \in \mathbb{R}^{r\times r}$ define its polar decomposition $B = U(B)H(B)$ with

$$H(B) := (B^\top B)^{1/2} \in \mathbb{S}_{++}^r, \qquad U(B) := B\, H(B)^{-1} \in O(r).$$

On Ω_I define

$$R_I(A) := U(A_I)^\top \in O(r), \qquad \tilde{s}_I(A) := A\, R_I(A) \in \mathbb{R}^{d\times r}.$$

Lemma F.14 (Local section for the right $O(r)$ action)**.** *Fix I and work on Ω_I. Then:*

(i) R_I and $\tilde{s}_I$ are continuous (hence Borel) on Ω_I;
(ii) $\tilde{s}_I$ is constant on right-orbits in Ω_I: $\tilde{s}_I(AR) = \tilde{s}_I(A)$ for all $R \in O(r)$;
(iii) $(\tilde{s}_I(A))_I \in \mathbb{S}_{++}^r$ (slice normalization);
(iv) consequently, $\tilde{s}_I$ factors through the local quotient $\Omega_I/O(r)$ to a measurable section.

Proof. **Step 1 (Continuity of the polar maps on $\mathrm{GL}(r)$).** On $\mathrm{GL}(r)$, the map $B \mapsto B^\top B$ is polynomial and hence continuous, and its image lies in $\mathbb{S}^r_{++}$. The principal square-root map $M \mapsto M^{1/2}$ is continuous on $\mathbb{S}^r_{++}$ (indeed smooth), and so is inversion $M \mapsto M^{-1}$. Therefore $H(B) = (B^\top B)^{1/2}$ and $H(B)^{-1}$ depend continuously on $B \in \mathrm{GL}(r)$. Consequently $U(B) = B\,H(B)^{-1}$ is continuous on $\mathrm{GL}(r)$.

Step 2 (Proof of (i)). On Ω_I, the map $A \mapsto A_I$ is linear (hence continuous). Since $A_I \in \mathrm{GL}(r)$ on Ω_I, Step 1 implies $A \mapsto U(A_I)$ is continuous on Ω_I. Thus $R_I(A) = U(A_I)^\top$ is continuous on Ω_I, and $\tilde{s}_I(A) = AR_I(A)$ is continuous as a product of continuous maps.

Step 3 (A conjugacy identity for square-roots). For $M \in \mathbb{S}^r_{++}$ and $R \in O(r)$,

$$(R^\top M R)^{1/2} = R^\top M^{1/2} R. \tag{348}$$

Indeed, $R^\top M^{1/2} R$ is symmetric positive definite and

$$\big(R^\top M^{1/2} R\big)^2 = R^\top M^{1/2}(RR^\top)M^{1/2}R = R^\top M R,$$

so by uniqueness of the principal square root on $\mathbb{S}^r_{++}$, (348) holds.

Step 4 (Equivariance of the polar orthogonal factor). Let $B \in \mathrm{GL}(r)$ and $R \in O(r)$. Then

$$(BR)^\top (BR) = R^\top (B^\top B) R.$$

Applying (348) with $M = B^\top B$ gives

$$H(BR) = \big((BR)^\top (BR)\big)^{1/2} = (R^\top (B^\top B) R)^{1/2} = R^\top H(B) R.$$

Therefore,

$$U(BR) = (BR)\,H(BR)^{-1} = BR\,(R^\top H(B) R)^{-1} = BR\,(R^\top H(B)^{-1} R) = B\,H(B)^{-1}\,R = U(B)\,R,$$

where we used $(R^\top H R)^{-1} = R^\top H^{-1} R$ for orthogonal R.

Step 5 (Proof of (ii): orbit constancy). Fix $A \in \Omega_I$ and $R \in O(r)$. Then $(AR)_I = A_I R$. By Step 4,

$$U((AR)_I) = U(A_I R) = U(A_I)\,R.$$

Hence

$$R_I(AR) = U((AR)_I)^\top = (U(A_I)R)^\top = R^\top\,U(A_I)^\top = R^{-1} R_I(A).$$

Therefore,

$$\tilde{s}_I(AR) = AR\,R_I(AR) = AR\,(R^{-1} R_I(A)) = AR_I(A) = \tilde{s}_I(A),$$

proving (ii).

Step 6 (Proof of (iii): slice normalization). Write the polar decomposition $A_I = U(A_I)H(A_I)$ with $H(A_I) \in \mathbb{S}^r_{++}$. Then

$$(\tilde{s}_I(A))_I = A_I R_I(A) = A_I U(A_I)^\top = U(A_I) H(A_I) U(A_I)^\top.$$

This matrix is symmetric positive definite because it is an orthogonal conjugate of $H(A_I) \in \mathbb{S}^r_{++}$. Thus $(\tilde{s}_I(A))_I \in \mathbb{S}^r_{++}$, proving (iii).

Step 7 (Proof of (iv): factorization through the local quotient). Let $q_I : \Omega_I \to \Omega_I/O(r)$ be the orbit map and equip $\Omega_I/O(r)$ with the quotient σ-field

$$\mathcal{B}(\Omega_I/O(r)) := \{B \subseteq \Omega_I/O(r) :\ q_I^{-1}(B) \in \mathcal{B}(\Omega_I)\}.$$

By (i), $\tilde{s}_I : \Omega_I \to \mathbb{R}^{d\times r}$ is Borel measurable; by (ii), it is constant on orbits. Define $s_I : \Omega_I/O(r) \to \mathbb{R}^{d\times r}$ by $s_I(\bar{A}) := \tilde{s}_I(A)$ for any A with $q_I(A) = \bar{A}$. This is well-defined by orbit constancy and satisfies $\tilde{s}_I = s_I \circ q_I$. To see s_I is measurable, let C be a Borel set in $\mathbb{R}^{d\times r}$. Then

$$q_I^{-1}\big(s_I^{-1}(C)\big) = \{A \in \Omega_I :\ s_I(q_I(A)) \in C\} = \{A \in \Omega_I :\ \tilde{s}_I(A) \in C\} = \tilde{s}_I^{-1}(C),$$

which is Borel since $\tilde{s}_I$ is Borel. By definition of the quotient σ-field, $s_I^{-1}(C) \in \mathcal{B}(\Omega_I/O(r))$. Thus s_I is measurable and provides a local measurable section. □

Remark F.15 (Uniqueness of orbit intersection with the slice). Let $S_I := \{B \in \mathbb{R}^{d\times r} :\ B_I \in \mathbb{S}^r_{++}\}$. Each orbit in Ω_I intersects S_I in exactly one point, namely $\tilde{s}_I(A)$.

Existence. For any $A \in \Omega_I$, Lemma F.14(iii) shows $\tilde{s}_I(A) \in S_I$ and $\tilde{s}_I(A) = AR_I(A)$ lies on the right orbit of A.

Uniqueness. Suppose $B_1, B_2 \in S_I \cap \Omega_I$ lie on the same orbit, so $B_2 = B_1R$ for some $R \in O(r)$. Restricting to the I-minor gives $B_{2,I} = B_{1,I}R$ with $B_{1,I}, B_{2,I} \in \mathbb{S}^r_{++}$. Then $R = B_{1,I}^{-1}B_{2,I}$ is the product of SPD matrices, hence has positive real spectrum and is itself SPD. But an orthogonal SPD matrix must be the identity: if $R \in O(r)$ and $R \succ 0$, then all eigenvalues of R lie in $\{1\}$, so $R = I_r$. Thus $B_1 = B_2$. Therefore the orbit intersection with S_I is unique.

F.6. **Signed-permutation symmetries (hyperoctahedral groups).** Finite sign and permutation symmetries arise whenever the latent prior is invariant under coordinate relabeling and sign flips. This includes many ICA/dictionary-learning style latent models (where identifiability is only up to signed permutation).

Signed permutation group. Let $r \geq 1$ and define the finite group

$$\mathsf{SP}(r) := \Big\{R \in \mathbb{R}^{r\times r} :\ R \text{ has exactly one nonzero entry in each row/column, equal to } \pm 1\Big\}.$$

Equivalently,

$$\mathsf{SP}(r) = \{D\Pi :\ D \in \{\pm 1\}^{r\times r} \text{ diagonal},\ \Pi \in S_r \text{ permutation matrix}\}.$$

Every $R \in \mathsf{SP}(r)$ is orthogonal: $R^\top R = I_r$ and $R^{-1} = R^\top$.

ICA / linear latent-factor template with i.i.d. symmetric priors. Let $\mathcal{Z} = \mathbb{R}^r$ with its Borel σ-field and let $\nu = \nu_0^{\otimes r}$, where ν_0 is a probability measure on $\mathbb{R}$ that is symmetric about 0, i.e. $\nu_0(B) = \nu_0(-B)$ for all Borel $B \subseteq \mathbb{R}$ (equivalently, $U \sim \nu_0 \Rightarrow -U \overset{d}{=} U$). Then ν is invariant under signed permutations in the following strong (measure-pushforward) sense.

Lemma F.16 (Invariance of the product prior under $\mathsf{SP}(r)$)**.** *Let $\nu = \nu_0^{\otimes r}$ with ν_0 symmetric. For each $R \in \mathsf{SP}(r)$ define $\tau_R : \mathbb{R}^r \to \mathbb{R}^r$ by $\tau_R(z) = Rz$. Then τ_R is a Borel bijection and*

$$\nu \circ \tau_R^{-1} = \nu \qquad \text{for all } R \in \mathsf{SP}(r).$$

Proof. **Step 1 (Borel bijection).** Since τ_R is linear and invertible, it is continuous (hence Borel measurable) and bijective with inverse $\tau_{R^{-1}}$.

Step 2 (Reduction to generators). It suffices to verify $\nu(\tau_R^{-1}(C)) = \nu(C)$ for all cylinder rectangles $C = \prod_{j=1}^r C_j$ with Borel $C_j \subseteq \mathbb{R}$, because such rectangles form a π-system generating the Borel σ-field on $\mathbb{R}^r$, and both ν and $\nu \circ \tau_R^{-1}$ are probability measures.

Step 3 (Explicit form of τ_R). Write $R = D\Pi$ with Π a permutation matrix and $D = \mathrm{diag}(\sigma_1, \ldots, \sigma_r)$ with $\sigma_j \in \{\pm 1\}$. Then $(\tau_R(z))_j = \sigma_j z_{\pi^{-1}(j)}$, i.e. τ_R permutes coordinates and flips signs of selected coordinates.

Step 4 (Compute $\nu(\tau_R^{-1}(C))$). We have

$$\tau_R^{-1}(C) = \Big\{z : \ \sigma_j z_{\pi^{-1}(j)} \in C_j \ \forall j\Big\} = \Big\{z : \ z_i \in \sigma_{\pi(i)}^{-1} C_{\pi(i)} \ \forall i\Big\},$$

where $\sigma^{-1}C := \{\sigma^{-1}x : x \in C\}$ equals C if $\sigma = 1$ and $-C$ if $\sigma = -1$. By product structure of ν,

$$\nu(\tau_R^{-1}(C)) = \prod_{i=1}^{r} \nu_0\big(\sigma_{\pi(i)}^{-1} C_{\pi(i)}\big).$$

By symmetry of ν_0, $\nu_0(-B) = \nu_0(B)$ for all B, hence $\nu_0(\sigma^{-1}B) = \nu_0(B)$ for $\sigma \in \{\pm 1\}$. Therefore

$$\nu(\tau_R^{-1}(C)) = \prod_{i=1}^{r} \nu_0(C_{\pi(i)}) = \prod_{j=1}^{r} \nu_0(C_j) = \nu(C),$$

since π is a permutation. This proves $\nu \circ \tau_R^{-1} = \nu$. □

Model. Let Θ contain a mixing matrix $A \in \mathbb{R}^{d\times r}$ and possibly nuisance parameters η (e.g. noise scale/covariance). Assume the conditional law of X given $Z = z$ depends on A only through Az; concretely,

$$X \mid Z = z \sim p_{\text{noise}}(\,\cdot\,; Az, \eta),$$

where for each (m, η) the map $x \mapsto p_{\text{noise}}(x; m, \eta)$ is a μ-density on $(\mathcal{X}, \mathcal{F})$, and $(x, z, A, \eta) \mapsto p_{\text{noise}}(x; Az, \eta)$ is jointly measurable. With ν as above, define the complete-data density w.r.t. $\mu \otimes \nu$ by

$$p_{(A,\eta)}(x, z) = p_{\text{noise}}(x; Az, \eta).$$

Action and transport. Let $G = \mathsf{SP}(r)$ act on Θ by *right* transformation on A:

$$R \cdot (A, \eta) := (AR^{-1}, \eta), \qquad R \in \mathsf{SP}(r),$$

and define latent transport $\tau_R(z) := Rz$ on $\mathcal{Z} = \mathbb{R}^r$.

Lemma F.17 (Signed-permutation transport and EM equivariance)**.** *In the ICA-style template above, Assumption C.1 holds with $G = \mathsf{SP}(r)$, $\tau_R(z) = Rz$, and prior $\nu = \nu_0^{\otimes r}$. Consequently:*

*(i) (*Observed invariance*) $P^X_{R\cdot(A,\eta)} = P^X_{(A,\eta)}$ for all $R \in \mathsf{SP}(r)$;*

*(ii) (*Posterior transport*) for all Borel $B \subseteq \mathbb{R}^r$ and $P^X_{(A,\eta)}$-a.e. x,*

$$\kappa_{R\cdot(A,\eta)}(x, B) = \kappa_{(A,\eta)}\big(x, \tau_R^{-1}(B)\big);$$

*(iii) (Q-*equivariance and M-step equivariance*) whenever Q is well-defined,*

$$Q(R\cdot\theta \mid R\cdot\theta') = Q(\theta \mid \theta') \qquad \forall \theta, \theta' \in \Theta, \ \forall R \in \mathsf{SP}(r),$$

and therefore the M-step correspondence $\mathcal{T}$ satisfies $\mathcal{T}(R\cdot\theta') = R\cdot\mathcal{T}(\theta')$.

