

“Post” Pre-Analysis Plans: Valid Inference for Non-Preregistered Specifications*

Reca Sarfati[†]

Vod Vilfort[‡]

October 2, 2025

Abstract

Pre-analysis plans (PAPs) have become standard in experimental economics research, but it is nevertheless common to see researchers deviating from their PAPs to supplement preregistered estimates with non-prespecified findings. While such ex-post analysis can yield valuable insights, there is broad uncertainty over how to interpret—or whether to even acknowledge—non-preregistered results. In this paper, we consider the case of a truth-seeking researcher who, after seeing the data, earnestly wishes to report additional estimates alongside those preregistered in their PAP. We show that, even absent “nefarious” behavior, conventional confidence intervals and point estimators are invalid due to the fact that non-preregistered estimates are only reported in a subset of potential data realizations. We propose inference procedures that account for this conditional reporting. We apply these procedures to [Bessone et al. \(2021\)](#), which studies the economic effects of increased sleep among the urban poor. We demonstrate that, depending on the reason for deviating, the adjustments from our procedures can range from having no difference to an economically significant difference relative to conventional practice. Finally, we consider the robustness of our procedure to certain forms of misspecification, motivating possible heuristic checks and norms for journals to adopt.

*We are grateful to Alberto Abadie, Isaiah Andrews, and Anna Mikusheva for their guidance and support throughout this project. We thank Sarah Moon, Ashesh Rambachan, and Jonathan Roth for valuable feedback and suggestions. We also thank Abhijit Banerjee, Esther Duflo, Ben Olken, and participants in lunch workshops at MIT for early conversations that helped situate this work within applied research settings. Finally, we are especially grateful to Frank Schilbach for generous and repeated feedback through the life of this project. Vod gratefully acknowledges support from the Jerry A. Hausman Fellowship.

[†]Department of Economics, Massachusetts Institute of Technology, sarfati@mit.edu.

[‡]Department of Economics, Massachusetts Institute of Technology, vod@mit.edu.

Introduction

Pre-analysis plans (PAPs) are time-stamped, publicly accessible documents that preregister hypotheses, experimental design, and planned statistical analyses before data are observed. Originating in clinical trials, PAPs have been championed as a key tool for improving credibility and transparency in experimental economics (De Angelis et al., 2004; Olken, 2015). Following the 2013 launch of the American Economic Association’s RCT Registry—a central repository tracking ongoing, completed, and withdrawn trials—the number of preregistrations with PAPs has grown remarkably (Ofosu and Posner, 2023). Despite this growth, uncertainty remains about their requisite detail and comprehensiveness, as well as the extent to which researchers should be allowed to deviate from preregistered choices (Olken, 2015; Banerjee et al., 2020a). In a 2023 survey of experimental economists, Imai et al. (2025) report that while 83% support deviations provided transparent disclosure, views diverge sharply on acceptable scope—ranging from unrestricted latitude to narrowly circumscribed departures. Overall, 74% expressed a desire for clearer norms and guidance for how PAPs should be drafted and reviewed.

At root, the core tension underlying arguments for and against PAPs turns on the permitted scope of deviation: On the one hand, strict adherence limits researchers’ degrees of freedom for data mining, specification searching, and ex post rationalization (Simmons et al., 2011; Wicherts et al., 2016). Formally, strict adherence sets the analysis rule ex ante, ensuring nominal frequentist guarantees (e.g., size, coverage) hold under identical replications of the experiment (Olken, 2015; Kasy and Spiess, 2023). On the other hand, strict adherence constrains researchers’ flexibility to handle unforeseen issues that arise during implementation, as well as ability to explore novel insights that emerge over the course of the study. Considering the high overhead costs to running a new experiment, not pursuing auxiliary analyses when the data suggest them also comes at a serious social cost (Miguel, 2021; Coffman and Niederle, 2015; Olken, 2015; Banerjee et al., 2020a). In practice, writing a PAP that is sufficiently exhaustive to anticipate all possible contingencies is extremely time and labor intensive, and deviations are extremely common (Ofosu and Posner, 2020, 2023; Brodeur et al., 2024). The pervasiveness of such deviations presents an urgent need to understand how the reporting of preregistered estimates (henceforth, *PAP estimates*) formally differs from that of non-preregistered estimates (henceforth, *post estimates*), and by extension, how the latter ought to be presented in papers and interpreted by readers.

