When "Normalization Without Loss of Generality" Loses Generality[*]

Wayne Yuan Gao[†]

March 29, 2026

Abstract

Normalization is ubiquitous in economics, and a growing literature shows that "normalizations" can matter for interpretation, counterfactual analysis, misspecification, and inference. This paper provides a general framework for these issues, based on the formalized notion of *modeling equivalence* that partitions the space of unknowns into equivalence classes, and defines normalization as a WLOG selection of one representative from each class. A counterfactual parameter is normalization-free if and only if it is constant on equivalence classes; otherwise any point identification is created by the normalization rather than by the model. Applications to discrete choice, demand estimation, and network formation illustrate the insights made explicit through this criterion. We then study two further sources of fragility: an extension trilemma establishes that fidelity, invariance, and regularity cannot simultaneously hold at a boundary singularity, while a normalization can itself introduce a coordinate singularity that distorts the topological and metric structures of the parameter space, with consequences for estimation and inference.

Keywords: identification, normalization, without loss of generality, equivalence class, counterfactual, discrete choice

1 Introduction

Normalization is standard practice in econometrics and applied economics. In binary choice models, the error variance or the norm of a structural parameter is set to unity. In demand estimation, the idiosyncratic taste shock follows a Type-I extreme value distribution, with the latent utility of the outside option often set to zero. In network formation models, parameter norms or quantiles of the unobserved shock are fixed. In entry games, the variance of the private shock is set to one. In each case, certain unknowns, typically the location and scale of an unobservable, are not separately identified from the structural parameters, and a normalization pins them down.

Normalization is not merely a notational convenience. A substantial literature now shows, in different ways, that restrictions described as normalizations can affect interpretation, counterfactuals, misspecification, and the geometry of inference. Point identification results are still routinely stated "under normalization," and the implications of that qualifier for counterfactual analysis are often left implicit. This creates two recurring problems for applied work. First, the distinction between "identified" and "identified under normalization" is easily blurred: once a normalization has been imposed, the resulting parameter values are treated as uniquely identified structural objects, when in fact they are only one labeling of an equivalence class. Second, counterfactual parameters such as welfare changes or percentage welfare changes may or may not be invariant to the normalization, and when they are not, any apparent point identification is normalization-relative rather than model-driven. The difficulty is not confined to one class of models, but it is usually diagnosed case by case.

Percentage changes can be particularly dangerous, because when the normalization involves a location shift, the denominator of the percentage is itself normalization-dependent. A simple daily-life analogy makes this concrete: consider a warming of 10°C in physical temperature from 1°C to 11°C across two days, and the same physical temperature change simultaneously corresponds to a 1,000% increase in Celsius, a 53% increase in Fahrenheit, but only a 3.6% increase in Kelvin. The percentage is an artifact of the zero point, not of the phenomenon, and a report of "1,000% increase in temperature" without proper disclosure of the underlying unit (and the affine transformation that defines the unit) can be misleading. Just as Celsius, Fahrenheit, and Kelvin are related by positive affine transformations, the indeterminacy in discrete response models is often generated (and then normalized away) by transformations of

the same form. Whenever a structural model identifies utility or surplus only up to a location-scale transformation, raw percentage welfare changes inherit the arbitrariness of the normalization's zero point.

A striking recent illustration of a related issue appears in Chen and Roth (2024), who show that for any "log-like" transformation such as $\log(1 + Y)$ or $\mathrm{arcsinh}(Y)$, the ATE is *arbitrarily* sensitive to the units of Y whenever the treatment moves individuals onto or off the boundary $Y = 0$. In the language of the present paper, the percentage effect has a singularity at $Y = 0$, and any completion that "patches" this singularity imports an auxiliary convention (units) rather than a well-defined property of the model. Section 4.1 formalizes this insight as a general extension trilemma.

Against this backdrop, the paper develops a general framework built on three interlocking concepts. *Modeling equivalence* is a bijective transformation of the unknowns that preserves the structural equation *and* all maintained assumptions. The notion of modeling equivalence is strictly stronger than classical observational equivalence, which allows arbitrary changes in the model so long as the observable data distribution is preserved. We then use modeling equivalence to define a partition of the space of unknowns in a model into equivalence classes. Proper *normalization* selects one representative from each class: it is a labeling device *without loss of generality*, not a source of identifying information.

We then articulate how normalization, even when imposed without loss of generality, affects the identifiability of a *counterfactual parameter*, defined as any known function of the unknowns. Our main results are a lemma on the *limit of identifiability*, showing that any counterfactual that varies within some equivalence class cannot be point-identified for all true parameter values, and a theorem on *WLOG normalization*, establishing that both the identifiability and the identified value of any counterfactual are invariant to the choice of normalization. Together, these provide a clean necessary condition: for a counterfactual to be point-identified, it must be constant on every modeling equivalence class. In other words, it must factor through the quotient space induced by the model's indeterminacy. Crucially, verifying this congruence condition requires an explicit statement of the modeling equivalence relation, which is encoded in the original unnormalized model but cannot be recovered from the normalized model alone.

We then turn to the interaction between normalization and singularity. An *extension trilemma* result shows that when a counterfactual is normalization-free on

a regular subdomain but the induced quotient functional has a singularity at the boundary, no completion can simultaneously preserve fidelity, invariance, and regularity. This provides a complementary perspective to the impossibility result in Chen and Roth (2024), framed in terms of arbitrary equivalence structures and quotient topologies: extensions that "patch" boundary singularities are not innocuous normalizations; they necessarily change the object of interest. The singularity story has a flip side: whereas domain augmentation attempts to repair a pre-existing singularity, a normalization can *create* one. Dividing by a distinguished coefficient maps the excluded hyperplane to infinity, introducing a coordinate singularity that distorts the topology and metric of the parameter space. Using binary choice scale normalization as a case study, we show that the special-coordinate normalization destroys compactness and connectedness and can reverse convergence verdicts, while the sphere normalization preserves these properties by construction. More broadly, whether a normalization is WLOG cannot be assessed in isolation, but must be assessed end-to-end, jointly with the maintained model, the target counterfactual, and any auxiliary structure used downstream.

We illustrate the framework through three applied settings: binary choice, multinomial choice (BLP-type demand models), and dyadic network formation, showing that marginal effects, market shares, and linking probabilities are normalization-free, while welfare levels and percentage welfare changes are generically normalization-dependent. Money-metric welfare *changes* occupy a middle ground: in the standard quasilinear logit model, the location shift cancels in the difference and the scale cancels against the price coefficient, making ΔCS normalization-free. That said, this cancellation is model-specific and does not extend to welfare levels or percentages. When normalizations or other conventional assumptions are adopted, we argue that the normalization choice, the underlying equivalence class, and the dependence of each target counterfactual on that choice should all be made explicit, especially to applied researchers and policy makers who rely on the reported quantities.

Literature Review

The literature most closely related to ours comes from three partly separate strands. An important starting point is Lewbel (2019), especially Section 6.3, which highlights that whether a restriction is genuinely without loss depends on the object of interest, and that normalization choice can matter for interpretation and precision even

when it is innocuous for point identification. Our paper formalizes and extends these insights by introducing modeling equivalence as a refinement of observational equivalence, defining the quotient space of equivalence classes, proving WLOG as a theorem rather than asserting it informally, and giving a single general invariance criterion for normalization-free counterfactuals.

The second strand studies normalization as a problem of parameter-space geometry and inference. Hamilton, Waggoner, and Zha (2007) show that poor normalizations can distort the geometry of the parameter space, producing multimodal posteriors, disjoint confidence regions, and misleading uncertainty statements in likelihood-based models. Their "identification principle" says that the boundary of a normalized parameter space should line up with loci of local nonidentification. Our analysis in Section 4.2 is closely related in spirit, but differs in scope and object: rather than start from a finite-dimensional likelihood, we work with modeling-equivalence classes in the space of unknowns and ask which counterfactuals, regularity conditions, and inferential objects descend to the quotient.

The third strand shows in concrete structural models that some restrictions described as normalizations are in fact substantive assumptions. Agostinelli and Wiswall (2025) show that period-by-period re-normalization in dynamic latent-factor models can be over-identifying and restrictive when the production technology already has known location and scale, and can bias estimation; their earlier working paper (Agostinelli and Wiswall, 2016) develops the normalization analysis in detail. Freyberger (2025) likewise shows in skill-formation models that seemingly innocuous scale and location restrictions may constrain identified parameters and alter counterfactuals, depending on the production technology, measurement system, and estimation method. Freyberger also provides a formal definition of normalization: a restriction $\tau \in \mathcal{T}_N$ is a normalization with respect to $g(\tau_0)$ if $\{g(\tau) : \tau \in \mathcal{T}_0 \cap \mathcal{T}_N\} = \{g(\tau) : \tau \in \mathcal{T}_0\}$, which is restriction-based and function-specific. Our contribution is to provide the general language that makes these findings portable across models: a restriction is a true normalization only if it amounts to choosing a representative within a modeling-equivalence class, while failures arise either because the target counterfactual is not invariant on the class or because the downstream object does not extend coherently to the quotient. In addition, our notion of modeling equivalence is strictly stronger than the observational equivalence used in all three strands, because it requires preservation of all maintained assumptions, not just the data distribution (see

Section 2.3).