Proof. Write $\theta = (A, \eta)$.

Step 1 (Group action and measurability). $\mathsf{SP}(r)$ is a finite group under matrix multiplication. The action $\theta \mapsto R\cdot\theta = (AR^{-1}, \eta)$ is continuous in A for fixed R, hence Borel measurable on any Borel parameter space $\Theta \subseteq \mathbb{R}^{d\times r} \times \mathsf{H}$ (with H Euclidean for η). The latent transport $\tau_R(z) = Rz$ is continuous and bijective, hence Borel measurable with measurable inverse.

Step 2 (Base-measure invariance on the latent space). By Lemma F.11, $\nu \circ \tau_R^{-1} = \nu$ for all $R \in \mathsf{SP}(r)$.

Step 3 (Complete-data equivariance identity). For every $(x, z) \in \mathcal{X} \times \mathbb{R}^r$,

$$p_{R\cdot\theta}\big(x, \tau_R(z)\big) = p_{(AR^{-1},\eta)}(x, Rz) = p_{\text{noise}}(x; AR^{-1}(Rz), \eta) = p_{\text{noise}}(x; Az, \eta) = p_\theta(x, z).$$

This is exactly the density-transport identity required in Assumption C.1.

Step 4 (Observed invariance by marginalization and change of variables). Let p_θ^X be the μ-density of P_θ^X, given by integrating out z against ν:

$$p_\theta^X(x) = \int_{\mathbb{R}^r} p_\theta(x,z)\,\nu(dz) = \int_{\mathbb{R}^r} p_{\text{noise}}(x; Az, \eta)\,\nu(dz).$$

Then, using Step 3 and Step 2,

$$\begin{aligned} p_{R\cdot\theta}^X(x) &= \int_{\mathbb{R}^r} p_{R\cdot\theta}(x,z)\,\nu(dz) = \int_{\mathbb{R}^r} p_{R\cdot\theta}\big(x,\tau_R(z')\big)\,\nu(d(\tau_R(z'))) \\ &= \int_{\mathbb{R}^r} p_{R\cdot\theta}\big(x,\tau_R(z')\big)\,\nu(dz') \qquad (\text{since } \nu\circ\tau_R^{-1}=\nu) \\ &= \int_{\mathbb{R}^r} p_\theta(x,z')\,\nu(dz') = p_\theta^X(x), \end{aligned}$$

so $P_{R\cdot\theta}^X = P_\theta^X$, proving (i).

Step 5 (Posterior transport). Whenever $p_\theta^X(x) > 0$, Bayes' formula yields the posterior as a probability kernel

$$\kappa_\theta(x,B) = \frac{\int_B p_\theta(x,z)\,\nu(dz)}{\int_{\mathbb{R}^r} p_\theta(x,z)\,\nu(dz)}.$$

Using Step 3 and the change of variables $z = \tau_R(z')$ plus Step 2,

$$\int_B p_{R\cdot\theta}(x,z)\,\nu(dz) = \int_{\tau_R^{-1}(B)} p_{R\cdot\theta}\big(x,\tau_R(z')\big)\,\nu(dz') = \int_{\tau_R^{-1}(B)} p_\theta(x,z')\,\nu(dz'),$$

and similarly $\int p_{R\cdot\theta}(x,z)\nu(dz) = \int p_\theta(x,z)\nu(dz)$. Therefore $\kappa_{R\cdot\theta}(x,B) = \kappa_\theta(x,\tau_R^{-1}(B))$ for all B and for all x with $p_\theta^X(x) > 0$. Since $P_{R\cdot\theta}^X = P_\theta^X$ (Step 4), this holds P_θ^X-a.e., proving (ii).

Step 6 (Q-equivariance and M-step equivariance). Assume $Q(\cdot \mid \cdot)$ is well-defined (integrability of $\log p_\theta(X,Z)$ under the relevant posteriors). By definition,

$$Q(\theta \mid \theta') = \mathbb{E}_{P^\star}\Big[\int_{\mathbb{R}^r} \log p_\theta(X,z)\,\kappa_{\theta'}(X,dz)\Big].$$

Apply this with (θ,θ') replaced by $(R\cdot\theta, R\cdot\theta')$ and use Steps 3 and 5 with the change of variables $z = \tau_R(z')$:

$$\begin{aligned} \int \log p_{R\cdot\theta}(X,z)\,\kappa_{R\cdot\theta'}(X,dz) &= \int \log p_{R\cdot\theta}\big(X,\tau_R(z')\big)\,\kappa_{\theta'}(X,dz') \\ &= \int \log p_\theta(X,z')\,\kappa_{\theta'}(X,dz'). \end{aligned}$$

Taking $\mathbb{E}_{P^\star}$ yields $Q(R\cdot\theta \mid R\cdot\theta') = Q(\theta \mid \theta')$. Equivariance of the M-step correspondence $\mathcal{T}(\theta') \in \arg\max_\theta Q(\theta \mid \theta')$ follows immediately: if $\theta \in \mathcal{T}(\theta')$, then for all ϑ,

$$Q(R\cdot\theta \mid R\cdot\theta') = Q(\theta \mid \theta') \geq Q(\vartheta \mid \theta') = Q(R\cdot\vartheta \mid R\cdot\theta'),$$

so $R\cdot\theta \in \mathcal{T}(R\cdot\theta')$. □

Measurable section. Since $\mathsf{SP}(r)$ is finite, Lemma F.5 yields a global Borel section on $\Theta/\mathsf{SP}(r)$ whenever Assumption F.4 holds.

F.7. **Rotational symmetries in linear-Gaussian factor models ($G = O(r)$).** This is the continuous-group analogue of label switching. It is canonical in factor analysis and probabilistic PCA, where A is only identified up to right orthogonal transformations.
Model and domination. Let $\mathcal{Z} = \mathbb{R}^r$ with its Borel σ-field and let $\nu = \mathcal{N}(0, I_r)$ be the standard Gaussian *measure* on $\mathbb{R}^r$. Let μ be Lebesgue measure on $\mathbb{R}^d$, fix $d \geq r$, and let $\Psi \in \mathbb{R}^{d\times d}$ be positive definite. Under P_A, sample $Z \sim \nu$ and then

$$X \mid Z = z \sim \mathcal{N}(Az, \Psi).$$

Write $\varphi_d(\cdot; m, \Psi)$ for the μ-density of $\mathcal{N}(m, \Psi)$. Since we take $\mu \otimes \nu$ as the dominating measure for (X, Z), the joint density is simply the conditional density:

$$p_A(x, z) \;:=\; \frac{dP_A^{X,Z}}{d(\mu\otimes\nu)}(x, z) \;=\; \varphi_d(x; Az, \Psi).$$

Action and transport. Let $G = O(r)$ act by right multiplication $R \cdot A := AR$ and define $\tau_R(z) := R^\top z$.

Lemma F.18 (Latent transport and observed invariance in the factor model)**.** *In the factor model above, Assumption C.1 holds with $G = O(r)$, $\tau_R(z) = R^\top z$, and $\nu = \mathcal{N}(0, I_r)$. Consequently, $P_{AR}^X = P_A^X$ for all $R \in O(r)$; posteriors transport by τ_R; and Q (hence the M-step) is $O(r)$-equivariant whenever defined.*

Proof. **Step 1 (Group action, transport map, and measurability).** The map $(A, R) \mapsto AR$ is continuous, hence the action $A \mapsto AR$ is Borel measurable on $\mathbb{R}^{d\times r}$ for each fixed R. For each $R \in O(r)$, $\tau_R(z) = R^\top z$ is linear and invertible, hence a Borel bijection with inverse $\tau_{R^\top}$.

Step 2 (Gaussian base-measure invariance). Let $Z \sim \mathcal{N}(0, I_r)$. For any $R \in O(r)$, $R^\top Z \overset{d}{=} Z$. Equivalently, $\nu \circ \tau_R^{-1} = \nu$. A quick verification: the characteristic function satisfies

$$\mathbb{E}[e^{i\langle t, R^\top Z\rangle}] = \mathbb{E}[e^{i\langle Rt, Z\rangle}] = \exp\Big(-\tfrac{1}{2}\|Rt\|^2\Big) = \exp\Big(-\tfrac{1}{2}\|t\|^2\Big),$$

using R orthogonal.

Step 3 (Complete-data equivariance identity). For all $(x, z) \in \mathbb{R}^d \times \mathbb{R}^r$ and $R \in O(r)$,

$$p_{AR}\big(x, \tau_R(z)\big) = \varphi_d\big(x; AR(R^\top z), \Psi\big) = \varphi_d(x; Az, \Psi) = p_A(x, z).$$

Thus the density-transport identity holds.

Step 4 (Observed invariance by marginalization). The observed density is

$$p_A^X(x) = \int_{\mathbb{R}^r} \varphi_d(x; Az, \Psi)\, \nu(dz).$$

Using Step 3 with $z = \tau_R(z')$ and Step 2,

$$p_{AR}^X(x) = \int \varphi_d(x; ARz, \Psi)\, \nu(dz) = \int \varphi_d\big(x; AR(\tau_{R^\top}(z')), \Psi\big)\, \nu(dz') = \int \varphi_d(x; Az', \Psi)\, \nu(dz') = p_A^X(x),$$

so $P_{AR}^X = P_A^X$.

Step 5 (Posterior transport). For $p_A^X(x) > 0$,

$$\kappa_A(x, B) = \frac{\int_B p_A(x, z)\, \nu(dz)}{\int p_A(x, z)\, \nu(dz)}.$$

Using Step 3 and change of variables $z = \tau_R(z')$ with Step 2,

$$\int_B p_{AR}(x,z)\,\nu(dz) = \int_{\tau_R^{-1}(B)} p_{AR}\big(x,\tau_R(z')\big)\,\nu(dz') = \int_{\tau_R^{-1}(B)} p_A(x,z')\,\nu(dz').$$

The denominator is invariant by the same argument, so $\kappa_{AR}(x,B) = \kappa_A(x,\tau_R^{-1}(B))$.

Step 6 (Q-equivariance). As in Lemma F.12, combine Step 3 (transport for $\log p$) and Step 5 (posterior transport) inside the defining conditional expectation for Q. This yields $Q(AR_1 \mid AR_2) = Q(A \mid A)$ with the appropriate pairing (more generally $Q(R\cdot A \mid R\cdot A') = Q(A \mid A')$), whenever both sides are finite. Equivariance of the M-step correspondence follows. □

Local measurable sections (minor charts and polar normalization). Global measurable sections for $A \mapsto AR$ generally do not exist on all of $\mathbb{R}^{d\times r}$, but local sections on an open cover of the full-column-rank set suffice for slice-based arguments.
Minor charts. Fix $I \subseteq \{1,\dots,d\}$ with $|I| = r$ and let A_I be the $r\times r$ submatrix formed by rows indexed by I. Define the open set

$$\Omega_I := \{A \in \mathbb{R}^{d\times r} :\ \det(A_I) \neq 0\}.$$

Every full-column-rank A lies in at least one Ω_I.
Polar factors and normalization. For invertible $B \in \mathbb{R}^{r\times r}$ define its polar decomposition $B = U(B)H(B)$ with

$$H(B) := (B^\top B)^{1/2} \in \mathbb{S}^r_{++}, \qquad U(B) := B\,H(B)^{-1} \in O(r).$$

On Ω_I define

$$R_I(A) := U(A_I)^\top \in O(r), \qquad \tilde{s}_I(A) := A\,R_I(A) \in \mathbb{R}^{d\times r}.$$

Lemma F.19 (Local section for the right $O(r)$ action)**.** *Fix I and work on Ω_I. Then:*

(i) R_I and $\tilde{s}_I$ are continuous (hence Borel) on Ω_I;
(ii) $\tilde{s}_I$ is constant on right-orbits in Ω_I: $\tilde{s}_I(AR) = \tilde{s}_I(A)$ for all $R \in O(r)$;
(iii) $(\tilde{s}_I(A))_I \in \mathbb{S}^r_{++}$ (slice normalization);
(iv) consequently, $\tilde{s}_I$ factors through the local quotient $\Omega_I/O(r)$ to a measurable section.

Proof. **Step 1 (Continuity of the polar maps on $\mathrm{GL}(r)$).** On $\mathrm{GL}(r)$, the map $B \mapsto B^\top B$ is polynomial and hence continuous, and its image lies in $\mathbb{S}^r_{++}$. The principal square-root map $M \mapsto M^{1/2}$ is continuous on $\mathbb{S}^r_{++}$ (indeed smooth), and so is inversion $M \mapsto M^{-1}$. Therefore $H(B) = (B^\top B)^{1/2}$ and $H(B)^{-1}$ depend continuously on $B \in \mathrm{GL}(r)$. Consequently $U(B) = B\,H(B)^{-1}$ is continuous on $\mathrm{GL}(r)$.

Step 2 (Proof of (i)). On Ω_I, the map $A \mapsto A_I$ is linear (hence continuous). Since $A_I \in \mathrm{GL}(r)$ on Ω_I, Step 1 implies $A \mapsto U(A_I)$ is continuous on Ω_I. Thus $R_I(A) = U(A_I)^\top$ is continuous on Ω_I, and $\tilde{s}_I(A) = AR_I(A)$ is continuous as a product of continuous maps.

Step 3 (A conjugacy identity for square-roots). For $M \in \mathbb{S}^r_{++}$ and $R \in O(r)$,

$$(R^\top M R)^{1/2} = R^\top M^{1/2} R. \tag{349}$$

Indeed, $R^\top M^{1/2} R$ is symmetric positive definite and

$$\big(R^\top M^{1/2} R\big)^2 = R^\top M^{1/2}(RR^\top)M^{1/2}R = R^\top M R,$$

so by uniqueness of the principal square root on $\mathbb{S}^r_{++}$, (348) holds.