In this paper, we develop a framework to formalize the statistical properties of deviations from the PAP, thereby allowing us to provide certain inferential guarantees for the resulting post estimates. In doing so, we offer a means to resolve the core tension above: Researchers can still enjoy the benefits of PAP estimates, while leaving open the option to validly report and interpret post estimates. Our approach frames deviations as instances of *selective reporting*, which differ from prespecified analyses specifically in that the decision to report can depend on realized

data. For example, a non-preregistered regression coefficient might be of interest only when it is found to be significant and large; a new specification pooling treatment arms might be reported only if the preregistered specification fails to reject the null; alternately, a non-preregistered economic model might be tested only if the data violate a core assumption of the preregistered model. In modeling deviations in this way, we formalize an existing norm where researchers are expected to provide a rationale for deviating. Here, this rationale is expressed through what we term the *deviation set*, which corresponds to the subset of data realizations satisfying the conditions for reporting a given post estimate. Using the deviation set, we show how one may obtain confidence intervals and point estimators that provide valid inference *conditional* on the event that the researcher deviates. We argue that researchers should leverage this deviation set to report corrected (i.e., conditional) inferences alongside conventional (i.e., unconditional) inferences. We therefore provide general inference procedures for obtaining confidence intervals that have correct conditional coverage and point estimators that are conditionally unbiased.

The first and primary contribution of this paper is the framing of PAP deviations as selective reporting, which provides a formal language for addressing the concerns faced by researchers when reporting post estimates. In viewing deviations through this lens, our framework is able to leverage results and algorithms from the conditional inference literature to correct for the statistical distortions arising from PAP deviations, ensuring valid inference for post estimates.

To develop our results, we consider a model where the estimates are normally distributed with a known covariance matrix. The inference procedures that we propose are based on the observation that, since a deviation only occurs in the data realizations from its deviation set, the corresponding post estimates are no longer normally distributed, owing to dependence between the post estimate and its deviation event. In the leading case where the deviation set is based on realizations of PAP estimates, this dependence is governed by the correlation between the PAP and post estimates. After accounting for this dependence, the distribution of post estimates corresponds to a class of truncated normal distributions, with truncation regions that depend on the deviation set. We show that for a broad class of polyhedral deviation sets, which encompass the deviations that occur in practice, this yields computationally tractable test statistics from which conditionally valid confidence intervals and point estimators can be derived—results from Pfanzagl (1994) further imply the optimality of these procedures for our setting.

As a second contribution, we relax the assumptions of normality and known covariance to show that feasible analogues of the proposed procedures—where one plugs in asymptotically normal estimates and consistent covariance matrix estimators—are asymptotically valid in large samples, uniformly over a broad class of data generating processes. Thus, researchers can safely implement these plug-in procedures, which we detail explicitly. While such uniformity results have been established in Markovic et al. (2017); Tian and Taylor (2017); Andrews et al. (2024); and McCloskey (2024) for the case of one polyhedra, our results establish uniformity for the gen-

eral case of multiple polyhedra.¹ In proving uniformity for unions of multiple polyhedra, we allow for the complexity and expressiveness of deviation sets that arise in practice—for instance, conditioning on two-sided significance requires conditioning on a union of two polyhedra. This uniformity result is of independent interest, beyond the PAP setting. For example, the selective inference literature is often interested in conditional inference for lasso model coefficients conditional on the lasso selection event, which can be represented as a union of polyhedra (Lee et al., 2016).

We apply our framework to Bessone et al. (2021), which examines the economic effects of increased sleep among the urban poor. We construct three plausible deviation sets, based directly on text from the paper, and show that—depending on the deviation set and the realized data—accounting for conditional reporting can range from no change to economically meaningful departures from conventional practice. Intuitively, when the realized data lie near a reporting boundary implied by the deviation set, conditioning can have a large impact; when the data lie comfortably inside the region, the adjustments to conventional inference are minimal.

As an additional contribution, we examine the robustness of our procedure to certain forms of misspecification of the deviation set, encompassing cases of both strategic misreporting and earnest mistakes. First, we show that inference remains valid—albeit less precise—when the researcher can only describe a subset of their deviation set local to the realized data. Second, we present sensitivity analysis to assess robustness across alternative deviation sets, motivating potential norms that journals and researchers should adopt, such as agreed-upon significance cutoffs for reporting non-prespecified findings (e.g., a default significance level of 5%).