Normalization has also been discussed in many specific settings: semiparametric binary response (Manski, 1988), panel binary choice (Chamberlain, 2010), demand estimation (Berry et al., 1995; Berry and Haile, 2014), entry games (Tamer, 2003; Ciliberto and Tamer, 2009), network formation (Graham, 2017; Gao, 2020), and dynamic discrete choice (Rust, 1994; Magnac and Thesmar, 2002), where the outside-option normalization is famously *not* free.

A related and growing literature studies the identification of counterfactual parameters in structural models. Kalouptsidi, Scott, and Souza-Rodrigues (2021) characterize which counterfactuals are identified in dynamic discrete choice models despite non-identification of the full model, and note that scale normalizations cease to be innocuous once cardinal objects such as value functions enter the counterfactual. Bhattacharya, Dupas, and Kanaya (2024) show that in standard discrete choice the outside-option normalization is innocuous for welfare analysis, whereas with social interactions it is not, because the outside-option utility itself becomes policy-dependent. Gu, Russell, and Stringham (2025) develop a general framework for counterfactual identification in discrete outcome models. These papers study what can be identified *after* a normalization has been imposed, and some point out in model-specific terms that a purported normalization is or is not innocuous. Our contribution is different: we isolate normalization itself as a general, cross-model source of counterfactual fragility, formalized through modeling-equivalence classes, quotient factorization, interior non-invariance, and boundary or chart singularities.

The rest of the paper is organized as follows. Section 2 presents the general framework. Section 3 develops three applied examples. Section 4 examines how normalization interacts with singularities, both in the target functional and in the normalized parameter space. Section 5 introduces end-to-end WLOG normalization and offers practical recommendations. Section 6 concludes.

2 General Framework

2.1 An Econometric Model

We write a generic econometric model, i.e., a collection of structural equations together with all maintained assumptions, abstractly as

$$f(\mathbf{X}, \mathbf{U}, \gamma_0) = 0, \tag{1}$$

where $\mathbf{X}$ collects all observable random elements (data), $\mathbf{U} \in \mathcal{U}$ collects all unobservable random elements (errors and latent variables), $\gamma_0 \in \Gamma$ is the true value of all unknown fixed parameters (possibly infinite-dimensional), and f is a known functional encoding the structural equation *and* all maintained assumptions, such as distributional assumptions, independence restrictions, support conditions, and any other maintained hypotheses. We emphasize that f incorporates every aspect of the model: there is no fundamental distinction between "model specification" and "model assumption."

The tuple $(\mathbf{X}, \mathbf{U}, \gamma_0, f)$ provides a completely general decomposition of an econometric model into its observable, unobservable, unknown, and known components. We call $\mathcal{U} \times \Gamma$ the *space of unknowns*.[1]

2.2 Modeling Equivalence

Definition 1 (Modeling Equivalent Transformation). A bijection $\psi : \mathcal{U} \times \Gamma \to \mathcal{U} \times \Gamma$ is a *modeling equivalent transformation* if

$$f\big(\mathbf{X}, \psi(\mathbf{U}, \gamma)\big) = f(\mathbf{X}, \mathbf{U}, \gamma), \quad \forall\, (\mathbf{U}, \gamma) \in \mathcal{U} \times \Gamma. \tag{2}$$

In other words, ψ preserves the entire model f, not merely the distribution of the data, but all maintained assumptions as well.

Definition 2 (Modeling Equivalence). Two points $(\mathbf{U}, \gamma)$ and $(\mathbf{U}', \gamma')$ in $\mathcal{U} \times \Gamma$ are *modeling equivalent*, written $(\mathbf{U}, \gamma) \sim (\mathbf{U}', \gamma')$, if there exists a modeling equivalent

---

[1]More abstractly, let Θ denote the space of unknowns and let $\mathcal{L} : \Theta \rightrightarrows \mathscr{P}(\mathcal{X})$ be the observable implication map, where $\mathcal{L}(\theta)$ collects all observable laws consistent with the model when the state equals θ. The structural representation $f(\mathbf{X}, \mathbf{U}, \gamma_0) = 0$ is one way to generate $\mathcal{L}$; the definitions below apply to any $\mathcal{L}$, and in particular accommodate incomplete models and models with nuisance variation.

transformation ψ such that $(\mathbf{U}', \gamma') = \psi(\mathbf{U}, \gamma)$.

Note that modeling equivalence is strictly stronger than the classical notion of *observational equivalence*, which requires only that two configurations of the unknowns generate the same data distribution. Observational equivalence allows a simultaneous change in the maintained assumptions (e.g., switching distributional families), while modeling equivalence does not. See Section 2.3 for more detail about this distinction.

Modeling equivalence is a well-defined equivalence relation on $\mathcal{U} \times \Gamma$: the identity map establishes reflexivity, the inverse of a bijection satisfying (2) again satisfies (2) (symmetry), and the composition of two such bijections satisfies (2) (transitivity).

The resulting quotient space is thus a partition, which we call the *modeling equivalence partition*, defined by

$$Q_\sim := (\mathcal{U} \times \Gamma)/\sim. \tag{3}$$

We write $[\mathbf{U}, \gamma]$ for the equivalence class containing $(\mathbf{U}, \gamma)$. Throughout, $Q_\sim$ carries the quotient topology induced by the canonical projection $\pi : \mathcal{U} \times \Gamma \to Q_\sim$, so that $V \subseteq Q_\sim$ is open if and only if $\pi^{-1}(V)$ is open in $\mathcal{U} \times \Gamma$.

2.3 Modeling Equivalence versus Observational Equivalence

The standard notion of observational equivalence requires only that two configurations generate the same data distribution: $(\mathbf{U}, \gamma) \sim_{OE} (\mathbf{U}', \gamma')$ if the observable distributions of data they induce are identical. Clearly, modeling equivalence implies observational equivalence, since (2) ensures the data distribution is preserved, but the converse is in general false, because observationally equivalent transformations may violate maintained assumptions. This distinction matters because normalization constrains the unknowns *within* the model's maintained assumptions, but it should not change the model itself. The relevant indeterminacy for normalization analysis is therefore characterized by modeling equivalence, not by observational equivalence.

We illustrate this difference in two stylized examples.

Example 1 (Binary choice: probit and logit)**.** In the binary choice model $Y_i = \mathbb{1}\{X_i'\boldsymbol{\beta}_0 \geq \varepsilon_i\}$, compare two distinct maintained specifications: one with $\varepsilon_i \sim \mathcal{N}(0, 1)$ (probit) and another with $\tilde{\varepsilon}_i \sim \text{Logistic}(0, 1)$ (logit), with suitably adjusted $\tilde{\boldsymbol{\beta}}_0$. Suppose they happen to generate the same choice probabilities on the support of X_i, say, taken to be a finite set on which one can solve the nonlinear equations that

assert the equality of choice probabilities. Then they are observationally equivalent on that support. However, they are *not* modeling equivalent: replacing a normal error distribution with a logistic one changes a distributional assumption of the stated model rather than relabeling the same underlying unknowns. Changing the error family, even when the observable distribution is unchanged, is therefore not a modeling equivalent transformation.[2]

Example 2 (Fixed Effects in Dyadic Network formation)**.** Consider the dyadic network formation model $D_{ij} = \mathbb{1}\{w(X_i, X_j) + A_i + A_j \geq U_{ij}\}$ under cross-sectional random sampling and i.i.d. pairwise shocks U_{ij} à la Graham (2017) and Gao (2020), which is analyzed in more detail in Section 3.3. One may define a transformation of the unknown error and fixed effects by $\hat{A}_i := A_i - A_1$ and $\hat{U}_{ij} := U_{ij} - 2A_1$, i.e., "normalize" with respect to individual 1's fixed effect A_1. Since $w(X_i, X_j) + \hat{A}_i + \hat{A}_j - \hat{U}_{ij} = w(X_i, X_j) + A_i + A_j - U_{ij}$, the transformed specification produces exactly the same data distribution; hence, the transformation preserves observational equivalence. However, while A_i is iid across i under the assumption of cross-sectional random sampling, $\hat{A}_i = A_i - A_1$ is clearly *not* iid (all $\hat{A}_i$ are correlated through A_1, and furthermore $\hat{A}_i = 0$ is degenerate), violating the assumption of cross-sectional random sampling. Hence this transformation does *not* preserve modeling equivalence.

These examples show that the set of modeling equivalent transformations is a proper subset of the set of observational-equivalence-preserving transformations. Using observational equivalence for normalization analysis would overstate the indeterminacy by including transformations that are inadmissible given the model's maintained assumptions.

2.4   Normalization

Definition 3 (Normalization)**.** A mapping $\psi_N : \mathcal{U} \times \Gamma \to \mathcal{U} \times \Gamma$ is a *normalization* if both the following hold:

   (i) Within-class collapse: $(\mathbf{U}, \gamma) \sim (\mathbf{U}', \gamma') \implies \psi_N(\mathbf{U}, \gamma) = \psi_N(\mathbf{U}', \gamma') \sim (\mathbf{U}, \gamma)$.
   (ii) Across-class separation: $(\mathbf{U}, \gamma) \not\sim (\mathbf{U}', \gamma') \implies \psi_N(\mathbf{U}, \gamma) \neq \psi_N(\mathbf{U}', \gamma')$.