Step 4 (Equivariance of the polar orthogonal factor). Let $B \in \mathrm{GL}(r)$ and $R \in O(r)$. Then

$$(BR)^\top (BR) = R^\top (B^\top B) R.$$

Applying (348) with $M = B^\top B$ gives

$$H(BR) = \big((BR)^\top (BR)\big)^{1/2} = (R^\top (B^\top B) R)^{1/2} = R^\top H(B) R.$$

Therefore,

$$U(BR) = (BR)\, H(BR)^{-1} = BR\, (R^\top H(B) R)^{-1} = BR\, (R^\top H(B)^{-1} R) = B\, H(B)^{-1}\, R = U(B)\, R,$$

where we used $(R^\top H R)^{-1} = R^\top H^{-1} R$ for orthogonal R.

Step 5 (Proof of (ii): orbit constancy). Fix $A \in \Omega_I$ and $R \in O(r)$. Then $(AR)_I = A_I R$. By Step 4,

$$U((AR)_I) = U(A_I R) = U(A_I)\, R.$$

Hence

$$R_I(AR) = U((AR)_I)^\top = (U(A_I) R)^\top = R^\top\, U(A_I)^\top = R^{-1} R_I(A).$$

Therefore,

$$\tilde{s}_I(AR) = AR\, R_I(AR) = AR\, (R^{-1} R_I(A)) = A R_I(A) = \tilde{s}_I(A),$$

proving (ii).

Step 6 (Proof of (iii): slice normalization). Write the polar decomposition $A_I = U(A_I) H(A_I)$ with $H(A_I) \in \mathbb{S}^r_{++}$. Then

$$(\tilde{s}_I(A))_I = A_I R_I(A) = A_I U(A_I)^\top = U(A_I) H(A_I) U(A_I)^\top.$$

This matrix is symmetric positive definite because it is an orthogonal conjugate of $H(A_I) \in \mathbb{S}^r_{++}$. Thus $(\tilde{s}_I(A))_I \in \mathbb{S}^r_{++}$, proving (iii).

Step 7 (Proof of (iv): factorization through the local quotient). Let $q_I : \Omega_I \to \Omega_I / O(r)$ be the orbit map and equip $\Omega_I / O(r)$ with the quotient σ-field

$$\mathcal{B}(\Omega_I / O(r)) := \{B \subseteq \Omega_I / O(r) :\ q_I^{-1}(B) \in \mathcal{B}(\Omega_I)\}.$$

By (i), $\tilde{s}_I : \Omega_I \to \mathbb{R}^{d \times r}$ is Borel measurable; by (ii), it is constant on orbits. Define $s_I : \Omega_I / O(r) \to \mathbb{R}^{d \times r}$ by $s_I(\bar{A}) := \tilde{s}_I(A)$ for any A with $q_I(A) = \bar{A}$. This is well-defined by orbit constancy and satisfies $\tilde{s}_I = s_I \circ q_I$. To see s_I is measurable, let C be a Borel set in $\mathbb{R}^{d \times r}$. Then

$$q_I^{-1}\big(s_I^{-1}(C)\big) = \{A \in \Omega_I :\ s_I(q_I(A)) \in C\} = \{A \in \Omega_I :\ \tilde{s}_I(A) \in C\} = \tilde{s}_I^{-1}(C),$$

which is Borel since $\tilde{s}_I$ is Borel. By definition of the quotient σ-field, $s_I^{-1}(C) \in \mathcal{B}(\Omega_I / O(r))$. Thus s_I is measurable and provides a local measurable section. □

Remark F.20 (Uniqueness of orbit intersection with the slice). Let $S_I := \{B \in \mathbb{R}^{d \times r} :\ B_I \in \mathbb{S}^r_{++}\}$. Each orbit in Ω_I intersects S_I in exactly one point, namely $\tilde{s}_I(A)$.

Existence. For any $A \in \Omega_I$, Lemma F.14(iii) shows $\tilde{s}_I(A) \in S_I$ and $\tilde{s}_I(A) = A R_I(A)$ lies on the right orbit of A.

Uniqueness. Suppose $B_1, B_2 \in S_I \cap \Omega_I$ lie on the same orbit, so $B_2 = B_1 R$ for some $R \in O(r)$. Restricting to the I-minor gives $B_{2,I} = B_{1,I} R$ with $B_{1,I}, B_{2,I} \in \mathbb{S}^r_{++}$. Then $R = B_{1,I}^{-1} B_{2,I}$ is the product of SPD matrices, hence has positive real spectrum and is itself SPD. But an orthogonal SPD matrix must be the identity: if $R \in O(r)$ and $R \succ 0$, then all eigenvalues of R lie in $\{1\}$, so $R = I_r$. Thus $B_1 = B_2$. Therefore the orbit intersection with S_I is unique.

F.8. **Pure sign symmetries ($G = \{\pm 1\}$).** This is the simplest nontrivial symmetry and is often combined with permutations (Appendix F.4).
Setup. Let $\Theta \subseteq \mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$ be Borel and assume it is stable under the sign-flip map

$$\iota : \Theta \to \Theta, \qquad \iota(u, v) = (-u, -v).$$

Let $G = \{\pm 1\}$ act on Θ by

$$g \cdot (u, v) := (gu, gv).$$

Let $q : \Theta \to \Theta/G$ be the orbit map, and equip Θ/G with the quotient σ-field

$$\mathcal{B}(\Theta/G) := \{A \subseteq \Theta/G : \ q^{-1}(A) \in \mathcal{B}(\Theta)\}.$$

Assumption F.21 (Sign-invariant parametrization)**.** There exists a measurable map $\Psi : \Theta \to \mathcal{W}$ into a measurable space $\mathcal{W}$ such that:

(i) (*Sign invariance*) $\Psi(u, v) = \Psi(-u, -v)$ for all $(u, v) \in \Theta$;
(ii) (*Model depends only on* Ψ) if $\Psi(u, v) = \Psi(u', v')$, then $P^X_{(u,v)} = P^X_{(u',v')}$.

Lemma F.22 (Observed invariance and projection-set invariance under sign flips)**.** *Under Assumption F.21, $P^X_{-1\cdot\theta} = P^X_\theta$ for all $\theta \in \Theta$. Consequently, $\Phi(-\theta) = \Phi(\theta)$ and $\bar{\Theta} = \arg\max_{\theta\in\Theta} \Phi(\theta)$ is G-invariant.*

Proof. Fix $\theta \in \Theta$. By Assumption F.21(i), $\Psi(\theta) = \Psi(-\theta)$. By Assumption F.21(ii), equality of Ψ implies equality of the observed laws, hence $P^X_{-\theta} = P^X_\theta$, i.e. observed invariance holds for the $G = \{\pm 1\}$ action.

Now apply Lemma F.7: observed invariance implies $\Phi(-\theta) = \Phi(\theta)$ (as extended real numbers), and the maximizer set $\bar{\Theta}$ is invariant under the group action. (Equivalently: $\theta \in \bar{\Theta} \Rightarrow -\theta \in \bar{\Theta}$.) □

A canonical Borel representative. A convenient global representative for each orbit is obtained by fixing a measurable tie-breaking rule. We use the "first nonzero coordinate is positive" convention.

To make this explicit, identify $\mathbb{R}^{d_1} \times \mathbb{R}^{d_2} \cong \mathbb{R}^d$ with $d = d_1 + d_2$ by concatenation: for $\theta = (u, v)$ write $\theta = (\theta_1, \ldots, \theta_d)$ where $(\theta_1, \ldots, \theta_{d_1}) = u$ and $(\theta_{d_1+1}, \ldots, \theta_d) = v$.

Define the *first-nonzero index* map $J : \Theta \to \{1, \ldots, d\} \cup \{\infty\}$ by

$$J(\theta) := \begin{cases} \min\{j \in \{1, \ldots, d\} : \ \theta_j \neq 0\}, & \theta \neq 0, \\ \infty, & \theta = 0, \end{cases}$$

and define the *canonical sign* $\sigma : \Theta \to \{\pm 1\}$ by

$$\sigma(\theta) := \begin{cases} +1, & J(\theta) = \infty, \\ +1, & J(\theta) < \infty \text{ and } \theta_{J(\theta)} > 0, \\ -1, & J(\theta) < \infty \text{ and } \theta_{J(\theta)} < 0. \end{cases}$$

Finally define the *canonical relabeling map* $\tilde{s} : \Theta \to \Theta$ by

$$\tilde{s}(\theta) := \sigma(\theta) \cdot \theta. \tag{350}$$

(When $\theta = 0$, this fixes the point: $\tilde{s}(0) = 0$.)

Lemma F.23 (A Borel section for the sign symmetry)**.** *Assume Θ is Borel and stable under $\theta \mapsto -\theta$. Let $\tilde{s}$ be defined by (350). Then:*

(i) $\tilde{s}$ is Borel measurable;
(ii) $\tilde{s}$ is constant on G-orbits: $\tilde{s}(-\theta) = \tilde{s}(\theta)$ for all $\theta \in \Theta$;

(iii) $\tilde{s}$ *factors through the quotient: there exists a unique measurable map* $s : \Theta/G \to \Theta$ *such that*

$$\tilde{s} = s \circ q;$$

(iv) s *is a measurable section:* $q(s(\bar{\theta})) = \bar{\theta}$ *for all* $\bar{\theta} \in \Theta/G$.

Proof. **Step 1 (Measurability of the index map J).** For each $j \in \{1, \dots, d\}$,

$$\{J = j\} = \Big(\bigcap_{i<j} \{\theta_i = 0\}\Big) \cap \{\theta_j \neq 0\},$$

which is Borel since each coordinate projection $\theta \mapsto \theta_i$ is continuous and $\{0\}$ is closed. Also,

$$\{J = \infty\} = \bigcap_{i=1}^{d} \{\theta_i = 0\},$$

which is Borel. Therefore J is measurable as a map into the countable measurable space $\{1, \dots, d\} \cup \{\infty\}$.

Step 2 (Measurability of the sign map σ). For $j \in \{1, \dots, d\}$ define the Borel sets

$$A_j^+ := \{J = j\} \cap \{\theta_j > 0\}, \qquad A_j^- := \{J = j\} \cap \{\theta_j < 0\}.$$

Then $\{J = \infty\} \cup (\bigcup_j A_j^+) = \{\sigma = +1\}$ and $\bigcup_j A_j^- = \{\sigma = -1\}$ are Borel. Hence σ is measurable.

Step 3 (Measurability of $\tilde{s}$). The map $\theta \mapsto -\theta$ is continuous, hence Borel. Since

$$\tilde{s}(\theta) = \mathbf{1}_{\{\sigma(\theta)=+1\}}\theta + \mathbf{1}_{\{\sigma(\theta)=-1\}}(-\theta),$$

and indicator functions of Borel sets are measurable, $\tilde{s}$ is Borel measurable.

Step 4 (Orbit constancy). Fix $\theta \in \Theta$. If $\theta = 0$ then $\tilde{s}(-\theta) = \tilde{s}(\theta) = 0$. Otherwise, $J(-\theta) = J(\theta)$ and $(-\theta)_{J(\theta)} = -\theta_{J(\theta)}$, so $\sigma(-\theta) = -\sigma(\theta)$. Therefore

$$\tilde{s}(-\theta) = \sigma(-\theta) \cdot (-\theta) = (-\sigma(\theta)) \cdot (-\theta) = \sigma(\theta) \cdot \theta = \tilde{s}(\theta).$$

Thus $\tilde{s}$ is constant on orbits.

Step 5 (Factorization through the quotient and measurability of s). By Step 4, $\tilde{s}$ is G-invariant. By Lemma F.1 (invariant functions factor through the quotient), there exists a unique measurable map $s : \Theta/G \to \Theta$ such that $\tilde{s} = s \circ q$.

Step 6 (Section property). Fix $\theta \in \Theta$. By construction, $\tilde{s}(\theta) \in \{\theta, -\theta\}$. Therefore $\tilde{s}(\theta)$ lies in the same orbit as θ, i.e. $q(\tilde{s}(\theta)) = q(\theta)$. Using $\tilde{s} = s \circ q$ and surjectivity of q, for every $\bar{\theta} \in \Theta/G$ choose θ with $q(\theta) = \bar{\theta}$ and compute

$$q(s(\bar{\theta})) = q(s(q(\theta))) = q(\tilde{s}(\theta)) = q(\theta) = \bar{\theta}.$$

This proves that s is a section. □

F.9. **How to read these instantiations in the main text.** The preceding examples cover common nonidentifiabilities in EM analyses:

- **Finite label-switching** (S_k): mixtures, HMMs, latent class and SBM-type models (Appendix F.1–F.3).
- **Signed permutations** (finite $\mathsf{SP}(r)$): ICA/dictionary-learning style latent models with i.i.d. symmetric priors (Appendix F.4).
- **Orthogonal rotations** ($O(r)$): factor analysis and probabilistic PCA with standard Gaussian priors (Appendix F.5).
- **Pure sign flips** ($\{\pm 1\}$): the simplest symmetry, often combined with permutations (Appendix F.8).

Across these cases, the verification steps used in the main text always follow the same template.
Step A: verify latent transport. One explicitly constructs a measurable transport $\tau_g : \mathcal{Z} \to \mathcal{Z}$ such that:

(i) $\tau_{gh} = \tau_g \circ \tau_h$ and $\tau_e = \mathrm{id}$ (compatibility with the group);
(ii) $\nu \circ \tau_g^{-1} = \nu$ (base latent measure is preserved);
(iii) the complete-data likelihood is equivariant:

$$p_{g\cdot\theta}\big(x, \tau_g(z)\big) = p_\theta(x, z) \qquad \text{for all } (x, z) \text{ (or for a.e. } (x, z) \text{ w.r.t. } \mu \otimes \nu).$$

These items are exactly what Assumption C.1 encodes.
Step B: deduce observed invariance by marginalization. From ν-preservation and the complete-data equivariance, one obtains

$$p^X_{g\cdot\theta}(x) = \int p_{g\cdot\theta}(x, z)\, \nu(dz) = \int p_\theta(x, z)\, \nu(dz) = p^X_\theta(x),$$

hence $P^X_{g\cdot\theta} = P^X_\theta$. This yields invariance of the population objective Φ and G-invariance of $\bar{\Theta} = \arg\max \Phi$ (Lemma F.7).
Step C: transport posteriors and obtain Q-equivariance. Bayes' rule plus the same change-of-variables argument implies posterior transport

$$\kappa_{g\cdot\theta}(x, B) = \kappa_\theta\big(x, \tau_g^{-1}(B)\big),$$

and inserting this into the definition of Q yields $Q(g \cdot \theta \mid g \cdot \theta') = Q(\theta \mid \theta')$ whenever defined. Thus the M-step correspondence is equivariant: $\mathcal{T}(g \cdot \theta') = g \cdot \mathcal{T}(\theta')$.
Step D: choose representatives (sections/slices) to enforce transverse geometry. To turn orbit-level statements into concrete Euclidean estimates on parameters, one chooses orbit representatives:

- for **finite groups**, a global measurable section exists by orbit minimization (Lemma F.5) or by explicit tie-breaking rules (Lemma F.23);
- for **continuous groups** like $O(r)$, one typically works on an open cover and uses local slices (Lemma F.14).