To ease the adoption of our approach, we give examples of how to specify deviation sets for the most common rationales, demonstrating that the costs to formal articulation are low. For example, if a post estimate is reported because it is significant at the 5% level, then the deviation set is characterized by the significance cutoff. We show that an immediate implication of our framework is that conventional and conditional inferences are equivalent when deviations occur (i) for reasons independent of the data (e.g., arrival of new econometric methods, honest mistakes, or oversights in PAP specification) or (ii) in events that are independent of the post estimates (e.g., deviations based on the results from a pilot study on an independent sample). In these special cases, there is no need to adjust one’s conventional inferences.

Related Literature. Our paper contributes to an ongoing dialogue among economists on the costs and benefits of PAPs; see, for instance, McKenzie (2012), Humphreys et al. (2013), Coffman and Niederle (2015), Olken (2015), Glennerster (2017), Christensen and Miguel (2018), Christensen et al. (2019), Banerjee et al. (2020a). These discussions are accompanied by an

¹Tibshirani et al. (2018) develop uniform asymptotics for a model where the events of interest take the form of unions of polyhedra, but (i) impose that underlying mean parameters shrink in accordance with the sample size and (ii) employ procedures that condition on individual polyhedra before aggregating. The latter can lead to efficiency loss in the normal model, relative to the optimal procedure that we employ (Fithian et al., 2014).

active literature meta-analyzing the empirical success of PAPs in achieving their stated goals of curbing publication bias and p -hacking, such as Brodeur et al. (2020), Brodeur et al. (2023), Ofosu and Posner (2023), and Brodeur et al. (2024). On the whole, existing work has primarily centered around questions of (i) whether to use a PAP in the first place and (ii) its appropriate length, detail, and contents. Less explored, however, are the specific consequences of deviating from a PAP *after* it is already preregistered (i.e., taken as fixed) and data observed. This question is the focus of our paper, isolated from considerations of publication bias and p -hacking.

There is also a literature discussing optimal PAP construction (Ludwig et al., 2019; Banerjee et al., 2020b; Anderson and Magruder, 2022; Kasy and Spiess, 2023). Our conditional inference framework instead takes the PAP as given and asks how one can validly interpret deviations from the PAP. That said, while we do not make formal statements in this paper about optimal PAP specification, we do provide suggestions to researchers and journals for potential norms to adopt. We leave formal consideration of the implications of our framework for optimal PAP construction to future work.

Methodologically, this paper draws from the conditional inference literature, particularly from work on optimal conditionally quantile-unbiased estimation for polyhedral conditioning events (Pfanzagl, 1979; Fithian et al., 2014; Lee et al., 2016; Andrews et al., 2024; McCloskey, 2024). The conditionally quantile-unbiased estimators form the building blocks for our proposed confidence intervals and point estimators. In terms of model setup, the asymptotic framework underlying our uniformity result most closely resembles the approaches of Andrews et al. (2024) and McCloskey (2024), whose results can be used to prove the case of a single polyhedra. Here we prove the general case of multiple polyhedra.

Outline. In Section 1, we begin with a description of the research process surrounding PAP construction and deviations therefrom, using Bessone et al. (2021) as an illustrative application. In Section 2, we introduce notation to formalizes this process; moreover, we interpret deviations through the lens of a simple Bayesian decision problem—though our inference results are valid without the Bayesian structure. In Section 3, we introduce our conditional inference objectives in a model with normally distributed estimates, and present general procedures for constructing confidence intervals and point estimators that are conditionally valid. In Section 4, we show how to implement these procedures in practice, allowing for non-normal estimates. In Section 5, we apply our results to Bessone et al. (2021). In Section 6, we discuss robustness of our proposed procedures to misspecification of the deviation set. We provide all proofs and supplementary results in the Appendix.