---

[2]Note that one can drop the distributional assumption, consider a semiparametric binary choice, and reasonably ask the question of whether the whole CDF can be "normalized" once other modeling assumptions are laid out. However, that would be a different model from a "probit" or a "logit" model, and the corresponding definition of modeling equivalence would change accordingly.

In other words, a normalization selects exactly one representative from each equivalence class. It is not bijective by design: all elements within a given class are mapped to a single element. Different global normalizations used for binary choice models, for example, $\varepsilon_i \sim \mathcal{N}(0,1)$ versus $\|\boldsymbol{\beta}_0\| = 1$, are simply different labelings of the same equivalence classes. By contrast, a coordinate restriction such as $|\beta_{0,1}| = 1$ is best viewed as a *local chart* on the two open hemispheres where $\beta_{0,1} \neq 0$, not as a globally available normalization; it breaks down as $\beta_{0,1} \to 0$ (see Section 4.2).

Proposition 1 (Normalization Induces a Bijection). *Every normalization ψ_N induces a bijection between $Q_\sim$ and the normalized parameter space $\psi_N(\mathcal{U} \times \Gamma)$.*

Proof. Define $s_N : Q_\sim \to \psi_N(\mathcal{U} \times \Gamma)$ by $s_N([\mathbf{U}, \gamma]) = \psi_N(\mathbf{U}, \gamma)$. This is well-defined by condition (i) of Definition 3. Surjectivity is by construction; injectivity follows from condition (ii). $\qquad\square$

Proposition 1 shows that the normalized parameter space is simply a relabeling of the quotient: every normalization picks a set of representatives that is in one-to-one correspondence with the equivalence classes. The choice among normalizations is therefore a choice among coordinate systems for the same underlying object.[3]

2.5   Identifiability of Counterfactual Parameters

Definition 4 (Counterfactual Parameter). A *counterfactual* is a known function $q : \mathcal{U} \times \Gamma \to \mathbb{R}$. The *counterfactual parameter* is the value $q(\mathbf{U}_0, \gamma_0)$ evaluated at the true unknowns $(\mathbf{U}_0, \gamma_0)$.

Counterfactual parameters are the ultimate objects of economic interest. They encode the answers to questions such as "what is the marginal effect of x_j?" or "what is the welfare change from a policy intervention?"

Technically, a counterfactual is a known *functional* defined on the space of the unknowns. We use the name "counterfactual" to make explicit its standard use in applied economics.

---

[3]Freyberger (2025) provides a complementary definition: a restriction $\tau \in \mathcal{T}_N$ is a normalization with respect to $g(\tau_0)$ if $\{g(\tau) : \tau \in \mathcal{T}_0 \cap \mathcal{T}_N\} = \{g(\tau) : \tau \in \mathcal{T}_0\}$. This is restriction-based and function-specific. Our definition is structure-based: the equivalence class partition governs all functions simultaneously, and a function is normalization-free if and only if it factors through the quotient $Q_\sim$.

Lemma 1 (Limit of Identifiability of A Counterfactual). *If a counterfactual $q : \mathcal{U} \times \Gamma \to \mathbb{R}$ is not constant on some modeling equivalence class, then there exist true values $(\mathbf{U}_0, \gamma_0)$ such that $q(\mathbf{U}_0, \gamma_0)$ cannot be point-identified.*

Proof. Since q is not $Q_\sim$-measurable, there exists a modeling equivalence class $p \in Q_\sim$ and two elements $(\mathbf{U}, \gamma), (\mathbf{U}', \gamma') \in p$ such that $q(\mathbf{U}, \gamma) \neq q(\mathbf{U}', \gamma')$.

Suppose the true values satisfy $(\mathbf{U}_0, \gamma_0) \in p$. Since $(\mathbf{U}_0, \gamma_0) \sim (\mathbf{U}, \gamma) \sim (\mathbf{U}', \gamma')$, there exist modeling equivalent transformations ψ and ψ' with $(\mathbf{U}_0, \gamma_0) = \psi(\mathbf{U}, \gamma) = \psi'(\mathbf{U}', \gamma')$. Define $\bar{f}(\mathbf{X}, \cdot) := f(\mathbf{X}, \psi(\cdot))$ and $\underline{f}(\mathbf{X}, \cdot) := f(\mathbf{X}, \psi'(\cdot))$. By (2), $\bar{f} \equiv f \equiv \underline{f}$. Under $\bar{f}$ the "true value" is $\psi^{-1}(\mathbf{U}_0, \gamma_0) = (\mathbf{U}, \gamma)$; under $\underline{f}$ it is $(\psi')^{-1}(\mathbf{U}_0, \gamma_0) = (\mathbf{U}', \gamma')$. Since $\bar{f} \equiv \underline{f} \equiv f$ and the observable data are the same, any valid identified set must contain both $q(\mathbf{U}, \gamma)$ and $q(\mathbf{U}', \gamma')$. Since these differ, $q(\mathbf{U}_0, \gamma_0)$ is not point-identified. $\square$

The intuition is straightforward: if two modeling-equivalent configurations $(\mathbf{U}, \gamma)$ and $(\mathbf{U}', \gamma')$ have $q(\mathbf{U}, \gamma) \neq q(\mathbf{U}', \gamma')$, then the model, which by construction cannot distinguish these two configurations, must include both values in any identified set.

This suggests a useful classification of counterfactual statements. We say that q is *normalization-free* if it is constant on every equivalence class, and *normalization-dependent* otherwise. Lemma 1 says that being normalization-free is *necessary* for point identification: a normalization-dependent counterfactual is provably non-point-identified for some true parameter values. Being normalization-free is not by itself *sufficient* for point identification, since there may be other, non-normalization sources of non-identification (which is model specific and not the focus of this paper). The paper's message is not about what *can* be identified after normalization, but about what *cannot* be identified regardless of normalization.

Corollary 1 (Necessary Condition for Point Identification). *If a counterfactual q is point-identified for all true values $(\mathbf{U}_0, \gamma_0)$, then q must be constant on every modeling equivalence class.*

The corollary is simply the contrapositive of Lemma 1. Its content can be restated as a factorization condition: q is normalization-free if and only if there exists a map $\bar{q} : Q_\sim \to \mathbb{R}$ such that

$$q = \bar{q} \circ \pi, \quad \text{for } \pi : \mathcal{U} \times \Gamma \to Q_\sim, \tag{4}$$

where π is the *canonical projection* that sends each $(\mathbf{U}, \gamma)$ to its equivalence class $[\mathbf{U}, \gamma]$. Informally, truly identified counterfactuals are functionals of equivalence classes rather than of arbitrary representatives.

The lemma and corollary above allow us to separate several cases regarding the identification of a counterfactual parameter:

1. *Normalization-free point identification*: the counterfactual is constant on every equivalence class and is point-identified (the ideal case).
2. *Normalization-dependent "point identification"*: the counterfactual varies within some equivalence class, yet is reported as point-identified after normalization.
3. *Point identification up to normalization* (or *normalizable set identification*): the identified set, prior to normalization, is a union of equivalence classes that collapses to a point once a normalization is imposed.[4]
4. *Unnormalizable partial identification*: the identified set does not collapse to a point even after a proper normalization.[5]

Generically, we view case (2) as an ill-posed notion of identification: a normalization-dependent counterfactual that is "point-identified" only under a specific and fundamentally arbitrary normalization. Such results should be interpreted as negative findings, or at the very least with great caution, since the reported value depends on a labeling convention rather than on the phenomenon itself.

Once the subtlety above is clarified, a researcher can then make a justified WLOG claim about normalization as formalized in the following theorem.

Theorem 1 (Normalization WLOG)**.** *Let q be any counterfactual and ψ_N any normalization in the sense of Definition 3. Then the identified set of q under the normalized model is identical to its identified set under the original model.*

Proof. Let ψ_N be any normalization in the sense of Definition 3. After imposing ψ_N, the model becomes $f_N(\mathbf{X}, \tilde{\mathbf{U}}, \tilde{\gamma}) := f(\mathbf{X}, \psi_N(\tilde{\mathbf{U}}, \tilde{\gamma}))$. By Definition 3(i), $\psi_N(\tilde{\mathbf{U}}, \tilde{\gamma}) \sim (\tilde{\mathbf{U}}, \tilde{\gamma})$ for all $(\tilde{\mathbf{U}}, \tilde{\gamma})$. By the definition of modeling equivalence (2):

$$f_N(\mathbf{X}, \tilde{\mathbf{U}}, \tilde{\gamma}) = f\big(\mathbf{X}, \psi_N(\tilde{\mathbf{U}}, \tilde{\gamma})\big) = f(\mathbf{X}, \tilde{\mathbf{U}}, \tilde{\gamma}), \quad \forall (\tilde{\mathbf{U}}, \tilde{\gamma}). \tag{5}$$

---

[4]Normalizable set identification refers to a set identification result that admits a normalization with respect to the relevant modeling equivalence class, after which it becomes a point identification result up to the normalization.

[5]Unnormalizable partial identification refers to a "standard" partial identification result that does not pin down a point even after a proper normalization.

Hence $f_N \equiv f$. The identified set of any counterfactual q depends only on the observable $\mathbf{X}$ and the known model f; since $f_N = f$, the identified set is unchanged. $\square$

In other words, normalization WLOG cannot create, destroy, enlarge, or shrink the identified set of any counterfactual q defined on the original space of unknowns.