These representatives are what make it meaningful to impose strong concavity/contractivity *transversely* to orbits, while still respecting the intrinsic nonidentifiability.
Summary of what the main results need. The main text uses only the following outputs of the appendices:

(i) orbit-level targets under misspecification are well-defined via $\bar{\Phi}$ on Θ/G;
(ii) EM dynamics are compatible with the symmetry (equivariance), so orbit-level convergence statements are canonical;
(iii) slice/section representatives allow Euclidean curvature and contraction estimates in directions orthogonal to orbits.

Appendix G. IPM choices: separation, moduli, and estimability

Throughout this appendix, let $(\mathsf{X}, \mathcal{F})$ be the observation space, and let $\mathcal{P}$ denote a class of probability measures on $(\mathsf{X}, \mathcal{F})$ that contains the model range

$$\mathcal{P}_\Theta := \{P^X_\theta : \theta \in \Theta\}.$$

For a class of measurable test functions $\mathcal{F}_0 \subseteq \{f : \mathsf{X} \to \mathbb{R}\}$ for which all expectations exist (possibly as extended reals when explicitly stated), the associated integral probability metric (IPM) is

$$d_{\mathcal{F}_0}(P, Q) := \sup_{f \in \mathcal{F}_0} \big|\mathbb{E}_P[f(X)] - \mathbb{E}_Q[f(X)]\big|. \tag{351}$$

We use $\mathsf{D}(\cdot,\cdot)$ as a generic IPM notation when the underlying class is clear.
Role in the EM theory. The main text uses IPMs in two logically distinct ways:

(1) *Statistical deviation control*: bounding $\mathsf{D}(\widehat{P}_n, P^\star)$ (one-sample) or $\mathsf{D}(\widehat{P}_n, \widehat{Q}_m)$ (two-sample) as a function of n and the complexity of $\mathcal{F}_0$;
(2) *Geometric translation*: relating $\mathsf{D}(P_\theta^X, P_{\theta'}^X)$ to orbit/quotient distances (Appendix B.4) via moduli of continuity and separation properties on Θ/G.

Accordingly, this appendix organizes IPM choices around (i) separation on the relevant range, (ii) modulus comparisons, and (iii) estimability.

G.1. **When an IPM is determining vs. pseudometric.**
Basic separation notions. Fix a class $\mathcal{P} \subseteq \mathcal{P}(\mathsf{X})$ and a measurable class $\mathcal{F}_0 \subseteq \{f : \mathsf{X} \to \mathbb{R}\}$ for which $\mathbb{E}_P f$ is well-defined for all $P \in \mathcal{P}$ and all $f \in \mathcal{F}_0$ (allowing extended values only when explicitly intended). Recall the induced integral probability metric (IPM)

$$d_{\mathcal{F}_0}(P,Q) \;:=\; \sup_{f\in\mathcal{F}_0} \Big|\mathbb{E}_P f - \mathbb{E}_Q f\Big|, \qquad P,Q \in \mathcal{P}.$$

We say:

- $\mathcal{F}_0$ is *measure-determining on* $\mathcal{P}$ (or $\mathcal{P}$*-determining*) if

$$d_{\mathcal{F}_0}(P,Q) = 0 \implies P = Q \qquad \forall P,Q \in \mathcal{P}.$$

- $d_{\mathcal{F}_0}$ is a *pseudometric on* $\mathcal{P}$ if it is symmetric, satisfies the triangle inequality, and $d_{\mathcal{F}_0}(P,P) = 0$, but it may fail to separate points.

By construction, $d_{\mathcal{F}_0}$ is always a pseudometric on any subcollection of measures for which the defining expectations are finite and the class $\mathcal{F}_0$ is closed under negation (or if we replace $\mathcal{F}_0$ by $\mathcal{F}_0^{\pm} := \mathcal{F}_0 \cup (-\mathcal{F}_0)$).
Elementary sufficient conditions for being determining. The following templates are standard.

Lemma G.1 (Determining classes via generating π-systems)**.** *Let $\mathcal{C} \subseteq \mathcal{F}$ be a π-system generating $\mathcal{F}$ (i.e. $\sigma(\mathcal{C}) = \mathcal{F}$), and suppose*

$$\{\mathbf{1}_C : \; C \in \mathcal{C}\} \;\subseteq\; \mathcal{F}_0.$$

Then $d_{\mathcal{F}_0}(P,Q) = 0$ implies $P = Q$ for all probability measures P, Q on $(\mathsf{X}, \mathcal{F})$.

Proof. Assume $d_{\mathcal{F}_0}(P,Q) = 0$. Since $\mathbf{1}_C \in \mathcal{F}_0$ for each $C \in \mathcal{C}$,

$$0 = d_{\mathcal{F}_0}(P,Q) \;\geq\; \big|\mathbb{E}_P \mathbf{1}_C - \mathbb{E}_Q \mathbf{1}_C\big| \;=\; |P(C) - Q(C)| \qquad \forall C \in \mathcal{C}.$$

Hence $P(C) = Q(C)$ for all $C \in \mathcal{C}$.

Define

$$\Lambda := \{A \in \mathcal{F} : \; P(A) = Q(A)\}.$$

We verify Λ is a λ-system:

- $\mathsf{X} \in \Lambda$ because $P(\mathsf{X}) = Q(\mathsf{X}) = 1$.
- If $A \in \Lambda$, then $P(A^c) = 1 - P(A) = 1 - Q(A) = Q(A^c)$, so $A^c \in \Lambda$.
- If $A_1, A_2, \ldots$ are disjoint with each $A_i \in \Lambda$, then by countable additivity, $P(\cup_i A_i) = \sum_i P(A_i) = \sum_i Q(A_i) = Q(\cup_i A_i)$, so $\cup_i A_i \in \Lambda$.

Thus Λ is a λ-system.

Since $P(C) = Q(C)$ for all $C \in \mathcal{C}$, we have $\mathcal{C} \subseteq \Lambda$. By the π–λ theorem, $\sigma(\mathcal{C}) \subseteq \Lambda$. But $\sigma(\mathcal{C}) = \mathcal{F}$, hence $\Lambda = \mathcal{F}$, meaning $P(A) = Q(A)$ for all $A \in \mathcal{F}$. Therefore $P = Q$. □

Lemma G.2 (Determining classes via dense function algebras on compact spaces)**.** *Let* (X, ρ) *be compact metric and let* $\mathcal{F}_0 \subseteq C(\mathsf{X})$ *be a subalgebra containing constants and separating points. Let* $\operatorname{span}(\mathcal{F}_0)$ *denote the linear span of* $\mathcal{F}_0$*. Then for all Borel probability measures* P, Q *on* X*,*

$$\big(\mathbb{E}_P f = \mathbb{E}_Q f \ \forall f \in \mathcal{F}_0\big) \implies P = Q,$$

equivalently $d_{\mathcal{F}_0}(P, Q) = 0 \Rightarrow P = Q$*.*

Proof. Assume $\mathbb{E}_P f = \mathbb{E}_Q f$ for all $f \in \mathcal{F}_0$. By linearity of the integral, the same equality holds for all $g \in \operatorname{span}(\mathcal{F}_0)$.

Step 1 (Uniform density). Because X is compact metric, $C(\mathsf{X})$ is a Banach space under $\|\cdot\|_\infty$. By the Stone–Weierstrass theorem, the hypotheses (subalgebra, constants, point-separating) imply that $\operatorname{span}(\mathcal{F}_0)$ is uniformly dense in $C(\mathsf{X})$: for every $h \in C(\mathsf{X})$ and every $\varepsilon > 0$, there exists $g \in \operatorname{span}(\mathcal{F}_0)$ with $\|h - g\|_\infty \le \varepsilon$.

Step 2 (Pass equality of integrals from the dense subspace to all of $C(\mathsf{X})$). Fix $h \in C(\mathsf{X})$ and choose $g_\varepsilon \in \operatorname{span}(\mathcal{F}_0)$ such that $\|h - g_\varepsilon\|_\infty \le \varepsilon$. Then

$$\big|\mathbb{E}_P h - \mathbb{E}_Q h\big| \le \big|\mathbb{E}_P(h - g_\varepsilon)\big| + \big|\mathbb{E}_P g_\varepsilon - \mathbb{E}_Q g_\varepsilon\big| + \big|\mathbb{E}_Q(g_\varepsilon - h)\big|.$$

The middle term is 0 by the assumption on $\operatorname{span}(\mathcal{F}_0)$. The outer terms are bounded by $\|h - g_\varepsilon\|_\infty \le \varepsilon$ since P, Q are probability measures:

$$\big|\mathbb{E}_P(h - g_\varepsilon)\big| \le \mathbb{E}_P|h - g_\varepsilon| \le \varepsilon, \qquad \big|\mathbb{E}_Q(g_\varepsilon - h)\big| \le \varepsilon.$$

Thus $|\mathbb{E}_P h - \mathbb{E}_Q h| \le 2\varepsilon$. Letting $\varepsilon \downarrow 0$ yields $\mathbb{E}_P h = \mathbb{E}_Q h$ for all $h \in C(\mathsf{X})$.

Step 3 (Continuous functions determine Borel measures on a compact metric space). On a compact metric space, equality of integrals against all continuous bounded functions implies equality of measures (equivalently: $C(\mathsf{X})$ determines the weak topology and separates probability measures). Hence $P = Q$. □

Canonical examples.

- **Total variation (TV).** Taking $\mathcal{F}_0 = \{f : \|f\|_\infty \le 1\}$ yields
$$d_{\mathcal{F}_0}(P, Q) = \sup_{\|f\|_\infty \le 1} \big|\mathbb{E}_P f - \mathbb{E}_Q f\big| = 2\,\|P - Q\|_{\mathrm{TV}},$$
which is determining on all probability measures.
- **Bounded-Lipschitz (BL) metric.** On a metric space (X, ρ), let
$$\|f\|_{\mathrm{BL}} := \|f\|_\infty + \operatorname{Lip}_\rho(f), \qquad \mathcal{F}_0 := \{f : \|f\|_{\mathrm{BL}} \le 1\}.$$
On Polish (X, ρ), d_{BL} is determining and metrizes weak convergence.
- **Wasserstein-1.** On (X, ρ), define $\mathcal{F}_0 = \{f : \operatorname{Lip}_\rho(f) \le 1\}$. Then $d_{\mathcal{F}_0} = W_1$ by Kantorovich–Rubinstein duality, but it is finite only on measures with finite first moment, so W_1 is determining on
$$\mathcal{P}_1 := \Big\{P : \int \rho(x, x_0)\, P(dx) < \infty \text{ for some (hence all) } x_0\Big\}.$$
- **MMD (kernel IPM).** Let k be a measurable positive definite kernel with RKHS $\mathcal{H}_k$, and take
$$\mathcal{F}_0 := \{f \in \mathcal{H}_k : \ \|f\|_{\mathcal{H}_k} \le 1\}.$$
Then $d_{\mathcal{F}_0}(P, Q) = \|\mu_P - \mu_Q\|_{\mathcal{H}_k}$ where $\mu_P := \mathbb{E}_P[k(X, \cdot)]$ is the kernel mean embedding (when k is bounded). This is determining on a class $\mathcal{P}$ iff k is *characteristic on* $\mathcal{P}$; otherwise it is only a pseudometric.

Determining *on the model range* (the relevant notion here). In applications one rarely needs $\mathcal{F}_0$ to separate *all* probability measures on X. What matters is separation on the *image of the model*, possibly modulo symmetries. Write

$$\mathcal{P}_\Theta \;:=\; \{P_\theta^X :\; \theta \in \Theta\} \;\subseteq\; \mathcal{P}(\mathsf{X}).$$

Even if $\mathcal{F}_0$ is not determining on $\mathcal{P}(\mathsf{X})$, it can be determining on $\mathcal{P}_\Theta$ (or determining on $\mathcal{P}_\Theta$ *up to the group action*). This is precisely the notion needed to translate distributional convergence into parameter/orbit convergence.

Definition G.3 (Orbit-determining IPMs)**.** Assume a group G acts on Θ and $P_{g\cdot\theta}^X = P_\theta^X$ for all g, θ. We say $\mathcal{F}_0$ (equivalently $d_{\mathcal{F}_0}$) is *orbit-determining on the model range* if

$$d_{\mathcal{F}_0}(P_\theta^X, P_{\theta'}^X) = 0 \implies \theta' \in [(]\,\theta).$$

Equivalently, the induced pseudometric on Θ/G (Lemma G.14 below) is a genuine metric.

Remark G.4 (Why this is the right notion). If the model is nonidentifiable, then even an "ideal" metric such as total variation cannot separate parameters on the same orbit. Orbit-determiningness is the sharp replacement: $d_{\mathcal{F}_0}$ should separate *orbits* (equivalence classes), not raw parameters.

A practical orbit-separation template: finite feature embeddings. A common situation is that an IPM controls finitely many moments/features.