1 Setting

We begin with a general description of how PAPs are used in the status quo, accompanied by examples from the experimental economics literature. This descriptive timeline will motivate the formal model presented in Section 2. Throughout this paper, we focus on PAPs in the context of randomized control trials (RCTs), as (i) this is where PAPs are most used in economics, and (ii) it offers a well-defined division between preregistration and data collection. That said, these discussions in principle also apply to empirical research using preexisting datasets (Burlig, 2018; Miguel, 2021), though such cases generally raise the additional challenge of verifying that researchers are indeed preregistering studies before analyzing the data² (Christensen and Miguel, 2018). To avoid distracting from the central focus of this paper, we take the delineation of “pre” and “post” data collection as given.

Priors. Researchers embark on projects with priors informed by sources such as economic theory, existing literature, outcomes of a pilot study, discussions with experts, or first-hand observation. For example, Bessone et al. (2021) benchmark the anticipated magnitude of their results as follows:

“While experimental evidence on the effect of increasing sleep in field settings is scarce, there is a widely held belief among researchers and the public that reducing sleep deprivation would lead to improvements in economic outcomes (Walker, 2017). To document these priors, we surveyed 119 experts from sleep science and economics who predicted sizable economic benefits, including a 7% increase in work output, of increasing sleep by half an hour a night from the low levels observed in our setting.”

At the highest level, priors can inform the researcher’s choice of which hypotheses to test, estimands to consider, and outcomes to measure given theoretical salience or expected precision and effect size. Because empirical projects are constrained in scale and scope, priors can further discipline downstream design choices: how large of a sample to field; how many waves of follow-up to fund; which populations to emphasize; whether to invest in costly data enhancements; and how to allocate statistical power between estimating overall effects and probing heterogeneity or mechanism analysis.

Preregistration. In advance of data collection, the researcher preregisters information about their study such as primary outcomes, experimental design, randomization method, randomization unit, clustering, and sample size (including total number of observations, number of

²There do exist non-experimental settings where researchers are required to prespecify analyses prior to data access. For instance, restricted-use microdata such as those accessed through the U.S. Census Bureau often require submission and approval of detailed research proposals before any data extracts are made, with each additional data request or modification of scope typically necessitating a new proposal.

clusters, and units per treatment arm). In addition to this information, the researcher may also register a PAP stipulating which specifications will be used to estimate quantities of interest. In practice, depending on the research topic and journal, there is variation in whether a PAP is required, and if so, to what level of specificity. At its simplest, a PAP might consist of a single list of linear regressions. On another extreme, a PAP might take the form of a complete contingent plan which maps every possible data realization to a particular set of specifications, mirroring the notion of a complete strategy in game theory.

To illustrate, consider a simple example, shown in Figure 1. After estimating the average treatment effect (ATE) of a deworming intervention on school attendance, a complete contingent PAP might stipulate that if the attendance effect is statistically insignificant, the researcher will analyze a particular set of secondary outcomes such as health measures (e.g., incidence of anemia, weight-for-age) to assess whether the intervention improved child well-being through channels other than schooling. By contrast, if the effect *is* significant, the PAP specifies a set of persistence checks to perform, such as attendance in the following academic year.