We should make an important clarification: the theorem concerns a fixed counterfactual q, the same function evaluated before and after normalization. Normalization does make a normalization-*relative* object single-valued, but that is a different quantity, namely q composed with the normalization map ψ_N, not q itself. The theorem says that for the original q, imposing normalization changes nothing.

It should also be emphasized that "without loss of generality" does not mean "powerful." It means normalization causes *no change* in the model's ability to inform about any counterfactual. Normalization is purely a labeling convenience.

Remark 1 (Auxiliary restrictions and end-to-end WLOG). Theorem 1 establishes that normalization is WLOG for identification, but this theorem is relative to a given model with all its specifications and assumptions held fixed. Any object introduced *after* normalization (a different counterfactual, a regularity condition, a metric, a topology, a smoothness requirement, or a prior, imposed beyond the original model) must itself be well-defined on the quotient space for the WLOG interpretation to extend to that object. Formally, let $R \subseteq \mathcal{U} \times \Gamma$ denote an auxiliary restriction imposed after normalization. If R is a union of modeling equivalence classes, then imposing R preserves the WLOG status. If R is not a union of equivalence classes, then R is a substantive restriction: it refines the model, alters the quotient structure, and may change which counterfactuals are admissible or identified. A notable applied example is the "re-normalization" studied by Agostinelli and Wiswall (2025) and Freyberger (2025) in dynamic latent factor models: fixing the scale of a latent skill variable each period is WLOG with fully flexible production technologies, but becomes over-identifying and restrictive when combined with CES production functions, where the relative scale across periods is already pinned down by functional form. The discussion in Section 5 develops this point through concrete examples.

3 Examples

Normalization issues arise across a wide range of structural models. The papers discussed in the literature review provide detailed analyses in several specific settings: Hamilton, Waggoner, and Zha (2007) study mixture models, structural VARs, and cointegration; Agostinelli and Wiswall (2025) and Freyberger (2025) study dynamic latent-factor models of skill formation. Rather than revisit those settings, we illustrate the general framework through three different econometric models that highlight the larger-picture issues: how modeling equivalence classes are determined, how normalization-free and normalization-dependent counterfactuals are distinguished, and how percentage welfare changes can be particularly misleading. For each, we identify the modeling equivalent transformations, characterize the equivalence classes, and classify key counterfactuals.

3.1 Binary Choice Models

Consider the standard binary response model $Y_i = \mathbb{1}\{X_i'\boldsymbol{\beta}_0 \geq \varepsilon_i\}$, where $X_i = (1, X_{i,-1}')' \in \mathbb{R}^d$ includes a constant, $\boldsymbol{\beta}_0 \in \mathbb{R}^d$, and ε_i has CDF F_ε that is continuous, strictly increasing, and differentiable with density f_ε. For any $a \in \mathbb{R}$ and $b > 0$, define the positive affine transformation

$$\tilde{\varepsilon}_i = a + b\varepsilon_i, \qquad \tilde{\beta}_{0,1} = a + b\beta_{0,1}, \qquad \tilde{\boldsymbol{\beta}}_{0,-1} = b\boldsymbol{\beta}_{0,-1}. \tag{6}$$

By construction, $X_i'\tilde{\boldsymbol{\beta}}_0 \geq \tilde{\varepsilon}_i$ if and only if $X_i'\boldsymbol{\beta}_0 \geq \varepsilon_i$, and the transformed error $\tilde{\varepsilon}_i$ inherits continuity and strict monotonicity from ε_i. Hence (6) is a modeling equivalent transformation with two free parameters (a, b), reflecting one location and one scale degree of freedom.

It is important to distinguish the researcher's choice of distributional family, such as probit and logit, from the normalization that operates *within* the chosen family. The choice of family is a substantive modeling decision that determines the shape of F_ε, and this choice is not a normalization in our sense. Once the family is chosen, however, the location and scale of F_ε are not separately identified from $\boldsymbol{\beta}_0$, and it is here that normalization enters: setting $\varepsilon_i \sim \mathcal{N}(0, 1)$ or $Logistic(0, 1)$. Each such restriction selects one representative from the positive affine equivalence class. In the semiparametric case where F_ε is left unspecified (assumed only continuous and

strictly increasing), only location-scale restrictions, such as median-zero plus unit interquartile range, are genuine normalizations.

Several standard quantities of interest are normalization-free. Coefficient ratios $\beta_{0,j}/\beta_{0,k}$ for $j, k \geq 2$ are invariant because $\tilde{\beta}_{0,j}/\tilde{\beta}_{0,k} = (b\beta_{0,j})/(b\beta_{0,k}) = \beta_{0,j}/\beta_{0,k}$.[6] Marginal effects are invariant because the density rescaling exactly offsets the coefficient rescaling:

$$f_{\tilde{\varepsilon}}(x'\tilde{\boldsymbol{\beta}}_0)\,\tilde{\beta}_{0,j} = \tfrac{1}{b}\,f_{\varepsilon}\big(\tfrac{x'\tilde{\boldsymbol{\beta}}_0 - a}{b}\big)\cdot b\beta_{0,j} = f_{\varepsilon}(x'\boldsymbol{\beta}_0)\,\beta_{0,j}. \tag{7}$$

Choice probabilities $\mathbb{P}(Y_i = 1 \mid X_i = x) = F_{\varepsilon}(x'\boldsymbol{\beta}_0)$ are invariant because $F_{\tilde{\varepsilon}}(x'\tilde{\boldsymbol{\beta}}_0) = F_{\varepsilon}\big(\tfrac{x'\tilde{\boldsymbol{\beta}}_0 - a}{b}\big) = F_{\varepsilon}(x'\boldsymbol{\beta}_0)$. Note that choice probabilities and marginal effects are also directly identified from data, so it is not surprising at all that they are invariant under normalization.

By contrast, the level of latent utility $x'\boldsymbol{\beta}_0$ is normalization-dependent: under the transformation, $x'\tilde{\boldsymbol{\beta}}_0 = bx'\boldsymbol{\beta}_0 + a$, which depends on both the scale b and the location a. Average utility changes $\mathbb{E}[\Delta x'\boldsymbol{\beta}_0]$ scale with b (and shift with a when Δx includes a change in the intercept). Percentage welfare changes are even more problematic. Writing $W(x) := x'\boldsymbol{\beta}_0$ for latent welfare under the original parameterization, the percentage change from x to x' is

$$\frac{W(x') - W(x)}{W(x)} = \frac{\Delta x'\boldsymbol{\beta}_0}{x'\boldsymbol{\beta}_0}. \tag{8}$$

Under the positive affine transformation (6), this becomes

$$\frac{\widetilde{W}(x') - \widetilde{W}(x)}{\widetilde{W}(x)} = \frac{b\,\Delta x'\boldsymbol{\beta}_0}{b\,x'\boldsymbol{\beta}_0 + a} \neq \frac{\Delta x'\boldsymbol{\beta}_0}{x'\boldsymbol{\beta}_0}$$

whenever $a \neq 0$: the location shift contaminates the denominator, making the percentage depend on the normalization's zero point, which corresponds exactly to the temperature-change example in the introduction.

In applied binary choice, marginal effects and coefficient ratios are the standard quantities of interest for a good reason: they are the normalization-free counterfactuals. Claims about welfare levels or willingness-to-pay in dollar terms (unless expressed

---

[6]The intercept $j = 1$ is excluded because the location shift prevents cancellation: $\tilde{\beta}_{0,1}/\tilde{\beta}_{0,k} = (a + b\beta_{0,1})/(b\beta_{0,k})$, which depends on a.

as a ratio) are normalization-dependent and should be interpreted with care.

3.2 Multinomial Choice and Demand Models

Multinomial discrete choice models, originating with McFadden (1974) and extended to differentiated products markets by Berry, Levinsohn, and Pakes (1995), are among the most widely used tools in empirical industrial organization, marketing, transportation, and public policy. They underpin merger simulations, antitrust analysis, and welfare evaluations of product introductions and regulatory interventions, settings where the distinction between normalization-free and normalization-dependent counterfactuals has direct policy consequences.

To illustrate, consider a random-coefficient discrete choice model of demand, where consumer i's utility from product j in market t is $u_{ijt} = \delta_{jt} + \nu_{ijt} + \varepsilon_{ijt}$, with mean utility $\delta_{jt} = x'_{jt}\boldsymbol{\beta}_0 - \alpha_0 p_{jt} + \xi_{jt}$, preference heterogeneity ν_{ijt}, and idiosyncratic shock ε_{ijt}. Consumer i chooses $\arg\max_j u_{ijt}$, where $\varepsilon_{ijt} \sim$ Type-I extreme value distribution. For any $a \in \mathbb{R}$ and $b > 0$, the positive affine transformation $\tilde{u}_{ijt} = a + bu_{ijt}$ applied to *all* alternatives (including the outside option) preserves the argmax. In the model's decomposition this corresponds to $\tilde{\delta}_{jt} = a + b\delta_{jt}$ (for all j, so the outside option's mean utility becomes $\tilde{\delta}_{0t} = a$ rather than 0), $\tilde{\nu}_{ijt} = b\nu_{ijt}$, and $\tilde{\varepsilon}_{ijt} = b\varepsilon_{ijt}$. The standard normalization $\delta_{0t} = 0$ absorbs the location degree of freedom a, and $\varepsilon_{ijt} \sim$ TIEV absorbs the scale b.