Lemma G.5 (Feature IPMs are norms of mean-feature differences)**.** *Let $\phi : \mathsf{X} \to \mathbb{R}^m$ be measurable and define the function class*

$$\mathcal{F}_\phi := \{f_u(x) := \langle u, \phi(x)\rangle :\; u \in \mathbb{R}^m,\; \|u\| \le 1\}.$$

Assume $\mathbb{E}_P\|\phi(X)\| < \infty$ and $\mathbb{E}_Q\|\phi(X)\| < \infty$. Then

$$d_{\mathcal{F}_\phi}(P, Q) = \big\|\mathbb{E}_P[\phi(X)] - \mathbb{E}_Q[\phi(X)]\big\|.$$

Consequently, $d_{\mathcal{F}_\phi}$ is orbit-determining on $\mathcal{P}_\Theta$ whenever the mean-feature map

$$\Theta/G \;\ni\; \bar\theta \;\longmapsto\; \mathbb{E}_{P_\theta^X}\phi(X) \in \mathbb{R}^m$$

is injective (well-definedness follows from observed invariance).

Proof. **Step 1 (Well-definedness and finiteness).** The integrability assumptions ensure the vector means $\mathbb{E}_P\phi$ and $\mathbb{E}_Q\phi$ exist in $\mathbb{R}^m$. Moreover for each $\|u\| \le 1$, $|\mathbb{E}_P\langle u, \phi\rangle| \le \mathbb{E}_P\|\phi\| < \infty$, so $d_{\mathcal{F}_\phi}(P, Q)$ is finite.

Step 2 (Compute the supremum). By linearity of expectation,

$$\mathbb{E}_P f_u - \mathbb{E}_Q f_u = \mathbb{E}_P\langle u, \phi(X)\rangle - \mathbb{E}_Q\langle u, \phi(X)\rangle = \Big\langle u, \mathbb{E}_P\phi - \mathbb{E}_Q\phi\Big\rangle.$$

Hence

$$d_{\mathcal{F}_\phi}(P, Q) = \sup_{\|u\|\le 1} \Big|\Big\langle u, \mathbb{E}_P\phi - \mathbb{E}_Q\phi\Big\rangle\Big|.$$

By Cauchy–Schwarz, $|\langle u, \Delta\rangle| \le \|u\|\,\|\Delta\| \le \|\Delta\|$ where $\Delta := \mathbb{E}_P\phi - \mathbb{E}_Q\phi$, so the supremum is at most $\|\Delta\|$.

To see equality, if $\Delta \neq 0$ choose $u^\star := \Delta/\|\Delta\|$, which satisfies $\|u^\star\| = 1$ and yields $|\langle u^\star, \Delta\rangle| = \|\Delta\|$. If $\Delta = 0$ the supremum is 0. Therefore $d_{\mathcal{F}_\phi}(P, Q) = \|\mathbb{E}_P\phi - \mathbb{E}_Q\phi\|$.

Step 3 (Orbit-determiningness criterion). If the mean-feature map is injective on Θ/G, then

$$d_{\mathcal{F}_\phi}(P_\theta^X, P_{\theta'}^X) = 0 \iff \mathbb{E}_{P_\theta^X}\phi = \mathbb{E}_{P_{\theta'}^X}\phi \iff q(\theta) = q(\theta'),$$

i.e. θ' lies in the orbit of θ. This is exactly Definition G.3. □

Remark G.6 (Feature IPMs as dual norms). The argument above is simply duality between the Euclidean norm and its dual: $\|\Delta\| = \sup_{\|u\|\le 1}\langle u, \Delta\rangle$. One can replace $\|\cdot\|$ by any norm $\|\cdot\|_*$ on $\mathbb{R}^m$ and take $\mathcal{F}_\phi$ indexed by the corresponding dual unit ball, obtaining $d_{\mathcal{F}_\phi}(P, Q) = \|\mathbb{E}_P\phi - \mathbb{E}_Q\phi\|_*$.

Local orbit-separation via transverse Jacobians. Orbit-determiningness is frequently required only locally on the stability region where EM is analyzed. The next lemma gives a standard differential criterion on a slice/section.

Lemma G.7 (Local orbit separation from a feature Jacobian on a slice)**.** *Assume there is a C^1 slice map (local section) $s : U \subseteq \mathbb{R}^p \to \Theta$ such that $q \circ s$ is a homeomorphism onto its image in Θ/G. Let $\phi : \mathsf{X} \to \mathbb{R}^m$ be measurable with $\mathbb{E}_{P_{s(u)}^X}\|\phi(X)\| < \infty$ for all $u \in U$, and define*

$$m(u) := \mathbb{E}_{P_{s(u)}^X}[\phi(X)] \in \mathbb{R}^m.$$

Assume m is differentiable on a neighborhood of $u_0 \in U$ and that its Jacobian at u_0 has full column rank, with

$$\sigma_{\min}(\nabla m(u_0)) \ge \alpha > 0.$$

Then there exists a neighborhood $U_0 \subseteq U$ of u_0 such that for all $u, v \in U_0$,

$$\|m(u) - m(v)\| \ \ge\ \frac{\alpha}{2}\,\|u - v\|.$$

Equivalently, with $\mathcal{F}_\phi$ as in Lemma G.5,

$$d_{\mathcal{F}_\phi}(P_{s(u)}^X, P_{s(v)}^X) \ \ge\ \frac{\alpha}{2}\,\|u - v\|, \qquad u, v \in U_0.$$

In particular, $d_{\mathcal{F}_\phi}$ separates points on the slice locally, hence is locally orbit-determining in the chart neighborhood.

Proof. **Step 1 (Quantitative lower bound for the linearization at u_0).** Let $A := \nabla m(u_0) \in \mathbb{R}^{m\times p}$. The condition $\sigma_{\min}(A) \ge \alpha$ means

$$\|Ah\| \ \ge\ \alpha\|h\| \qquad \forall h \in \mathbb{R}^p. \tag{352}$$

Step 2 (Control the Jacobian uniformly near u_0). Since m is differentiable in a neighborhood of u_0, ∇m exists there. Assume (as stated) that ∇m is continuous at u_0; this is automatic if m is C^1, and it also holds under standard assumptions used to justify differentiation under the expectation defining m. By continuity, choose a neighborhood U_0 of u_0 such that

$$\|\nabla m(w) - A\|_{\mathrm{op}} \ \le\ \alpha/2 \qquad \forall w \in U_0. \tag{353}$$

Step 3 (Mean value theorem in integral form). Fix $u, v \in U_0$ and consider the segment $w(t) := v + t(u - v)$ for $t \in [0, 1]$. For vector-valued differentiable maps, the fundamental theorem of calculus gives

$$m(u) - m(v) = \int_0^1 \nabla m(w(t))\,(u - v)\,dt. \tag{354}$$

Step 4 (Lower bound by the linear part minus a perturbation). Add and subtract $A(u-v)$ inside the integral:

$$m(u)-m(v)=A(u-v)+\int_0^1 \big(\nabla m(w(t))-A\big)(u-v)\,dt.$$

Take norms and use the triangle inequality in the form $\|a+b\|\ge\|a\|-\|b\|$:

$$\begin{aligned}\|m(u)-m(v)\| &\ge \|A(u-v)\|-\left\|\int_0^1 (\nabla m(w(t))-A)(u-v)\,dt\right\|\\ &\ge \|A(u-v)\|-\int_0^1 \|(\nabla m(w(t))-A)(u-v)\|\,dt\\ &\ge \|A(u-v)\|-\int_0^1 \|\nabla m(w(t))-A\|_{\mathrm{op}}\,\|u-v\|\,dt.\end{aligned}$$

By (352), $\|A(u-v)\|\ge\alpha\|u-v\|$. By (353), $\|\nabla m(w(t))-A\|_{\mathrm{op}}\le\alpha/2$ for all t. Therefore

$$\|m(u)-m(v)\|\ge\alpha\|u-v\|-\int_0^1(\alpha/2)\|u-v\|\,dt=\frac{\alpha}{2}\|u-v\|.$$

Step 5 (Translate to an IPM lower bound). By Lemma G.5,

$$d_{\mathcal{F}_\phi}(P^X_{s(u)},P^X_{s(v)})=\|m(u)-m(v)\|,$$

so the same lower bound holds for the IPM.

Step 6 (Local orbit-determiningness on the slice). If $d_{\mathcal{F}_\phi}(P^X_{s(u)},P^X_{s(v)})=0$ for $u,v\in U_0$, then the inequality forces $\|u-v\|=0$, hence $u=v$ and $s(u)=s(v)$. Since $q\circ s$ is injective on U_0 (homeomorphism onto its image), this implies the corresponding orbits are distinct unless $u=v$. Thus the IPM separates orbits locally in the chart neighborhood. □

Remark G.8 (Where continuity of ∇m comes from in EM applications). In EM-style settings $m(u)=\mathbb{E}_{P^X_{s(u)}}\phi(X)$ is a parameterized expectation. One typically justifies differentiability and continuity of ∇m by dominated convergence: if $u\mapsto p_{s(u)}(x)$ is smooth and $\phi(x)$ has an integrable envelope uniform over u in a neighborhood, then m is C^1 with $\nabla m(u)=\mathbb{E}[\nabla_u(\phi(X)\text{ under }P^X_{s(u)})]$. Lemma D.6 records the same dominated-convergence mechanism for second derivatives.

Remark G.9 (Interpretation: a local bi-Lipschitz embedding of the quotient). The lemma says that near u_0 the map $u\mapsto m(u)$ is locally injective with a quantitative inverse Lipschitz constant $\alpha/2$. Composed with the local section s, this yields a locally injective map on Θ/G. This is the precise sense in which finitely many moments/features can identify orbits *locally*: the feature mean embedding is transverse to the orbits in the neighborhood where EM contracts.

G.2. **Statistical estimation of IPMs and induced rates.**
Empirical IPMs. Given i.i.d. $X_1,\dots,X_n\sim P$, write $P_n:=\widehat{P}_n:=\frac{1}{n}\sum_{i=1}^n\delta_{X_i}$ and define the plug-in IPM

$$d_{\mathcal{F}_0}(P_n,P)\;:=\;\sup_{f\in\mathcal{F}_0}\left|\frac{1}{n}\sum_{i=1}^n f(X_i)-\mathbb{E}_P f(X)\right|\;=\;\sup_{f\in\mathcal{F}_0}\big|(P_n-P)f\big|.$$

Its statistical behavior is governed by: (i) the "size" of $\mathcal{F}_0$ (Rademacher complexity, VC/entropy, chaining), and (ii) envelope/tail conditions ensuring concentration.

Throughout, whenever we take $\sup_{f\in\mathcal{F}_0}$ inside expectations or probabilities, we tacitly impose a standard *measurability/separability hypothesis* (e.g. Assumption E.1) so that the supremum

is measurable; otherwise, one may interpret all expectations/probabilities as *outer* expectation/probability, which yields identical inequalities.

A generic complexity–concentration template. The next lemma packages the two standard steps: *symmetrization* (to control expectations via Rademacher averages) and *Talagrand–Bousquet concentration* (to upgrade to high probability). We state it in a form tailored to empirical IPMs.

Lemma G.10 (Symmetrization and Bousquet concentration for bounded IPMs)**.** *Assume* $\mathcal{F}_0 \subseteq \{f : \mathsf{X} \to \mathbb{R}\}$ *has bounded envelope*

$$\sup_{f\in\mathcal{F}_0} \|f\|_\infty \;\le\; b < \infty.$$

Let $\epsilon_1, \dots, \epsilon_n$ *be i.i.d. Rademacher signs independent of* $X_{1:n}$ *and define the conditional empirical Rademacher complexity*

$$\widehat{\mathfrak{R}}_n(\mathcal{F}_0; X_{1:n}) := \mathbb{E}_\epsilon\Big[\sup_{f\in\mathcal{F}_0} \frac{1}{n}\sum_{i=1}^n \epsilon_i f(X_i) \;\Big|\; X_{1:n}\Big].$$

Let the variance proxy *be*

$$v \;:=\; \sup_{f\in\mathcal{F}_0} \mathrm{Var}\big(f(X)\big) \;\le\; b^2.$$

Then:

*(1) (*Expectation bound*)*

$$\mathbb{E}\big[d_{\mathcal{F}_0}(P_n, P)\big] \;\le\; 2\,\mathbb{E}\big[\widehat{\mathfrak{R}}_n(\mathcal{F}_0^{\pm}; X_{1:n})\big], \qquad \mathcal{F}_0^{\pm} := \mathcal{F}_0 \cup (-\mathcal{F}_0).$$

In particular, if $\mathcal{F}_0$ *is symmetric (*$\mathcal{F}_0 = -\mathcal{F}_0$*), then* $\mathbb{E}[d_{\mathcal{F}_0}(P_n, P)] \le 2\,\mathbb{E}[\widehat{\mathfrak{R}}_n(\mathcal{F}_0; X_{1:n})]$.

*(2) (*High-probability bound*) for all* $t > 0$*, with probability at least* $1 - e^{-t}$*,*

$$d_{\mathcal{F}_0}(P_n, P) \;\le\; \mathbb{E}\big[d_{\mathcal{F}_0}(P_n, P)\big] \;+\; \sqrt{\frac{2vt}{n}} \;+\; \frac{bt}{3n}. \tag{355}$$

In particular, using $v \le b^2$*,*

$$d_{\mathcal{F}_0}(P_n, P) \;\le\; \mathbb{E}\big[d_{\mathcal{F}_0}(P_n, P)\big] \;+\; b\sqrt{\frac{2t}{n}} \;+\; \frac{bt}{3n}.$$

Combining with (i) yields the common "complexity + concentration" form

$$d_{\mathcal{F}_0}(P_n, P) \;\le\; 2\,\mathbb{E}\big[\widehat{\mathfrak{R}}_n(\mathcal{F}_0^{\pm}; X_{1:n})\big] \;+\; \sqrt{\frac{2vt}{n}} \;+\; \frac{bt}{3n}.$$

Proof. **Step 0 (Absolute supremum as a one-sided supremum over** $\mathcal{F}_0^{\pm}$**).** By definition,

$$d_{\mathcal{F}_0}(P_n, P) = \sup_{f\in\mathcal{F}_0} \big|(P_n - P)f\big| = \sup_{g\in\mathcal{F}_0^{\pm}} (P_n - P)g.$$

Thus it suffices to work with the one-sided supremum functional over $\mathcal{F}_0^{\pm}$.