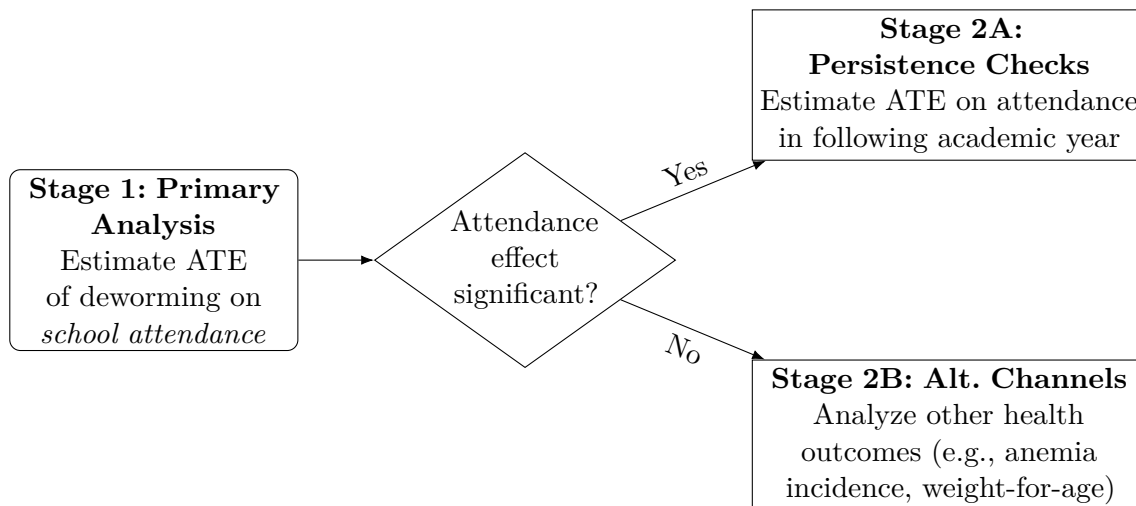


Figure 1: Complete contingent PAP for deworming intervention.

In the sample PAP above, the researcher not only specifies their primary analysis, but also their intended secondary analysis for both possible contingencies based on their initial findings. Even the simplest of empirical exercises, however, can be far more complex than the contingencies shown in Figure 1. Operationally, most PAPs in practice lie somewhere in the middle: The researcher has a core set of primary specifications and also stipulates how these specifications will be adjusted for some—but not all—contingencies.

An example of a partial contingent plan can be found in the (31-page) PAP for [Bessone et al. \(2021\)](#). One component of their analysis tests whether sleep affects how attentive individuals are to earnings-related incentives. The authors approached this question by introducing a “salience variation” in a particular data-entry task. The plan was first to establish that par-

ticipants respond more strongly when incentives are presented in a high-salience rather than a low-salience way, and then to assess whether the sleep intervention reduces this gap.³ However, it was unknown *ex ante* whether (i) the salience manipulation would succeed and (ii) what form the response would take. As directly stated in their PAP:

“We do not fully pre-specify this analysis since the appropriate analysis will depend upon first establishing that the salience variation works “as intended”—that is, the response to incentives is stronger under high-salience relative to low-salience. Moreover, the precise form of the response to salience—e.g. whether individuals notice and respond to the incentive change in 5 minutes versus 30 minutes—remains unknown at the time of this pre-registration, making it difficult to write down the full contingent analysis plan.”

Data Collection and Analysis. After completing the preregistered analyses, a researcher may wish to go beyond the PAP and compute additional results. Common motivations might include exploratory work prompted by unexpected findings, knowledge gained during data collection, new statistical tools, or unanticipated issues with the original specification. As an instance of the latter, [Bessone et al. \(2021\)](#) write:

“We preregistered daily net savings as our main variable of interest. However, this measure suffers from an unanticipated design issue: participants make large one-time withdrawals right before the study ends, which mechanically drives down net savings. We believe deposits more accurately reflect differences in savings behavior, and the accrued interest captures the benefit of savings.”

More generally, because PAPs reflect prior beliefs about where the most interesting or strongest effects are likely to lie, the desire for additional analysis is particularly likely when the realized data substantially shift the researcher’s prior. The next section formalizes this idea.

Reporting of Results and Deviations from PAP. The researcher publicly reports the PAP and the results of all preregistered analyses.⁴ At this time, there is broad uncertainty in the profession about whether and how to report *non*-preregistered results. Prescriptions range from strict “no reporting” to more permissive approaches that allow reporting conditional on clearly flagging non-PAP findings as exploratory or suggestive.

³If individuals are fully attentive, their responses to incentives should be the same whether the incentives are displayed in a high- or low-salience way. A larger gap between the two conditions therefore signals inattention. As sleep is hypothesized to improve alertness, if the treatment reduces this gap, that provides evidence that sleep mitigates inattention to incentives.

⁴We assume that PAP results are always reported (i.e., the researcher does not *hide* any findings). For simplicity, we also do not model degrees of emphasis in reporting. In our framework, placing PAP results in an appendix is equivalent to including them in the main text.

We refer to the reporting of a non-preregistered result as a *deviation* from the PAP. Uncertainty over their interpretation notwithstanding, deviations are extremely common in practice and arise for many reasons: understanding surprising results (exploration or mechanism probes), referee and editor requests (e.g., robustness or placebo checks), or the availability of improved statistical tools since PAP submission. We use [Bessone et al. \(2021\)](#) as a running example of such deviations.