As expected, market shares s_{jt}, own-price elasticities, and diversion ratios are all functions of choice probabilities and hence normalization-free.

The consumer surplus case is more nuanced and illustrates the framework's value. To obtain a closed form, consider the logit special case (no random coefficients: $\nu_{ijt} = 0$). Under logit with scale parameter μ (the standard normalization sets $\mu = 1$), the expected consumer surplus is $CS = (\mu/\alpha_0)\log\left(\sum_{j=0}^{J}\exp(\delta_{jt}/\mu)\right) + C$, where the sum runs over all alternatives including the outside option ($\delta_{0t} = 0$), α_0 is the marginal utility of income, and C is an unknown constant. The CS *level* is normalization-dependent: the constant C absorbs the location shift a, and the scale μ enters the formula. However, the money-metric CS *change* from a policy that moves mean utilities from δ to δ' is

$$\Delta CS = \frac{\mu}{\alpha_0}\left[\log\sum_{j=0}^{J}\exp(\delta'_{jt}/\mu) - \log\sum_{j=0}^{J}\exp(\delta_{jt}/\mu)\right], \tag{9}$$

in which the unknown constant C cancels. Under the transformation $\tilde{\mu} = b\mu$, $\tilde{\alpha}_0 = b\alpha_0$, and $\tilde{\delta}_{jt} = a + b\delta_{jt}$, one can verify that $\widetilde{\Delta CS} = \Delta CS$: the location a cancels in the difference of log-sums, and the scale b cancels between the numerator $\tilde{\mu} = b\mu$ and the denominator $\tilde{\alpha}_0 = b\alpha_0$. Money-metric CS changes in the standard quasilinear logit model are therefore normalization-free. This cancellation relies on the closed-form log-sum structure of logit; in the full random-coefficients model, where ΔCS involves integrals over the heterogeneity distribution ν_{ijt}, the scale parameter b enters the integrand nonlinearly and the cancellation is no longer guaranteed without additional structure (e.g., conditions ensuring that scale enters the welfare integral only through the price coefficient).

Percentage welfare changes, however, are again *not* normalization-free. Since $\Delta CS/CS$ involves the CS level in the denominator, which retains the arbitrary additive constant from the location normalization, the percentage change can again suffer from the same pathology of the temperature-change example in the introduction: the level depends on the zero point, and any ratio or percentage involving that level inherits the arbitrariness.

The framework thus draws a sharp line within the standard demand toolkit: market shares, elasticities, diversion ratios, and money-metric welfare changes are normalization-free; utility levels and percentage welfare changes are not.

This distinction is consequential for applied work and may be non-obvious at first. For adaptations of the standard multinomial choice or BLP model to specialized contexts, where additional model ingredients, structures, and assumptions are introduced, the boundary between normalization-free and normalization-dependent quantities can shift in subtle and complicated ways, making explicit verification through the equivalence-class criterion all the more important.

3.3 Dyadic Network Formation Models

Dyadic network formation models have been important tools for the empirical analysis of social and economic networks. The interplay between individual heterogeneity, homophily, and unobserved shocks in the network setting also creates a richer equivalence class structure than in the binary or multinomial choice settings, making the normalization question both more consequential and more subtle.

Consider the following dyadic network formation model à la Graham (2017), more

specifically in the semiparametric form as considered in Gao (2020):

$$D_{ij} = \mathbb{1}\{w(X_i, X_j) + A_i + A_j \geq U_{ij}\}, \qquad U_{ij} \sim_{\text{iid}} F, \qquad (10)$$

where $w(\cdot)$ is a (possibly nonparametric) homophily effect, A_i are individual fixed effects, and F is an unknown CDF. For any $a, b \in \mathbb{R}$ and $c > 0$, the positive affine transformation

$$\hat{A}_i = cA_i + a, \quad \hat{w}(x_i, x_j) = cw(x_i, x_j) + b, \quad \hat{U}_{ij} = cU_{ij} + 2a + b \qquad (11)$$

preserves the link indicator and all maintained assumptions (symmetry of w, the i.i.d. structure of U_{ij}, independence). The equivalence class has *three* degrees of freedom: two location (a, b) and one scale (c). Here we have one more degree of freedom than in binary choice, because the homophily function w and the fixed effects A_i can absorb separate location shifts.

Following Gao (2020), one can normalize by setting $w(\bar{x}, \bar{x}) = 0$ (absorbing b), $F^{-1}(\alpha) = 0$ (absorbing a), and $F^{-1}(1 - \alpha) - F^{-1}(\alpha) = 1$ (absorbing c) for some $\alpha \in (0, 1/2)$. This *two-quantile normalization*, fixing two quantiles of the unobserved distribution rather than its mean and variance, has notable advantages over the standard alternatives. Variance normalization $(\text{Var}(U_{ij}) = 1)$ requires finite variance, excluding heavy-tailed distributions such as Cauchy; mean-zero normalization $(\mathbb{E}[U_{ij}] = 0)$ requires finite first moments. By contrast, the two-quantile normalization requires only that F be continuous and strictly increasing on the relevant range, a condition already maintained in the model, and is therefore applicable under minimal regularity. This illustrates a broader principle: the choice among WLOG-equivalent normalizations, while irrelevant for identification, can matter for the regularity conditions available to the subsequent estimation and inference program.

Linking probabilities $\mathbb{P}(D_{ij} = 1 \mid X_i, X_j, A_i, A_j)$ are normalization-free, as is the shape of w (the equivalence class preserves w up to $cw + b$, so curvature and relative patterns survive) and the ranking of fixed effects across individuals (since $\hat{A}_i = cA_i + a$ preserves orderings). The absolute levels of A_i and $w(x_i, x_j)$, by contrast, depend on a, b, and c, and similarly percentage changes in $w(x_i, x_j)$ are normalization-dependent.

A useful practice, adopted in Gao (2020), is to carry out identification arguments with the normalization constants left as free unknowns and present results as equivalence classes under positive affine transformations. This "gives back" the normal-

ization and makes transparent which properties of (w, A, F) are genuinely identified and which are artifacts of coordinate choice.

The analysis in Gao (2020) also illustrates a subtle but important point: while normalization is "useless" for identification in the sense of Theorem 1 (i.e., it cannot create or destroy identification of any counterfactual), a judicious choice of normalization can make the identification *argument* substantially easier. In the network setting, conditional linking probabilities are naturally expressed as CDF values, which are in turn linked to quantiles via the inverse transformation. The two-quantile normalization $F^{-1}(\alpha) = 0$ and $F^{-1}(1 - \alpha) = 1$ provides two "free starting points" for the unknown CDF, as known anchor values from which the rest of F can be constructed step by step. These anchors are later "given back" when the final result is stated as an equivalence class, but their temporary availability streamlines the constructive identification proof. This analytical convenience is fully consistent with the WLOG theorem: normalization changes no identification *result*, but it can simplify the *path* to that result.

Two extensions of the network model considered in Gao (2020) further illustrate how richer structure can reduce indeterminacy. First, under the parametric specification $w(x_i, x_j) = \tilde{w}(x_i, x_j)'\boldsymbol{\beta}_0$ with $\tilde{w}(x, x) := |x - x| = 0$ by construction, the first location normalization $w(\bar{x}, \bar{x}) = 0$ becomes automatic and no longer consumes a degree of freedom. Only the two normalizations on F are needed, reducing the equivalence class from three to two degrees of freedom. Second, if the additive coupling $A_i + A_j$ is replaced by a known nonlinear function $\phi(A_i, A_j)$, the curvature of ϕ restricts the class of admissible transformations. The positive affine transformation $\hat{A}_i = cA_i + a$ induces $\phi(cA_i + a, cA_j + a)$, which generally differs from $c\phi(A_i, A_j) + (\text{affine terms})$ unless ϕ is additively separable.

This illustrates the general principle that strengthening model structure shrinks the equivalence classes and makes more counterfactuals normalization-free. Of course, stronger structure comes at a cost: identification becomes more reliant on the model, and concerns about model misspecification become more salient. The framework clarifies this trade-off: richer assumptions buy smaller equivalence classes and more normalization-free counterfactuals, but the validity of those counterfactuals is only as good as the maintained assumptions that define the classes.

4 Normalization, Singularity, and Model Modification

Theorem 1 shows that genuine normalization, which selects one representative from each modeling equivalence class, is always WLOG for identification. In practice, however, operations that go beyond representative selection are sometimes performed under the same label.

Section 4.1 considers singularities in the target functional: extending a counterfactual beyond the domain where it is naturally well-defined, typically to "patch" a boundary singularity, changes the estimand, not just the coordinates. Section 4.2 considers singularities in the normalized parameter space: the regularity conditions (topological, metric, and compactness assumptions) imposed for estimation and inference are stated on the normalized space, not on the equivalence classes. Because these conditions are properties of the normalization rather than the equivalence classes, different normalizations that are equally valid for identification can lead to different, and sometimes contradictory, asymptotic conclusions.

The two cases are, in a sense, flip sides of the same singularity story: domain augmentation attempts to repair a pre-existing singularity in the target functional, while certain normalizations *create* a new singularity in the parameter space. In both cases, the framework provides a clean diagnostic: the operation goes beyond genuine normalization whenever it fails to preserve the equivalence-class structure.