Step 1 (Expectation: symmetrization). Let $X_1', \dots, X_n' \overset{iid}{\sim} P$ be an independent ghost sample, with empirical measure $P_n' := \frac{1}{n}\sum_{i=1}^n \delta_{X_i'}$. For each $g \in \mathcal{F}_0^{\pm}$,

$$(P_n - P)g \;=\; (P_n - P_n')g \;+\; (P_n' - P)g.$$

Taking $\mathbb{E}[\cdot \mid X_{1:n}]$ and using $\mathbb{E}[P_n' g \mid X_{1:n}] = Pg$ gives

$$\mathbb{E}\big[(P_n - P)g \mid X_{1:n}\big] = \mathbb{E}\big[(P_n - P_n')g \mid X_{1:n}\big].$$

Now apply Jensen to the convex functional $u(\cdot) \mapsto \sup_{g\in\mathcal{F}_0^{\pm}} u(g)$:

$$\mathbb{E}\Big[\sup_{g\in\mathcal{F}_0^{\pm}}(P_n-P)g\Big] = \mathbb{E}\Big[\sup_{g\in\mathcal{F}_0^{\pm}}\mathbb{E}\big[(P_n-P_n')g \mid X_{1:n}\big]\Big] \le \mathbb{E}\Big[\sup_{g\in\mathcal{F}_0^{\pm}}(P_n-P_n')g\Big].$$

Conditioning on $(X_{1:n}, X'_{1:n})$, the random variables $\{g(X_i)-g(X_i')\}_{i=1}^n$ are symmetrized by Rademachers:

$$\mathbb{E}\Big[\sup_{g\in\mathcal{F}_0^{\pm}}(P_n-P_n')g \;\Big|\; X_{1:n}, X'_{1:n}\Big] = \mathbb{E}_\epsilon\Big[\sup_{g\in\mathcal{F}_0^{\pm}}\frac{1}{n}\sum_{i=1}^n \epsilon_i\big(g(X_i)-g(X_i')\big)\Big].$$

Use $\sup(a+b)\le \sup a + \sup b$:

$$\mathbb{E}_\epsilon\Big[\sup_g \frac{1}{n}\sum_i \epsilon_i(g(X_i)-g(X_i'))\Big] \le \mathbb{E}_\epsilon\Big[\sup_g \frac{1}{n}\sum_i \epsilon_i g(X_i)\Big] + \mathbb{E}_\epsilon\Big[\sup_g \frac{1}{n}\sum_i (-\epsilon_i) g(X_i')\Big].$$

The two terms are equal after taking expectation over $(X_{1:n}, X'_{1:n})$ (symmetry of ϵ and i.i.d. samples), so

$$\mathbb{E}\Big[\sup_{g\in\mathcal{F}_0^{\pm}}(P_n-P)g\Big] \le 2\,\mathbb{E}\Big[\sup_{g\in\mathcal{F}_0^{\pm}}\frac{1}{n}\sum_{i=1}^n \epsilon_i g(X_i)\Big] = 2\,\mathbb{E}\big[\widehat{\mathfrak{R}}_n(\mathcal{F}_0^{\pm}; X_{1:n})\big],$$

which proves (i).

Step 2 (High probability: apply Bousquet to the absolute supremum). Define the random variable

$$Z \;:=\; \sup_{g\in\mathcal{F}_0^{\pm}}(P_n-P)g \;=\; d_{\mathcal{F}_0}(P_n,P).$$

Since $\|g\|_\infty \le b$ for all $g\in\mathcal{F}_0^{\pm}$ and $\mathrm{Var}(g(X))\le v$, Lemma E.15 applies and yields: for all $t>0$, with probability at least $1-e^{-t}$,

$$Z \le \mathbb{E}[Z] + \sqrt{\frac{2vt}{n}} + \frac{bt}{3n}.$$

This is exactly (355). The simplification $v\le b^2$ is immediate. □

Remark G.11 (A common tightening and a data-dependent alternative). Lemma G.10 cleanly separates: (a) complexity (through $\mathbb{E}Z$) and (b) concentration. In many analyses one further bounds $\mathbb{E}Z$ by deterministic entropy integrals (Appendix E.3). If one prefers *data-dependent* bounds, there are also standard results of the form

$$d_{\mathcal{F}_0}(P_n,P) \;\lesssim\; 2\,\widehat{\mathfrak{R}}_n(\mathcal{F}_0^{\pm}; X_{1:n}) \;+\; b\sqrt{\frac{t}{n}}$$

with probability $\ge 1-e^{-t}$ (obtained by combining symmetrization with concentration around the empirical Rademacher complexity). We keep the present form because it is modular and pairs directly with Bousquet/Talagrand.

Vector-valued test statistics. Many applications (e.g. gradient/Hessian control) produce vector classes $\mathcal{H} \subseteq \{h : \mathsf{X} \to \mathbb{R}^d\}$ and consider $\sup_{h\in\mathcal{H}} \|(P_n - P)h\|$. Appendix E.1 reduces this to finitely many scalar IPMs by sphere nets:

$$\sup_{h\in\mathcal{H}} \|(P_n - P)h\| \;\le\; \frac{1}{1-\eta} \max_{v\in\mathcal{N}_\eta} \sup_{h\in\mathcal{H}} (P_n - P)\langle v, h\rangle.$$

One then applies Lemma G.10 to each induced scalar class and union-bounds over the net.

Two canonical estimable IPMs with explicit bounds.
Finite-dimensional feature IPMs. Lemma G.5 identifies $d_{\mathcal{F}_\phi}(P_n, P)$ with a Euclidean mean deviation.

Lemma G.12 (Estimation of feature IPMs)**.** *Let $\phi : \mathsf{X} \to \mathbb{R}^m$ be measurable with $\|\phi(x)\| \le B$ for all x. Let $\mathcal{F}_\phi := \{x \mapsto \langle u, \phi(x)\rangle : \ \|u\| \le 1\}$. Then for i.i.d. $X_1, \dots, X_n \sim P$,*

$$d_{\mathcal{F}_\phi}(P_n, P) = \big\|(P_n - P)\phi\big\|.$$

Moreover:

*(1) (*Expectation*) $\mathbb{E}[d_{\mathcal{F}_\phi}(P_n, P)] \le 2B/\sqrt{n}$.*
*(2) (*Concentration*) for all $t > 0$, with probability at least $1 - e^{-t}$,*

$$d_{\mathcal{F}_\phi}(P_n, P) \;\le\; \mathbb{E}[d_{\mathcal{F}_\phi}(P_n, P)] \;+\; B\sqrt{\frac{2t}{n}} \;\le\; \frac{2B}{\sqrt{n}} \;+\; B\sqrt{\frac{2t}{n}}.$$

Proof. **Step 1 (Identity with a Euclidean mean deviation).** By Lemma G.5,

$$d_{\mathcal{F}_\phi}(P_n, P) = \big\|\mathbb{E}_{P_n}\phi - \mathbb{E}_P\phi\big\| = \|(P_n - P)\phi\|.$$

Step 2 (Expectation bound via symmetrization and a Khintchine–Jensen step). Apply symmetrization (Lemma E.6) to the *vector* mean deviation via scalarization by duality:

$$\mathbb{E}\|(P_n - P)\phi\| = \mathbb{E} \sup_{\|u\|\le 1} \big\langle u, (P_n - P)\phi\big\rangle = \mathbb{E} \sup_{\|u\|\le 1} (P_n - P)\langle u, \phi\rangle \le 2\,\mathbb{E} \sup_{\|u\|\le 1} \frac{1}{n}\sum_{i=1}^n \epsilon_i \langle u, \phi(X_i)\rangle.$$

The inner supremum is again a dual norm:

$$\sup_{\|u\|\le 1} \frac{1}{n}\sum_{i=1}^n \epsilon_i \langle u, \phi(X_i)\rangle = \Big\|\frac{1}{n}\sum_{i=1}^n \epsilon_i \phi(X_i)\Big\|.$$

Thus

$$\mathbb{E}\|(P_n - P)\phi\| \le 2\,\mathbb{E}\Big\|\frac{1}{n}\sum_{i=1}^n \epsilon_i \phi(X_i)\Big\|.$$

Now condition on $X_{1:n}$ and use Jensen in the form $\mathbb{E}\|S\| \le (\mathbb{E}\|S\|^2)^{1/2}$:

$$\mathbb{E}_\epsilon\Big\|\frac{1}{n}\sum_{i=1}^n \epsilon_i \phi(X_i)\Big\| \le \left(\mathbb{E}_\epsilon\Big\|\frac{1}{n}\sum_{i=1}^n \epsilon_i \phi(X_i)\Big\|^2\right)^{1/2}.$$

Compute the second moment (cross-terms vanish because $\mathbb{E}[\epsilon_i\epsilon_j] = 0$ for $i \ne j$):

$$\mathbb{E}_\epsilon\Big\|\sum_{i=1}^n \epsilon_i a_i\Big\|^2 = \mathbb{E}_\epsilon \sum_{i,j} \epsilon_i\epsilon_j \langle a_i, a_j\rangle = \sum_{i=1}^n \|a_i\|^2.$$

With $a_i := \phi(X_i)$ this gives

$$\mathbb{E}_\epsilon\Big\|\frac{1}{n}\sum_{i=1}^n \epsilon_i\phi(X_i)\Big\| \le \frac{1}{n}\Big(\sum_{i=1}^n \|\phi(X_i)\|^2\Big)^{1/2} \le \frac{B}{\sqrt{n}}.$$

Taking expectation over $X_{1:n}$ yields $\mathbb{E}[d_{\mathcal{F}_\phi}(P_n, P)] \le 2B/\sqrt{n}$.

Step 3 (Concentration via bounded differences / McDiarmid). Consider the function of the sample

$$F(X_1, \dots, X_n) := \|(P_n - P)\phi\| = \left\|\frac{1}{n}\sum_{i=1}^n \phi(X_i) - \mathbb{E}\phi(X)\right\|.$$

If we replace X_i by X_i', then $P_n\phi$ changes by $\frac{1}{n}(\phi(X_i) - \phi(X_i'))$, so

$$|F(X_1, \dots, X_n) - F(X_1, \dots, X_i', \dots, X_n)| \le \frac{1}{n}\|\phi(X_i) - \phi(X_i')\| \le \frac{2B}{n}.$$

Thus McDiarmid's inequality gives: for all $t > 0$, with probability at least $1 - e^{-t}$,

$$F \le \mathbb{E}[F] + B\sqrt{\frac{2t}{n}}.$$

Insert the expectation bound from Step 2. □

Kernel IPMs (MMD).. For bounded kernels, MMD is naturally Hilbert-valued mean estimation and admits dimension-free root-n bounds.

Lemma G.13 (MMD estimation under a bounded kernel)**.** *Let $(\mathsf{X}, \mathcal{F})$ be a measurable space and let $k : \mathsf{X} \times \mathsf{X} \to \mathbb{R}$ be a measurable positive definite kernel with RKHS $\mathcal{H}_k$. Assume the diagonal is uniformly bounded:*

$$\sup_{x\in\mathsf{X}} k(x, x) \;\le\; \kappa^2 < \infty.$$

Let

$$\mathcal{F}_0 := \{f \in \mathcal{H}_k : \; \|f\|_{\mathcal{H}_k} \le 1\}, \qquad d_k := d_{\mathcal{F}_0}.$$

Assume (as is standard in learning theory, and automatic e.g. when X is standard Borel and k is measurable) that the canonical feature map $\Phi : \mathsf{X} \to \mathcal{H}_k$, $\Phi(x) := k(x, \cdot)$ is strongly measurable, so Bochner expectations below are well-defined. Then for i.i.d. $X_1, \dots, X_n \sim P$ and $P_n := \frac{1}{n}\sum_{i=1}^n \delta_{X_i}$:

(i) (Mean embedding identity). *The mean element $\mu_P := \mathbb{E}[\Phi(X)] \in \mathcal{H}_k$ exists and*

$$d_k(P, Q) = \|\mu_P - \mu_Q\|_{\mathcal{H}_k} \qquad \text{whenever } \mu_P, \mu_Q \text{ exist.}$$

In particular, $\mu_{P_n} = \frac{1}{n}\sum_{i=1}^n \Phi(X_i)$ and $d_k(P_n, P) = \|\mu_{P_n} - \mu_P\|_{\mathcal{H}_k}$.

(ii) (Expectation). $\mathbb{E}[d_k(P_n, P)] \le 2\kappa/\sqrt{n}$. *(Indeed, one can sharpen this to $\mathbb{E}[d_k(P_n, P)] \le \kappa/\sqrt{n}$.)*

(iii) (Concentration). *For all $t > 0$, with probability at least $1 - e^{-t}$,*

$$d_k(P_n, P) \;\le\; \mathbb{E}[d_k(P_n, P)] + \kappa\sqrt{\frac{2t}{n}} \;\le\; \frac{2\kappa}{\sqrt{n}} + \kappa\sqrt{\frac{2t}{n}}.$$

Proof. **Step 0 (Basic RKHS facts and integrability).** For each $x \in \mathsf{X}$, $\Phi(x) = k(x, \cdot) \in \mathcal{H}_k$ and the reproducing property gives

$$f(x) = \langle f, \Phi(x)\rangle_{\mathcal{H}_k} \qquad \forall f \in \mathcal{H}_k.$$

Moreover,

$$\|\Phi(x)\|_{\mathcal{H}_k}^2 = \langle \Phi(x), \Phi(x)\rangle_{\mathcal{H}_k} = k(x,x) \le \kappa^2,$$

so $\|\Phi(X)\|_{\mathcal{H}_k} \le \kappa$ almost surely. Hence $\mathbb{E}\|\Phi(X)\|_{\mathcal{H}_k} < \infty$, and the Bochner mean element $\mu_P := \mathbb{E}[\Phi(X)] \in \mathcal{H}_k$ exists. Likewise $\mu_{P_n} = \frac{1}{n}\sum_{i=1}^n \Phi(X_i)$ exists almost surely.