Example 1. [Bessone et al. \(2021\)](#) study the economic effects of different policies targeted at increasing sleep among the urban poor. The authors cross-randomize workers in Chennai, India, to receive (i) one of two treatments designed to improve nighttime sleep and (ii) a treatment that provides nap breaks during workdays. [Bessone et al. \(2021\)](#) preregister a set of specifications regressing indices of work, well-being, and cognition outcomes on interactions of the night sleep and nap treatments. However, as the authors describe in their paper,

“Those who received a night sleep treatment in addition to naps had very similar effects to those with naps only. This pattern of results suggests that naps have an overall positive effect on outcomes, whereas increases in night sleep do not. To increase statistical power and streamline the discussion of effects on the individual outcomes, we turn to an analysis that pools the two night sleep treatments and does not allow for an interaction effect of night sleep and nap treatments.”

This pooling decision is substantively motivated: When two night-sleep arms are practically indistinguishable, combining them yields a more precisely estimated quantity. However, at this stage there is uncertainty in the profession over how to proceed: A blanket ban on deviations would preclude learning from the pooled results, yet reporting conventional (unconditional) point estimators and confidence intervals would ignore that the pooling decision was made selectively after seeing the data, risking an overstatement of precision. While the researcher can flag such results as exploratory and meant to be taken with a grain of salt, it is still ambiguous how to interpret deviations from the PAP. In the coming sections, we express the stages outlined above with a formal model, then characterize the adjustments that should be made to account for the selective reporting of non-preregistered results.

2 Model

We begin by describing the environment and introducing terms in [Section 2.1](#), formalizing the researcher’s timeline discussed in the previous section. To ground motivation, in [Section 2.2](#) we contextualize the researcher’s choice of specifications in terms of a Bayesian decision problem and show that even a fully Bayesian decision maker may wish to deviate if constrained

-
1. Register PAP \mathcal{S}_{pre} consisting of specifications S_{pre}
 2. Observe data X
 3. Announce deviations \mathcal{S}_{post}^X consisting of non-preregistered specifications S_{post}
 4. Report PAP estimates $\hat{\beta}_{pre} = S_{pre}(X)$ and post estimates $\hat{\beta}_{post} = S_{post}(X)$
-

Figure 2: Timeline of Researcher’s Process

to choose from a limited set of specifications. We emphasize, however, that our subsequent inferential framework in Section 3 is valid without this Bayesian structure.

2.1 Environment

The researcher observes data $X \in \mathcal{X}$, for \mathcal{X} a sample space. The researcher considers specifications $S: X \rightarrow \mathbb{R}^d$ from some set \mathcal{S} ; that is, each specification $S \in \mathcal{S}$ maps data X to estimates $S(X)$. An example of \mathcal{S} could be the set of all linear regressions.⁵ In advance of observing the data, the researcher preregisters a set of specifications $\mathcal{S}_{pre} \subseteq \mathcal{S}$, where \mathcal{S}_{pre} represents the pre-analysis plan (PAP). We denote specifications from the PAP by $S_{pre} \in \mathcal{S}_{pre}$. For example, \mathcal{S}_{pre} might be the specific set of linear regressions the researcher intends to run, and S_{pre} one such regression.⁶ After observing X , the researcher announces a set of non-preregistered specifications $\mathcal{S}_{post}^X \subseteq \mathcal{S} \setminus \mathcal{S}_{pre}$. We refer to \mathcal{S}_{post}^X as *deviations from the PAP*, which may or may not depend on the observed data X . The researcher then reports (i) PAP estimates $\hat{\beta}_{pre} = S_{pre}(X)$ for all PAP specifications $S_{pre} \in \mathcal{S}_{pre}$ and (ii) post estimates $\hat{\beta}_{post} = S_{post}(X)$ for all deviations $S_{post} \in \mathcal{S}_{post}^X$. We depict this timeline in Figure 2. To ground this notation, we revisit the example of Bessone et al. (2021).

Example 1 (Continued). Bessone et al. (2021) cross-randomize workers to receive (i) one of two treatments designed to improve nighttime sleep (“Incentives” and “Encouragement”) and (ii) a treatment that provides nap breaks during workdays (“Nap”). They preregister a set of regression specifications \mathcal{S}_{pre} regressing indices of work, well-being, and cognition outcomes on interactions of the two night sleep and nap treatments.