4.1 Singularity in the Target Functional

A *boundary singularity* arises when a counterfactual q_0 is normalization-free on a "regular" subset of the parameter space but cannot be coherently extended to the boundary. When the target functional is only partially defined, well-defined on a regular domain but singular at the boundary, any attempt to extend it to the full space faces a tension among three desiderata: *fidelity* (the extension agrees with the original on the regular domain), *invariance* (the extension is $Q_\sim$-measurable on the full space), and *regularity* (the extension is finite-valued and continuous at the boundary). The following result shows this tension is unavoidable whenever the quotient-space functional has no well-behaved limit at the boundary.

Theorem 2 (Boundary Extension Trilemma). *Let $\Theta_0 \subsetneq \mathcal{U} \times \Gamma$ be an open invariant subset (a union of modeling equivalence classes). Let $q_0 : \Theta_0 \to \mathbb{R}$ be continuous and $Q_\sim$-measurable, so that $q_0 = \bar{q}_0 \circ \pi$ for a map $\bar{q}_0 : \pi(\Theta_0) \to \mathbb{R}$, where π is the canonical projection in* (4). *If $\bar{q}_0$ does not admit a* **finite continuous extension** *to some $p^* \in \overline{\pi(\Theta_0)} \setminus \pi(\Theta_0)$, then* **no** *map $\tilde{q} : \mathcal{U} \times \Gamma \to \mathbb{R}$ can simultaneously satisfy:*

  (a) Fidelity*: $\tilde{q}\big|_{\Theta_0} = q_0$.*
  (b) Invariance*: $\tilde{q}$ is $Q_\sim$-measurable on $\mathcal{U} \times \Gamma$.*
  (c) Regularity*: the induced map on the quotient $\bar{\tilde{q}}$ is finite and continuous at p^*.*

Proof. If $\tilde{q}$ satisfies (b), it factors as $\tilde{q} = \bar{\tilde{q}} \circ \pi$ for some $\bar{\tilde{q}} : Q_\sim \to \mathbb{R}$. By (a), $\bar{\tilde{q}} = \bar{q}_0$ on $\pi(\Theta_0)$. By (c), $\bar{\tilde{q}}$ is finite and continuous at p^*. Hence $\bar{\tilde{q}}$ is a finite continuous extension of $\bar{q}_0$ to p^*, contradicting the hypothesis. $\qquad\square$

A practical sufficient condition for the hypothesis is that there exist sequences $p_n, r_n \in \pi(\Theta_0)$ with $p_n \to p^*$ and $r_n \to p^*$ but $\lim \bar{q}_0(p_n) \neq \lim \bar{q}_0(r_n)$ (including the case where one limit diverges), which we call the *non-unique limit test* on the quotient. One response to the trilemma is to shrink the domain to exclude p^*, sacrificing generality rather than any of the three properties. Augmentation takes the opposite approach: it extends the functional to the boundary, but typically sacrifices invariance (b) and may also sacrifice exact fidelity (a), producing a well-defined number for a different, normalization-dependent object. Proposition 2 below identifies one algebraic mechanism behind the singularity: a fixed point of the transformation family that generates the non-extendability condition.

The "logs with zeros" problem studied by Chen and Roth (2024) provides a vivid illustration. The log transformation $\log(Y)$, commonly used to express percentage effects, is well-defined for strictly positive Y and yields a scale-invariant ATE (since $\log(aY_1) - \log(aY_0) = \log Y_1 - \log Y_0$), but has a singularity at $Y = 0$. Any extension, such as choosing c in $\log(c + Y)$, or switching to $\text{arcsinh}(Y)$, attempts to repair this singularity, but imports a new functional that does not coincide with $\log Y$ on the positive reals and whose induced ATE is not scale-invariant. The following result provides a complementary viewpoint of the algebraic mechanism behind the singularity.

Proposition 2 (Fixed-Point Singularity). *Let $\mathcal{G}$ be a collection of transformations on a space $\mathcal{Y}$, let $\rho : \mathcal{G} \to \mathbb{R}$ assign a real number to each transformation, and let*

$m : \mathcal{Y} \setminus \{p\} \to \mathbb{R}$ *satisfy*

$$m(g \cdot y) = \rho(g) + m(y), \qquad \forall\, g \in \mathcal{G},\ y \in \mathcal{Y} \setminus \{p\}. \tag{12}$$

Suppose $\rho(g) \neq 0$ for some $g \in \mathcal{G}$, and that $p \in \mathcal{Y}$ is a fixed point in the sense of $g \cdot p = p$ for all $g \in \mathcal{G}$. Then:

(i) *m cannot be extended to p: no value $m(p) \in \mathbb{R}$ is consistent with (12).*

(ii) *For any extension $\tilde{m} : \mathcal{Y} \to \mathbb{R}$ that agrees with m on $\mathcal{Y} \setminus \{p\}$, the induced difference $\tilde{q}(y_1, y_0) := \tilde{m}(y_1) - \tilde{m}(y_0)$ is not invariant under $\mathcal{G}$ whenever one argument equals p.*

Proof. For (i): if $m(p)$ were defined, then $m(p) = m(g \cdot p) = \rho(g) + m(p)$ for all $g \in \mathcal{G}$, giving $\rho(g) = 0$ for all g, contradicting the assumption that $\rho(g) \neq 0$ for some g. For (ii): set $y_0 = p$ and let $g \in \mathcal{G}$ satisfy $\rho(g) \neq 0$. Then $\tilde{q}(g \cdot y_1, g \cdot p) = \tilde{q}(g \cdot y_1, p) = \tilde{m}(g \cdot y_1) - \tilde{m}(p) = [\rho(g) + \tilde{m}(y_1)] - \tilde{m}(p) = \rho(g) + \tilde{q}(y_1, p) \neq \tilde{q}(y_1, p)$. $\qquad\square$

In the setting of Chen and Roth (2024), $\mathcal{G}$ is the collection of scaling transformations $y \mapsto ay$ for $a > 0$ on $\mathcal{Y} = [0, \infty)$, the fixed point is $p = 0$, and $m = \log$ satisfies (12) on $(0, \infty)$ with $\rho(a) = \log a$. Proposition 2(i) says that no value can be assigned to $\log(0)$ consistently with scale equivariance. Part (ii) says that even in the best case, where $\tilde{m}$ agrees exactly with $\log$ on $(0, \infty)$ and assigns an arbitrary finite value $\tilde{m}(0) = c$, the induced ATE $\mathbb{E}[\tilde{m}(Y_1) - \tilde{m}(Y_0)]$ is not scale-invariant whenever the treatment moves some individuals onto or off the boundary $Y = 0$. The commonly used "log-like" transformations such as $\log(1 + y)$ and $\operatorname{arcsinh}(y)$ do not even satisfy this best case: they differ from $\log y$ on the entire positive real line, so they sacrifice fidelity in addition to invariance. This is the further critique made in Chen and Roth (2024), which our Proposition 2 does not directly capture.

That said, our extension trilemma (Theorem 2) approaches the issue from a complementary perspective based on our broader framework. In terms of Theorem 2, the quotient functional $\bar{q}_0$ (the ATE restricted to strictly positive outcomes) is well-defined on the interior of the quotient but does not extend continuously to the boundary class where $P(Y = 0) > 0$: approximating the atom at 0 by mass at $\delta > 0$ contributes $P(Y_0 = 0) \cdot \log(a\delta)$ to the ATE of the a-scaled element, which shifts by $P(Y_0 = 0) \cdot \log a$ across elements of the same boundary class, so the non-unique limit test is satisfied. Any completion therefore achieves regularity (c) at the boundary but necessarily sacrifices fidelity (a), invariance (b), or both.

The structural models of Section 3 further illustrate the broader relevance of Theorem 2 beyond Proposition 2 and fixed-point singularities. Percentage welfare changes are normalization-dependent through a different mechanism: an *interior* failure rather than a boundary failure. The positive affine family $y \mapsto a + by$ has no fixed point, so Proposition 2 does not apply. Instead, the location shift a enters the utility *level*, making the denominator of any percentage normalization-dependent. The percentage is not $Q_\sim$-measurable even on its natural domain: it varies within equivalence classes, so Lemma 1 already implies it cannot be point-identified. The two settings thus illustrate complementary challenges for identification and normalization: Theorem 2 governs the *boundary* case (a normalization-free functional that resists extension), while Lemma 1 governs the *interior* case.

The takeaway from this subsection is that if a purported "normalization" is needed to make the target functional well-defined, the problem lies in the definition of the target, or in the model, not in the parameterization. This warning is not merely abstract. In dynamic latent-factor models of skill formation, Agostinelli and Wiswall (2025) show that period-by-period re-normalization can be over-identifying and can bias estimation when the production technology already has known location and scale, while Freyberger (2025) shows more broadly that seemingly innocuous scale and location restrictions can constrain identified parameters and alter counterfactuals in a model-dependent way. In the language of the present paper, these are precisely cases in which an operation marketed as a normalization is actually doing substantive modeling work.