Step 1 (Identity $d_k(P,Q) = \|\mu_P - \mu_Q\|$). Fix probability measures P, Q such that μ_P, μ_Q exist. For any $f \in \mathcal{H}_k$,

$$\mathbb{E}_P f - \mathbb{E}_Q f = \mathbb{E}_P\langle f, \Phi(X)\rangle_{\mathcal{H}_k} - \mathbb{E}_Q\langle f, \Phi(X)\rangle_{\mathcal{H}_k} = \Big\langle f, \mathbb{E}_P\Phi(X) - \mathbb{E}_Q\Phi(X)\Big\rangle_{\mathcal{H}_k} = \langle f, \mu_P - \mu_Q\rangle_{\mathcal{H}_k},$$

where we used linearity and continuity of the inner product to exchange $\mathbb{E}$ and $\langle\cdot,\cdot\rangle$ (justified by the bound $|\langle f, \Phi(X)\rangle| \le \|f\|\,\|\Phi(X)\| \le \kappa\|f\|$). Therefore,

$$d_k(P,Q) = \sup_{\|f\|_{\mathcal{H}_k}\le 1} \big|\langle f, \mu_P - \mu_Q\rangle_{\mathcal{H}_k}\big| = \|\mu_P - \mu_Q\|_{\mathcal{H}_k},$$

by Hilbert-space duality: $\sup_{\|f\|\le 1}|\langle f, h\rangle| = \|h\|$, achieved at $f = h/\|h\|$ when $h \ne 0$. Taking $Q = P_n$ yields the claimed identity for $d_k(P_n, P)$.

Step 2 (Expectation bound: symmetrization route). Let $\Phi_i := \Phi(X_i)$ and $\mu := \mu_P = \mathbb{E}[\Phi(X)]$. Then $\mu_{P_n} - \mu = \frac{1}{n}\sum_{i=1}^n(\Phi_i - \mu)$, so by Step 1,

$$\mathbb{E}[d_k(P_n, P)] = \mathbb{E}\|\mu_{P_n} - \mu\|.$$

Introduce i.i.d. Rademacher signs $\epsilon_1, \dots, \epsilon_n$, independent of $X_{1:n}$. A standard symmetrization inequality for Banach-valued means (apply the usual ghost-sample symmetrization to the scalar class $\{x \mapsto \langle h, \Phi(x)\rangle : \|h\| \le 1\}$ and then take the supremum over $\|h\| \le 1$) yields

$$\mathbb{E}\|\mu_{P_n} - \mu\| \le 2\,\mathbb{E}\left\|\frac{1}{n}\sum_{i=1}^n \epsilon_i\Phi_i\right\|.$$

Condition on $X_{1:n}$. In a Hilbert space,

$$\mathbb{E}_\epsilon\left\|\sum_{i=1}^n \epsilon_i\Phi_i\right\|^2 = \mathbb{E}_\epsilon\left\langle \sum_{i=1}^n \epsilon_i\Phi_i,\ \sum_{j=1}^n \epsilon_j\Phi_j\right\rangle = \sum_{i=1}^n \|\Phi_i\|^2 \quad (\text{since } \mathbb{E}[\epsilon_i\epsilon_j] = \mathbf{1}\{i=j\}).$$

By Jensen,

$$\mathbb{E}_\epsilon\left\|\frac{1}{n}\sum_{i=1}^n \epsilon_i\Phi_i\right\| \le \left(\mathbb{E}_\epsilon\left\|\frac{1}{n}\sum_{i=1}^n \epsilon_i\Phi_i\right\|^2\right)^{1/2} = \frac{1}{n}\left(\sum_{i=1}^n \|\Phi_i\|^2\right)^{1/2} \le \frac{1}{n}\left(\sum_{i=1}^n \kappa^2\right)^{1/2} = \frac{\kappa}{\sqrt{n}}.$$

Taking expectation over $X_{1:n}$ gives $\mathbb{E}[d_k(P_n, P)] \le 2\kappa/\sqrt{n}$.

(Sharper constant, optional). One may avoid symmetrization and compute the second moment directly:

$$\mathbb{E}\|\mu_{P_n} - \mu\|^2 = \mathbb{E}\left\|\frac{1}{n}\sum_{i=1}^n(\Phi_i - \mu)\right\|^2 = \frac{1}{n^2}\sum_{i=1}^n \mathbb{E}\|\Phi_i - \mu\|^2 \le \frac{1}{n^2}\cdot n\,\mathbb{E}\|\Phi(X)\|^2 \le \frac{\kappa^2}{n},$$

hence $\mathbb{E}\|\mu_{P_n} - \mu\| \le (\mathbb{E}\|\mu_{P_n} - \mu\|^2)^{1/2} \le \kappa/\sqrt{n}$.

Step 3 (Concentration via bounded differences / McDiarmid). Define the functional of the sample

$$F(X_{1:n}) \ :=\ \|\mu_{P_n} - \mu\|_{\mathcal{H}_k} = \left\|\frac{1}{n}\sum_{i=1}^n \Phi(X_i) - \mu\right\|.$$

Let X_i' be an independent copy of X_i and write $X^{(i)} := (X_1, \dots, X_{i-1}, X_i', X_{i+1}, \dots, X_n)$. Then

$$\mu_{P_n}(X_{1:n}) - \mu_{P_n}(X^{(i)}) = \frac{1}{n}\big(\Phi(X_i) - \Phi(X_i')\big),$$

so by the reverse triangle inequality,

$$|F(X_{1:n}) - F(X^{(i)})| \le \Big\|\mu_{P_n}(X_{1:n}) - \mu_{P_n}(X^{(i)})\Big\| = \frac{1}{n}\|\Phi(X_i) - \Phi(X_i')\| \le \frac{1}{n}\big(\|\Phi(X_i)\| + \|\Phi(X_i')\|\big) \le \frac{2\kappa}{n}.$$

Thus F satisfies bounded differences with constants $c_i = 2\kappa/n$. McDiarmid's inequality gives, for all $t > 0$, with probability at least $1 - e^{-t}$,

$$F \le \mathbb{E}[F] + \sqrt{\frac{t}{2}\sum_{i=1}^{n} c_i^2} = \mathbb{E}[F] + \sqrt{\frac{t}{2} \cdot n \cdot \frac{4\kappa^2}{n^2}} = \mathbb{E}[F] + \kappa\sqrt{\frac{2t}{n}}.$$

Substitute $\mathbb{E}[F] = \mathbb{E}[d_k(P_n, P)]$ and apply Step 2. □

Qualitative comparison of estimability across IPMs. The lemmas above illustrate the standard geometry–statistics tradeoff:

- Smaller/more regular test classes $\mathcal{F}_0$ yield smaller empirical-process complexity and hence tighter estimation error for $d_{\mathcal{F}_0}(P_n, P)$.
- But smaller $\mathcal{F}_0$ typically weakens separation on the model range (or the modulus linking $d_{\mathcal{F}_0}(P_\theta, P_{\theta'})$ to orbit/parameter distances).

In EM analyses, one typically chooses $\mathcal{F}_0$ to be just strong enough to be (locally) orbit-determining on the stability region, while remaining estimable at the sample size and dimension of interest.

G.3. **Compatibility with quotient geometry and induced moduli.**

Descending an IPM to the quotient. Assume the observed marginal laws are G-invariant:

$$P^X_{g\cdot\theta} = P^X_\theta \qquad \forall g \in G,\ \forall \theta \in \Theta. \tag{356}$$

Let D be any pseudometric on probability laws on $(\mathsf{X}, \mathcal{F})$ (e.g. an IPM $d_{\mathcal{F}_0}$), possibly taking values in $[0, \infty]$. Throughout this subsection we work on a region where $\mathsf{D}(P^X_\theta, P^X_{\theta'}) < \infty$ for the pairs of parameters considered, so that the triangle inequality is meaningful.
Orbit-invariance of $\theta \mapsto P^X_\theta$ implies descent. Define the induced parameter-level pseudometric

$$\mathsf{D}_\Theta(\theta, \theta') \ := \ \mathsf{D}(P^X_\theta, P^X_{\theta'}).$$

By (356), D_Θ is constant on orbits in each argument: for all $g, h \in G$,

$$\mathsf{D}_\Theta(g\cdot\theta, h\cdot\theta') = \mathsf{D}(P^X_{g\cdot\theta}, P^X_{h\cdot\theta'}) = \mathsf{D}(P^X_\theta, P^X_{\theta'}) = \mathsf{D}_\Theta(\theta, \theta').$$

Hence D_Θ canonically defines a pseudometric on the orbit space Θ/G.

Lemma G.14 (Well-defined quotient distance induced by an IPM)**.** *Let $q : \Theta \to \Theta/G$ be the orbit map. Under (356), the function*

$$\overline{\mathsf{D}}(\bar\theta, \bar\theta') \ := \ \mathsf{D}(P^X_\theta, P^X_{\theta'}), \qquad \bar\theta = q(\theta),\ \bar\theta' = q(\theta'),$$

is well-defined on Θ/G (i.e. independent of the chosen representatives) and is a pseudometric there.

Proof. **Step 1 (Well-definedness).** Let $\theta, \tilde{\theta} \in \Theta$ with $q(\theta) = q(\tilde{\theta})$. Then $\tilde{\theta} = g \cdot \theta$ for some $g \in G$. Using (356),

$$\mathsf{D}(P^X_{\tilde{\theta}}, P^X_{\theta'}) = \mathsf{D}(P^X_{g\cdot\theta}, P^X_{\theta'}) = \mathsf{D}(P^X_{\theta}, P^X_{\theta'}).$$

The same argument applies in the second argument. Thus $\overline{\mathsf{D}}$ does not depend on representatives.

Step 2 (Pseudometric properties). Nonnegativity and the identity $\overline{\mathsf{D}}(\bar{\theta}, \bar{\theta}) = 0$ follow from the corresponding properties of D. Symmetry follows from symmetry of D. For the triangle inequality, pick representatives $\theta, \theta', \theta''$ of $\bar{\theta}, \bar{\theta}', \bar{\theta}''$ and use

$$\overline{\mathsf{D}}(\bar{\theta}, \bar{\theta}'') = \mathsf{D}(P^X_{\theta}, P^X_{\theta''}) \le \mathsf{D}(P^X_{\theta}, P^X_{\theta'}) + \mathsf{D}(P^X_{\theta'}, P^X_{\theta''}) = \overline{\mathsf{D}}(\bar{\theta}, \bar{\theta}') + \overline{\mathsf{D}}(\bar{\theta}', \bar{\theta}'').$$

□

Equivalent "inf-over-representatives" form. Even though representative-independence makes the following trivial under (356), it is sometimes conceptually useful to note:

$$\overline{\mathsf{D}}(\bar{\theta}, \bar{\theta}') = \inf\left\{\mathsf{D}(P^X_{\vartheta}, P^X_{\vartheta'}) : \ q(\vartheta) = \bar{\theta}, \ q(\vartheta') = \bar{\theta}'\right\}.$$

Under (356), the infimum is attained by *any* representatives and has the same value.

G.4. **Compatibility with quotient geometry and induced moduli.**

Descending an IPM to the quotient. Assume the model has a G-symmetry on the observed marginal laws:

$$P^X_{g\cdot\theta} = P^X_{\theta} \qquad \forall g \in G, \ \forall \theta \in \Theta. \tag{357}$$

Let $q : \Theta \to \Theta/G$ denote the orbit map, $q(\theta) =: \bar{\theta}$.

Lemma G.15 (Well-defined quotient distance induced by a law-level pseudometric)**.** *Let* D *be any pseudometric on probability measures such that* $\mathsf{D}(P^X_{\theta}, P^X_{\theta'})$ *is well-defined (possibly* $+\infty$*) for the pairs* (θ, θ') *under consideration. Under* (356)*, the function*

$$\overline{\mathsf{D}}(\bar{\theta}, \bar{\theta}') \ := \ \mathsf{D}(P^X_{\theta}, P^X_{\theta'}), \qquad \bar{\theta} = q(\theta), \ \bar{\theta}' = q(\theta'),$$

is well-defined on Θ/G *and is a pseudometric there.*

Proof. Well-definedness. Suppose $q(\theta) = q(\tilde{\theta})$. Then $\tilde{\theta} = g \cdot \theta$ for some $g \in G$, hence by (356), $P^X_{\tilde{\theta}} = P^X_{g\cdot\theta} = P^X_{\theta}$. Therefore

$$\mathsf{D}(P^X_{\tilde{\theta}}, P^X_{\theta'}) = \mathsf{D}(P^X_{\theta}, P^X_{\theta'}).$$

The same argument applies to the second argument, proving that the right-hand side depends only on the orbits.

Pseudometric properties. Symmetry and the triangle inequality transfer directly from D because $\overline{\mathsf{D}}$ is defined by composing D with the model map $\theta \mapsto P^X_{\theta}$. □

Lemma G.16 (Orbit separation $\Longleftrightarrow$ metric on Θ/G)**.** *Assume* (356) *and let* $\overline{\mathsf{D}}$ *be defined as in Lemma G.14. Then* $\overline{\mathsf{D}}$ *is a genuine metric on* Θ/G *(i.e. it separates points) if and only if*

$$\mathsf{D}(P^X_{\theta}, P^X_{\theta'}) = 0 \implies \theta' \in [(]\, \theta). \tag{358}$$

Proof. By Lemma G.14, $\overline{\mathsf{D}}$ is a pseudometric on Θ/G, so the only issue is separation.

($\Rightarrow$). Assume $\overline{\mathsf{D}}$ is a metric. If $\mathsf{D}(P^X_{\theta}, P^X_{\theta'}) = 0$, then by definition of $\overline{\mathsf{D}}$,

$$\overline{\mathsf{D}}\big(q(\theta), q(\theta')\big) = 0.$$

Since $\overline{\mathsf{D}}$ separates points on Θ/G, it follows that $q(\theta) = q(\theta')$, i.e. $\theta' \in [(]\, \theta)$.