After conducting the study and observing data X , the authors compute the treatment effects of the five fully disaggregated treatment arms on preregistered outcomes, $\hat{\beta}_{pre} = S_{pre}(X)$.

⁵This setup accommodates cases where the researcher considers regression specifications with differing numbers of regressors, since vectors in $\mathbb{R}^{\tilde{d}}$ for $\tilde{d} < d$ can also be represented as vectors in \mathbb{R}^d .

⁶Observe that one could alternately define \mathcal{S} to be the set of all possible PAPs, from which the researcher chooses a single element $S_{pre} \in \mathcal{S}$. Our choice of notation, while mathematically equivalent, is more practical for when we discuss inference in Section 3, and seems better aligned with how researchers view PAPs.

However, after observing that “those who received a night sleep treatment in addition to naps had very similar effects to those with naps only,” the authors reason that “naps have an overall positive effect on outcomes, whereas increases in night sleep do not.” To increase statistical power in discussion of effects on outcomes, the authors estimate an additional set of specifications $S_{post} \in \mathcal{S}_{post}^X$ which (i) pools the two night sleep treatments and (ii) removes the interaction term for night sleep and nap treatments, yielding $\hat{\beta}_{post} = S_{post}(X)$. The authors supplement their report of the PAP estimates $\hat{\beta}_{pre}$ (disaggregated) with post estimates $\hat{\beta}_{post}$ (pooled), found in Tables III and IV of their paper, respectively.

2.2 Decision Problem and Deviations

There are many reasons that deviations arise. In this section, we show that even a fully Bayesian researcher in this framework will sometimes want to deviate when they are limited to a constrained set of specifications for constructing their PAP.⁷ We offer discussion on where such constraints arise in practice, such as through labor and communication costs, tractability, cognitive limitations, and norms in the profession.

The distribution of X is governed by the parameter $\theta \in \Theta$, $X \sim P_\theta \in \Delta(\mathcal{X})$, for Θ a parameter space and $\Delta(\mathcal{X})$ the set of probability distributions on \mathcal{X} . The researcher has a loss function $L(S(X), \theta) \geq 0$ that quantifies the consequences of reporting estimates $S(X) \in \mathbb{R}^d$ when the true parameter is $\theta \in \Theta$. For example, if the goal is to estimate a treatment effect $\tau(\theta) \in \mathbb{R}$, then a standard loss function is squared error:

$$L(S(X), \theta) = (S(X) - \tau(\theta))^2. \quad (1)$$

The *risk* of specification S is the expected loss from reporting $S(x)$ across possible realizations $x \in \mathcal{X}$ of the data:

$$R(S, \theta) = \int_{\mathcal{X}} L(S(x), \theta) dP_\theta(x).$$

Note that the risk depends on the unknown parameter $\theta \in \Theta$.

PAP Construction. We consider the researcher’s problem of choosing the set of specifications $\mathcal{S}_{pre}, \mathcal{S}_{post}^X \subseteq \mathcal{S}$ to estimate, focusing on the case where $|\mathcal{S}_{post}^X| \leq |\mathcal{S}_{pre}| = 1$ for simplicity. We suppose the researcher has prior beliefs about the unknown parameter θ , represented by a prior distribution $\pi \in \Delta(\Theta)$. The researcher determines which specification S_{pre} to preregister in their PAP by minimizing the average risk of S under their prior π . Stated formally, given a

⁷This model is for motivation alone: The inference procedures in Section 3 are valid without this Bayesian structure.

set of specifications \mathcal{S} , loss function L , and prior π , the researcher solves:

$$S_{pre} \in \operatorname{argmin}_{S \in \mathcal{S}} \bar{R}(S, \pi), \quad \bar{R}(S, \pi) = \int_{\Theta} R(S, \theta) d\pi(\theta), \quad (2)$$

where $\bar{R}(S, \pi)$ is the average risk of S under π . After solving this optimization problem, the researcher preregisters S_{pre} in their PAP.

Deviations from PAP. The researcher observes X and reports PAP estimates $\hat{\beta}_{pre} = S_{pre}(X)$, as planned. However, the data also lead the researcher to update their prior beliefs π to form posterior beliefs, represented by the posterior distribution $\pi_X \in \Delta(\Theta)$ of θ conditional on X . Having gleaned new information about θ , the researcher may now find their preregistered S_{pre} to be suboptimal, leading the researcher to also report post estimates $\hat{\beta}_{post} = S_{post}(X)$, where $S_{post} \neq S_{pre}$. We formalize this decision to deviate below.