4.2   Singularity in the Normalized Parameter Space

Domain augmentation confronts a pre-existing singularity in the *target functional*. In addition, there are cases where the singularity is not inherited but *introduced* by the "normalization", for example, dividing by a designated coefficient creates a coordinate singularity, distorting the geometry of the parameter space. The induced singularity may interact with further model modification in the form of regularity conditions (topological, metric, smoothness) when they are imposed for estimation and inference purposes. These conditions are substantive restrictions typically stated on the normalized parameter space rather than on the equivalence classes, and they need not be preserved across different normalizations of the same classes. Two normalizations

that are both valid for identification purposes can therefore induce sharply different parameter-space geometries, with direct consequences for the validity of standard asymptotic arguments. This point is closely related to Hamilton, Waggoner, and Zha (2007), who show in likelihood-based models that poor normalizations can create multimodal posteriors, disjoint confidence regions, and misleading descriptions of statistical uncertainty. Our contribution here is complementary: even when two normalizations are equally WLOG for identification, a coordinate chart can still distort the topology and metric imposed downstream on the quotient space.

We illustrate this using the semiparametric binary choice model à la Manski (1975). Specifically, consider

$$Y_i = \mathbb{1}\{X_i'\boldsymbol{\beta}_0 \geq \varepsilon_i\}, \quad \boldsymbol{\beta}_0 \in \mathbb{R}^D \setminus \{0\},$$

which has scale indeterminacy: each equivalence class is a ray $\{c\boldsymbol{\beta}_0 : c > 0\}$ in $\mathbb{R}^D \setminus \{0\}$, and any normalization that selects one representative per ray is WLOG for identification by Theorem 1. In the econometric literature, there are two commonly used approaches for scale normalization: the *coordinate normalization* and the *sphere normalization*, which we describe below.

4.2.1 Special-Coordinate versus Sphere Normalization

The most commonly adopted approach in the semiparametric binary choice literature, or what we call *special-coordinate normalization* here, normalizes the coefficient on a "special regressor": assume $\beta_{0,1} \neq 0$ in addition and set $|\beta_{0,1}| = 1$, so that the normalized parameter space becomes

$$\tilde{\boldsymbol{\beta}}_0 := \left(\text{sgn}(\beta_{0,1}),\ \frac{\beta_{0,2}}{|\beta_{0,1}|},\ \ldots,\ \frac{\beta_{0,D}}{|\beta_{0,1}|}\right)' \in \tilde{B} := \{1, -1\} \times \mathbb{R}^{D-1}. \tag{13}$$

The normalized parameter space $\tilde{B}$ consists of two copies of $\mathbb{R}^{D-1}$ (one for each sign of $\beta_{0,1}$), and the natural metric is the Euclidean norm on $\mathbb{R}^{D-1}$ applied to the remaining coordinates, with the sign handled separately. It is better understood as a pair of branchwise affine charts on the open set $\{\beta_{0,1} \neq 0\}$ than as a globally innocuous normalization of the quotient space. If the sign of $\beta_{0,1}$ is separately identified, one may work on a single branch such as $\beta_{0,1} = 1$ or $\beta_{0,1} = -1$; if not, one can keep both branches through $|\beta_{0,1}| = 1$. Either way, the deeper issue is unchanged: the chart

becomes singular as $\beta_{0,1} \to 0$. This approach has been widely adopted since Horowitz (1992) in the literature on semiparametric binary and discrete choice models (see also, e.g., Lewbel, 2000), because it effectively reduces the parameter space to $\mathbb{R}^{D-1}$, on which standard asymptotic theory applies without the need to deal with any spherical geometry.

An alternative normalization approach, which was used in Manski (1975, 1985) and is referred to as *sphere normalization* here, is to project the parameter space onto the unit sphere:

$$\bar{\boldsymbol{\beta}}_0 := \frac{\boldsymbol{\beta}_0}{\|\boldsymbol{\beta}_0\|} \in \mathbb{S}^{D-1} := \{v \in \mathbb{R}^D : \|v\| = 1\}.$$

This normalization was adopted in more recent work by Gao and Li (2026) in a panel multinomial choice setting. The relevant normalized parameter space is the sphere itself, equipped with the *great-circle (angular) metric*:

$$\rho_{GC}(\bar{\boldsymbol{\beta}}, \underline{\boldsymbol{\beta}}) := \arccos(\bar{\boldsymbol{\beta}}'\underline{\boldsymbol{\beta}}) \in [0, \pi], \tag{14}$$

which measures the angle between two unit vectors and reflects the intrinsic geometry of the quotient space of rays. The great-circle metric and the Euclidean metric on $\mathbb{S}^{D-1}$ are *strongly equivalent*: for all $\bar{\boldsymbol{\beta}}, \underline{\boldsymbol{\beta}} \in \mathbb{S}^{D-1}$,

$$\|\bar{\boldsymbol{\beta}} - \underline{\boldsymbol{\beta}}\| \leq \rho_{GC}(\bar{\boldsymbol{\beta}}, \underline{\boldsymbol{\beta}}) \leq \frac{\pi}{2}\|\bar{\boldsymbol{\beta}} - \underline{\boldsymbol{\beta}}\|, \tag{15}$$

so the two metrics generate the same topology, the same notions of convergence, and the same uniform continuity.

4.2.2 Implications for Estimation and Inference

Both normalizations are WLOG for identification under their stated assumptions, but some of their consequences diverge, especially once one imposes additional regularity conditions needed for estimation and inference. Four differences are particularly consequential.

Restrictiveness. The sphere $\mathbb{S}^{D-1}$ represents all directions in $\mathbb{R}^D \setminus \{0\}$, while the special-coordinate normalization excludes every $\boldsymbol{\beta}_0$ with $\beta_{0,1} = 0$. Even when $\beta_{0,1} \neq 0$ is assumed, the chart becomes singular as $\beta_{0,1} \to 0$: the excluded hyperplane is a set of measure zero, but its influence extends to an entire neighborhood through the

divergence of the normalized coordinates, as illustrated below.

Topological Structures. The two normalizations induce fundamentally different topologies on their parameter spaces. The coordinate-normalized space $\tilde{B} = \{1, -1\} \times \mathbb{R}^{D-1}$ is *disconnected* (the two sign components are disjoint) and *unbounded* (the coordinates $\beta_{0,k}/|\beta_{0,1}|$ diverge as the true parameter approaches the excluded hyperplane $\beta_{0,1} = 0$). Since compactness requires boundedness, the coordinate normalization destroys it. To restore it, as typically required for consistency of extremum estimators, one imposes an ad hoc compactness restriction $\tilde{\boldsymbol{\beta}}_0 \in \{1, -1\} \times [-M, M]^{D-1}$. However, this excludes not only directions with $\beta_{0,1} = 0$ but an entire neighborhood with $|\beta_{0,1}| < 1/M$, a substantive restriction masquerading as a regularity condition. By contrast, the sphere $(\mathbb{S}^{D-1}, \rho_{GC})$ is *connected* (for $D \geq 2$), *bounded*, and *closed*, and hence compact by construction, requiring no auxiliary trimming or boundedness assumptions.

Convergence Properties. Most importantly, the coordinate singularity at $\beta_{0,1} = 0$ can reverse convergence verdicts. We give two illustrations with $D = 2$.

Across sign components. Consider $\tilde{\boldsymbol{\beta}}^{(M)} = (1, M)'$ and $\tilde{\boldsymbol{\gamma}}^{(M)} = (-1, M)'$ in $\tilde{B}$, approaching the excluded hyperplane $\beta_{0,1} = 0$ from opposite sign components. Because $\tilde{\boldsymbol{\beta}}^{(M)}$ and $\tilde{\boldsymbol{\gamma}}^{(M)}$ lie in different connected components of $\tilde{B} = \{1, -1\} \times \mathbb{R}^{D-1}$, the chart topology treats them as permanently separated: no sequence in one component can converge to a point in the other. Yet both represent rays that approach the same direction $(0, 1)'$ as $M \to \infty$. On the sphere, let $\bar{\boldsymbol{\beta}}^{(M)} := \tilde{\boldsymbol{\beta}}^{(M)}/\|\tilde{\boldsymbol{\beta}}^{(M)}\|$ and $\underline{\boldsymbol{\beta}}^{(M)} := \tilde{\boldsymbol{\gamma}}^{(M)}/\|\tilde{\boldsymbol{\gamma}}^{(M)}\|$. Then $\rho_{GC}(\bar{\boldsymbol{\beta}}^{(M)}, \underline{\boldsymbol{\beta}}^{(M)}) \to 0$ as $M \to \infty$: the two directions converge on the sphere, but the disconnectedness of the chart makes this invisible.

Within a sign component. Now consider $\tilde{\boldsymbol{\beta}}^{(M)} = (1, M)'$ and $\tilde{\boldsymbol{\gamma}}^{(M)} = (1, 2M)'$, both in the $\{1\} \times \mathbb{R}$ component. Their coordinate distance is $\rho_{\text{Euc}}(\tilde{\boldsymbol{\beta}}^{(M)}, \tilde{\boldsymbol{\gamma}}^{(M)}) = M \to \infty$, suggesting divergence. Yet both rays converge to $(0, 1)'$ on the sphere, and the great-circle distance between their projections tends to zero: the chart inflates distances near the singularity without bound.

In both cases, what is nonconvergent under the special-coordinate normalization is convergent under the sphere normalization. Any asymptotic result formulated in $(\tilde{B}, \rho_{\text{Euc}})$ carries an implicit dependence on the coordinate chosen for normalization, a form of nonuniformity that the sphere normalization avoids entirely.