(⇐). Assume (358). If $\overline{\mathsf{D}}(\bar\theta,\bar\theta')=0$, choose representatives θ,θ' with $q(\theta)=\bar\theta$ and $q(\theta')=\bar\theta'$. Then

$$0=\overline{\mathsf{D}}(\bar\theta,\bar\theta')=\mathsf{D}(P_\theta^X,P_{\theta'}^X),$$

so by (358) we have $\theta'\in[(]\,\theta)$, hence $q(\theta')=q(\theta)$ and therefore $\bar\theta'=\bar\theta$. Thus $\overline{\mathsf{D}}$ separates points, i.e. it is a metric. □

Remark G.17 (What the quotient pseudometric measures). Under (356), the model map $\theta\mapsto P_\theta^X$ is constant on orbits and therefore factors through Θ/G. The quantity $\overline{\mathsf{D}}(\bar\theta,\bar\theta')$ is simply the law-level distance D evaluated on the corresponding observed laws, hence it is the intrinsic statistical distance between orbit-parameters.

Upper moduli: continuity of the moment map. Let $\mathsf{D}=d_{\mathcal{F}_0}$ be an IPM induced by a scalar class $\mathcal{F}_0$:

$$d_{\mathcal{F}_0}(P,Q):=\sup_{f\in\mathcal{F}_0}\big|\mathbb{E}_P f-\mathbb{E}_Q f\big|.$$

Lemma G.18 (Upper modulus from uniform Lipschitz dependence of expectations)**.** *Fix any pseudometric d_Θ on Θ and assume there exists $L<\infty$ such that*

$$\sup_{f\in\mathcal{F}_0}\big|\mathbb{E}_{P_\theta^X}f-\mathbb{E}_{P_{\theta'}^X}f\big|\ \le\ L\,d_\Theta(\theta,\theta')\qquad\forall\theta,\theta'\in\Theta. \tag{359}$$

Then for all $\theta,\theta'\in\Theta$,

$$d_{\mathcal{F}_0}(P_\theta^X,P_{\theta'}^X)\ \le\ L\,d_\Theta(\theta,\theta').$$

If, moreover, the observed invariance (356) holds and d_G is the orbit pseudometric induced by d_Θ,

$$d_G(\theta,\theta'):=\inf_{g\in G}d_\Theta(\theta,g\cdot\theta'),$$

then

$$d_{\mathcal{F}_0}(P_\theta^X,P_{\theta'}^X)\ \le\ L\,d_G(\theta,\theta').$$

Proof. The first claim is immediate from the definition of $d_{\mathcal{F}_0}$ as the supremum of the left-hand side of (359).

For the orbit bound, fix θ,θ' and any $g\in G$. By (356), $P_{g\cdot\theta'}^X=P_{\theta'}^X$, hence

$$d_{\mathcal{F}_0}(P_\theta^X,P_{\theta'}^X)=d_{\mathcal{F}_0}(P_\theta^X,P_{g\cdot\theta'}^X)\le L\,d_\Theta(\theta,g\cdot\theta').$$

Taking the infimum over $g\in G$ yields $d_{\mathcal{F}_0}(P_\theta^X,P_{\theta'}^X)\le L\,d_G(\theta,\theta')$. □

Remark G.19 (Upper moduli on the quotient). If $\bar d(\bar\theta,\bar\theta'):=d_G(\theta,\theta')$ denotes the induced pseudometric on Θ/G, then Lemma G.18 implies $\overline{\mathsf{D}}\le L\,\bar d$ on Θ/G.

Lower moduli: local injectivity and transverse identifiability.

Definition G.20 (Local lower modulus around an orbit)**.** Let D be any pseudometric on observed laws and define the induced pseudometric on parameters by

$$\mathsf{D}_\Theta(\theta,\theta'):=\mathsf{D}(P_\theta^X,P_{\theta'}^X).$$

Fix $\theta_0\in\Theta$ and an orbit pseudometric d_G on Θ (e.g. induced from a base pseudometric d_Θ). We say that D admits a *local lower modulus* $\underline{\omega}$ around $[(]\,\theta_0)$ if there exist $r>0$ and a nondecreasing function $\underline{\omega}:[0,r)\to[0,\infty)$ with $\underline{\omega}(0)=0$ such that

$$\underline{\omega}\big(d_G(\theta,\theta_0)\big)\ \le\ \mathsf{D}_\Theta(\theta,\theta_0)\qquad\text{whenever } d_G(\theta,\theta_0)<r. \tag{360}$$

If $\underline{\omega}$ is strictly increasing on $[0,r)$, define its right-continuous generalized inverse by

$$\underline{\omega}^{-1}(y)\ :=\ \inf\{u\in[0,r):\ \underline{\omega}(u)\ge y\}.$$

Remark G.21 (Interpretation). The inequality (360) states that motion *transverse to the orbit* (measured by d_G) forces a nontrivial movement of the observed law (measured by D). When $\underline{\omega}(u) \asymp u$ locally, this is a local bi-Lipschitz identifiability condition on the quotient.

Rate propagation through moduli.

Lemma G.22 (From D-estimation to an orbit rate via the inverse modulus)**.** *Assume the model is well-specified: $P^\star = P^X_{\theta_0}$ for some $\theta_0 \in \Theta$. Let $\widehat{\theta}_n$ be any (possibly randomized) estimator such that*

$$\mathsf{D}(P^X_{\widehat{\theta}_n}, P^X_{\theta_0}) \;=\; O_{\mathbb{P}}(r_n) \qquad \text{for some } r_n \downarrow 0.$$

If D admits a strictly increasing local lower modulus $\underline{\omega}$ around $[(]\,\theta_0)$ as in Definition G.20, then

$$d_G(\widehat{\theta}_n, \theta_0) \;=\; O_{\mathbb{P}}\big(\underline{\omega}^{-1}(r_n)\big).$$

Proof. Let r be the radius in Definition G.20. By (360), on the event $\{d_G(\widehat{\theta}_n, \theta_0) < r\}$ we have

$$\underline{\omega}\big(d_G(\widehat{\theta}_n, \theta_0)\big) \le \mathsf{D}(P^X_{\widehat{\theta}_n}, P^X_{\theta_0}).$$

Applying the (monotone) generalized inverse $\underline{\omega}^{-1}$ gives

$$d_G(\widehat{\theta}_n, \theta_0) \le \underline{\omega}^{-1}\big(\mathsf{D}(P^X_{\widehat{\theta}_n}, P^X_{\theta_0})\big) \qquad \text{on } \{d_G(\widehat{\theta}_n, \theta_0) < r\}.$$

Since $\mathsf{D}(P^X_{\widehat{\theta}_n}, P^X_{\theta_0}) = O_{\mathbb{P}}(r_n)$ and $r_n \downarrow 0$, the right-hand side converges to 0 in probability, hence $\Pr(d_G(\widehat{\theta}_n, \theta_0) < r) \to 1$; thus the inequality above holds with probability tending to one. Finally, monotonicity of $\underline{\omega}^{-1}$ and $O_{\mathbb{P}}(r_n)$ yield the claimed rate. □

Misspecification (distance to an orbit-invariant target set). Under misspecification, one often targets an orbit-invariant set $\bar{\Theta} \subseteq \Theta$ (e.g. $\arg\max \Phi$). Define the orbit-distance-to-set

$$d_G(\theta, \bar{\Theta}) := \inf_{\vartheta \in \bar{\Theta}} d_G(\theta, \vartheta),$$

and the approximation error (in D) to that set:

$$\varepsilon_\star(\bar{\Theta}) \;:=\; \inf_{\vartheta \in \bar{\Theta}} \mathsf{D}(P^\star, P^X_\vartheta).$$

(If $\bar{\Theta} = \Theta$ this reduces to the model-mismatch radius $\inf_{\vartheta \in \Theta} \mathsf{D}(P^\star, P^X_\vartheta)$.)

Lemma G.23 (Misspecified rate propagation to distance-to-set)**.** *Fix an orbit-invariant set $\bar{\Theta} \subseteq \Theta$ and an orbit pseudometric d_G. Assume there exist $r > 0$ and a strictly increasing $\underline{\omega} : [0, r) \to [0, \infty)$ such that for every $\vartheta \in \bar{\Theta}$,*

$$\underline{\omega}\big(d_G(\theta, \vartheta)\big) \;\le\; \mathsf{D}(P^X_\theta, P^X_\vartheta) \qquad \text{whenever } d_G(\theta, \vartheta) < r. \tag{361}$$

Let $\widehat{\theta}_n$ satisfy $\mathsf{D}(P^X_{\widehat{\theta}_n}, P^\star) = O_{\mathbb{P}}(r_n)$ with $r_n \downarrow 0$. Then

$$d_G(\widehat{\theta}_n, \bar{\Theta}) \;=\; O_{\mathbb{P}}\Big(\underline{\omega}^{-1}\big(r_n + \varepsilon_\star(\bar{\Theta})\big)\Big),$$

provided $r_n + \varepsilon_\star(\bar{\Theta}) < \underline{\omega}(r)$ eventually (so that the inverse is evaluated within its domain).

Proof. Fix $\delta > 0$ and choose $\vartheta_\delta \in \bar{\Theta}$ such that

$$\mathsf{D}(P^\star, P^X_{\vartheta_\delta}) \le \varepsilon_\star(\bar{\Theta}) + \delta.$$

By the triangle inequality,

$$\mathsf{D}(P^X_{\widehat{\theta}_n}, P^X_{\vartheta_\delta}) \le \mathsf{D}(P^X_{\widehat{\theta}_n}, P^\star) + \mathsf{D}(P^\star, P^X_{\vartheta_\delta}) \le \mathsf{D}(P^X_{\widehat{\theta}_n}, P^\star) + \varepsilon_\star(\bar{\Theta}) + \delta.$$

On the event $\{d_G(\widehat{\theta}_n, \vartheta_\delta) < r\}$, the uniform lower modulus (361) gives

$$\underline{\omega}\big(d_G(\widehat{\theta}_n, \vartheta_\delta)\big) \le \mathsf{D}(P^X_{\widehat{\theta}_n}, P^X_{\vartheta_\delta}).$$

Apply $\underline{\omega}^{-1}$ and use $d_G(\widehat{\theta}_n, \bar{\Theta}) \le d_G(\widehat{\theta}_n, \vartheta_\delta)$ to obtain

$$d_G(\widehat{\theta}_n, \bar{\Theta}) \le \underline{\omega}^{-1}\Big(\mathsf{D}(P^X_{\widehat{\theta}_n}, P^\star) + \varepsilon_\star(\bar{\Theta}) + \delta\Big) \qquad \text{on } \{d_G(\widehat{\theta}_n, \vartheta_\delta) < r\}.$$

Since $\mathsf{D}(P^X_{\widehat{\theta}_n}, P^\star) = O_{\mathbb{P}}(r_n)$, the right-hand side is $O_{\mathbb{P}}(\underline{\omega}^{-1}(r_n + \varepsilon_\star(\bar{\Theta}) + \delta))$. Letting $\delta \downarrow 0$ yields the claim. □

Remark G.24 (What must be checked in concrete EM problems). To use Lemmas G.22–G.23 effectively, one typically verifies: (i) an estimable control $\mathsf{D}(P^X_{\widehat{\theta}_n}, P^\star) = O_{\mathbb{P}}(r_n)$ (Appendix E), and (ii) a lower modulus on a slice (transverse identifiability), e.g. via a Jacobian nonsingularity condition for a finite feature map (Lemma G.7) or via curvature of a likelihood-type criterion restricted to a slice.

References

1. S. Balakrishnan, Martin J. Wainwright, and Bin Yu, *Statistical guarantees for the EM algorithm: From population to sample-based analysis*, Ann. Statist. **45** (2017), no. 1, 77–120.
2. Jérôme Bolte, Shoham Sabach, and Marc Teboulle, *Proximal alternating linearized minimization for nonconvex and nonsmooth problems*, Math. Programming **146** (2014), no. 1–2, 459–494.
3. Olivier Bousquet and André Elisseeff, *Stability and generalization*, J. Mach. Learn. Res. **2** (2002), 499–526.
4. A. P. Dempster, N. M. Laird, and D. B. Rubin, *Maximum likelihood from incomplete data via the EM algorithm*, J. Roy. Statist. Soc. Ser. B **39** (1977), no. 1, 1–38.
5. Raaz Dwivedi, Nhat Ho, Koulik Khamaru, Martin J. Wainwright, Michael I. Jordan, and Bin Yu, *Theoretical guarantees for EM under misspecified Gaussian mixture models*, Advances in Neural Information Processing Systems 31 (NeurIPS 2018), 2018.
6. ———, *Singularity, misspecification and the convergence rate of EM*, Ann. Statist. **48** (2020), no. 6, 3161–3182.
7. Koustav Mallik, *Tensor algebra toolkit for folded mixture models: Symmetry-aware moments, orbit-space estimation, and poly-LAN rates*, arXiv preprint, 2025.
8. A. Müller, *Integral probability metrics and their generating classes of functions*, Adv. in Appl. Probab. **29** (1997), 429–443.
9. Bharath K. Sriperumbudur, Kenji Fukumizu, Arthur Gretton, Bernhard Schölkopf, and Gert R. G. Lanckriet, *On the empirical estimation of integral probability metrics*, Electron. J. Statist. **6** (2012), 1550–1599.
10. Michel Talagrand, *The generic chaining: Upper and lower bounds of stochastic processes*, Springer, 2005.
11. Aad W. van der Vaart and Jon A. Wellner, *Weak convergence and empirical processes: With applications to statistics*, Springer, 1996.
12. C. F. J. Wu, *On the convergence properties of the EM algorithm*, Ann. Statist. **11** (1983), no. 1, 95–103.

M.Stat 2nd year student, Indian Staistical Institute,Kolkata
Email address: koustavmallik2020@gmail.com