Metric Structures. The metric distortion documented above has direct consequences for inference. Standard errors and confidence intervals inherit the Euclidean metric of the normalized coordinate space. Under the special-coordinate normalization, a 95% confidence interval for a parameter is an interval in $\mathbb{R}^{D-1}$. However, the Euclidean metric on $\tilde{B}$ is not the intrinsic metric on the space of equivalence classes (rays from the origin). A given length in $\tilde{B} = \{1, -1\} \times \mathbb{R}^{D-1}$ can correspond to a large arc near the poles ($|\tilde{\boldsymbol{\beta}}_{0,-1}|$ small) but a vanishingly small arc near the equator ($|\tilde{\boldsymbol{\beta}}_{0,-1}| \to \infty$). Consequently, the magnitude of the standard error and the "width" of a confidence interval reported in coordinate-normalized units are chart-dependent and not uniformly comparable to the intrinsic quotient geometry, especially near the singularity. By contrast, on the sphere $(\mathbb{S}^{D-1}, \rho_{GC})$, the strong metric equivalence (15) guarantees that Euclidean and great-circle distances are uniformly comparable, so standard errors retain their geometric meaning across the entire parameter space.

4.2.3 Summary

The binary choice example illustrates a more general lesson: a normalization that introduces a coordinate singularity can have consequences for estimation and inference that go far beyond the choice of representative, because the regularity conditions that underpin asymptotic theory (topological structure, metric structure, uniformity, and smoothness) are imposed on the *normalized parameter space*, not on the equivalence classes themselves.

This is not merely a theoretical curiosity. In any model with scale, location-scale, or more complex indeterminacy, the choice of normalization determines the topology in which estimators converge, the compactness conditions available to the analyst, and the uniformity properties of the resulting inference. Together with Section 4.1, the two subsections illustrate the two faces of the singularity problem that normalization can encounter. Domain augmentation attempts to repair a pre-existing singularity: $\log(0)$ is undefined, and any completion that extends the functional to the boundary necessarily sacrifices invariance or fidelity. The coordinate normalization, by contrast, *creates* a singularity: dividing by $\beta_{0,1}$ maps the excluded hyperplane $\{\beta_{0,1} = 0\}$ to infinity, so that the normalized coordinates $\beta_{0,k}/|\beta_{0,1}|$ diverge as $\beta_{0,1} \to 0$, distorting the topology and metric of the parameter space. The framework of this paper, and the concept of end-to-end WLOG normalization introduced in Section 5, provide the language for diagnosing when a normalization that is WLOG for identification may

nonetheless fail to be WLOG for the full statistical analysis.

5   Discussion

The preceding sections show that singularities can arise in two complementary ways: a pre-existing singularity in the target functional may resist extension to the boundary (Section 4.1), or a normalization may introduce a coordinate singularity that distorts the geometry of the parameter space (Section 4.2).

This motivates what we call *end-to-end WLOG normalization*: assessing normalization not in isolation at the identification stage, but jointly with the target counterfactual, the regularity conditions, and any auxiliary structure used downstream. A normalization is end-to-end WLOG only if every downstream object, whether a counterfactual parameter or a regularity condition for estimation and inference, is coherent with respect to the relevant quotient space. When verifying this for every downstream object is impractical, the researcher should at minimum select one specific object of primary interest and check whether the chosen normalization is end-to-end WLOG with respect to that object.

Seen through this lens, the papers discussed in the literature review line up naturally. Hamilton, Waggoner, and Zha (2007) emphasize the downstream inferential geometry of normalization, while Agostinelli and Wiswall (2025) and Freyberger (2025) emphasize that some restrictions advertised as normalizations are in fact over-identifying or misspecifying assumptions once combined with richer model structure. The present framework unifies these insights: a proper normalization is a representative choice on a fixed equivalence class, whereas an operation that changes the target counterfactual, shrinks the class, or distorts the downstream geometry belongs to modeling, not normalization.

The results extend naturally to partial identification: the identified set is a union of equivalence classes, and the WLOG Theorem applies equally. More broadly, richer model structure shrinks the equivalence classes and makes more counterfactuals normalization-free, a useful design principle for model specification.

The framework suggests the following checklist:
1. What are the modeling equivalent transformations, and how many degrees of freedom does each equivalence class have?
2. Is each target counterfactual constant on the relevant equivalence classes?

3. If not, is the number interpreted explicitly as normalization-relative?

Three practical recommendations follow. First, report the equivalence class: if the main result is "point identification under normalization," the reader should be told what the unnormalized class looks like. Second, separate identification from parameterization: the choice of normalization may be irrelevant for economic content yet matter for estimation, inference, and computation (Section 4.2). Third, favor normalization-free counterfactuals: policy questions stated in terms of choice probabilities, entry probabilities, or network statistics are more likely to be invariant and hence structurally identified. This last recommendation aligns with the findings of Freyberger (2025).

6 Conclusion

We have presented a general framework for analyzing the interplay between identification, normalization, and counterfactual parameters in econometric models. The framework provides both general theoretical results and concrete illustrations across binary choice, demand estimation, and network formation, clarifying when a normalization is genuinely without loss and when it quietly changes the object of interest.

The singularity perspective developed in Section 4 suggests that the consequences of normalization extend well beyond identification. A normalization that is innocuous for point identification may nonetheless introduce coordinate singularities that distort the topology and metric of the parameter space, with implications for the validity of regularity conditions, consistency of estimators, and uniformity of inference. A systematic study of which normalizations are "safe" for downstream statistical analysis, and whether singularity-free alternatives exist in broader classes of models, is an important direction for future work.

References

AGOSTINELLI, F. AND M. WISWALL (2016): "Identification of Dynamic Latent Factor Models: The Implications of Re-Normalization in a Model of Child Development," Working Paper 22441, National Bureau of Economic Research.

——— (2025): "Estimating the Technology of Children's Skill Formation," *Journal of Political Economy*, 133, 1168–1217.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

BERRY, S. T. AND P. A. HAILE (2014): "Identification in Differentiated Products Markets Using Market Level Data," *Econometrica*, 82, 1749–1797.

BHATTACHARYA, D., P. DUPAS, AND S. KANAYA (2024): "Demand and Welfare Analysis in Discrete Choice Models with Social Interactions," *Review of Economic Studies*, 91, 748–784.

CHAMBERLAIN, G. (2010): "Binary Response Models for Panel Data: Identification and Information," *Econometrica*, 78, 159–168.

CHEN, J. AND J. ROTH (2024): "Logs with Zeros? Some Problems and Solutions," *Quarterly Journal of Economics*, 139, 891–936.

CILIBERTO, F. AND E. TAMER (2009): "Market Structure and Multiple Equilibria in Airline Markets," *Econometrica*, 77, 1791–1828.

FREYBERGER, J. (2025): "Normalizations and Misspecification in Skill Formation Models," *Review of Economic Studies*, accepted for publication. First version: April 2020.

GAO, W. Y. (2017): "Nonparametric Identification in Index Models of Link Formation," ArXiv:1710.11230v5.

——— (2020): "Nonparametric Identification in Index Models of Link Formation," *Journal of Econometrics*, 215, 399–413.

GAO, W. Y. AND M. LI (2026): "Identification of Semiparametric Panel Multinomial Choice Models with Infinite-Dimensional Fixed Effects," *Review of Economics and Statistics*, forthcoming. Online February 2026.

GRAHAM, B. S. (2017): "An econometric model of network formation with degree heterogeneity," *Econometrica*, 85, 1033–1063.

GU, J., T. M. RUSSELL, AND T. STRINGHAM (2025): "Counterfactual Identification and Latent Space Enumeration in Discrete Outcome Models," *Review of Economic Studies*, forthcoming.

HAMILTON, J. D., D. F. WAGGONER, AND T. ZHA (2007): "Normalization in Econometrics," *Econometric Reviews*, 26, 221–252.

HOROWITZ, J. L. (1992): "A smoothed maximum score estimator for the binary response model," *Econometrica: journal of the Econometric Society*, 505–531.

KALOUPTSIDI, M., P. T. SCOTT, AND E. SOUZA-RODRIGUES (2021): "Identification of Counterfactuals in Dynamic Discrete Choice Models," *Quantitative Economics*, 12, 351–403.

LEWBEL, A. (2000): "Semiparametric Qualitative Response Model Estimation with Unknown Heteroscedasticity or Instrumental Variables," *Journal of Econometrics*, 97, 145–177.

——— (2019): "The Identification Zoo: Meanings of Identification in Econometrics," *Journal of Economic Literature*, 57, 835–903.

MAGNAC, T. AND D. THESMAR (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70, 801–816.

MANSKI, C. F. (1975): "Maximum score estimation of the stochastic utility model of choice," *Journal of econometrics*, 3, 205–228.

——— (1985): "Semiparametric analysis of discrete response: Asymptotic properties of the maximum score estimator," *Journal of econometrics*, 27, 313–333.

——— (1988): "Identification of Binary Response Models," *Journal of the American Statistical Association*, 83, 729–738.

MCFADDEN, D. (1974): "Conditional Logit Analysis of Qualitative Choice Behavior," *Frontiers in Econometrics*, 105–142.

RUST, J. (1994): "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics*, Elsevier, vol. 4, 3081–3143.

TAMER, E. (2003): "Incomplete Simultaneous Discrete Response Model with Multiple Equilibria," *Review of Economic Studies*, 70, 147–165.