The analogous problem to minimizing prior average risk (2) is to minimize posterior average loss:

$$S_{post} \in \operatorname{argmin}_{S \in \mathcal{S}} \bar{L}(S(X), \pi_X), \quad \bar{L}(S(X), \pi_X) = \int_{\Theta} L(S(X), \theta) d\pi_X(\theta), \quad (3)$$

where $\bar{L}(S(X), \pi_X)$ is the average loss of $S(X)$ under π_X . The researcher deviates from their PAP when $S_{post} \neq S_{pre}$. This decision to deviate depends on the coarseness of \mathcal{S} .

In a conventional Bayesian decision problem, \mathcal{S} is taken to be the set of all functions from \mathcal{X} to \mathbb{R}^d . In this case, a specification that minimizes prior average risk (2) also minimizes posterior average loss (3) almost surely, and vice versa (Lehmann and Casella, 2006, Theorem 1.1). Intuitively, when \mathcal{S} is unrestricted, a solution to (2) perfectly prepares for every possible data realization $x \in \mathcal{X}$. In such cases, there is never a need to deviate.⁸

We depart from the conventional Bayesian setup by allowing for a restricted \mathcal{S} . That \mathcal{S} is constrained is not a modeling convenience, but a description of practice: At the preregistration stage, researchers face explicit and implicit constraints that narrow the permissible set of specifications. In general, researchers cannot write PAPs sophisticated and lengthy enough to prepare for every $x \in \mathcal{X}$. For one, researchers face *time*, *cognitive*, and *communication* constraints: For most experimental designs, it is not feasible to anticipate, formulate, and communicate optimal responses to every contingency (Olken, 2015; Banerjee et al., 2020a). As an example, the PAP for Finkelstein et al. (2012)—a 50 page paper—is 159 pages long. And, despite being exceptionally detailed, the initial PAP still fell short of exhausting all relevant analyses: From 2012 to 2020,

⁸In Appendix A, we provide a variant of the decision problem where, instead of considering posterior average loss, the researcher solves (2) with π_X in the place of π . In this setup, S_{post} is the *counterfactual PAP* that the researcher would have constructed had they entered the experiment with their posterior beliefs. If one takes \mathcal{S} to be the relevant action space, this counterfactual PAP approach aligns with the conditional Bayes principle, which says to choose an action that minimizes average loss under one’s current beliefs (Berger, 2013, Section 1.5.1). Interestingly, deviations arise in this setup even when \mathcal{S} is unconstrained.

ten additional PAPs (totaling an additional 368 pages) were pre-registered to explore new outcome measures in follow-up papers after collecting additional data (Baicker et al., 2025).

Norms in the profession also constrain \mathcal{S} , favoring parsimonious, interpretable specifications tied to the design and estimand. By default, this includes conventional functional forms and links (e.g., linear or low-order polynomials, logit/probit for binary outcomes, etc.), standard control sets motivated by identification, and prespecified clustering and variance estimators. Idiosyncratic specifications—such as cube-root transformations, high-order polynomials, post-hoc cutpoints, data-driven bins or knots, bespoke weighting schemes, stepwise selection, or open-ended machine learning screens—are typically viewed with skepticism unless accompanied by a compelling ex-ante justification grounded in theory, measurement scale, or experimental design (and appropriately powered).

The constraints on \mathcal{S} lead to a classic dynamic inconsistency problem (Strotz, 1955; Kydland and Prescott, 1977), wherein a plan that is optimal ex ante under prior beliefs π may no longer be optimal after updating to the posterior π_X .⁹ In a standard Bayesian decision problem with a rich \mathcal{S} , the researcher would be allowed to preregister a contingent rule (e.g., “after X is observed, take whichever action meets criterion C ”), so the plan chosen ex ante already pins down the ex-post action. When \mathcal{S} excludes such contingent rules, however, this flexibility is lost. Concretely, if a researcher is restricted to naming one primary outcome in advance (employment or consumption), a contingent rule like “highlight whichever outcome best meets our welfare/precision criterion after we see X ” is not in \mathcal{S} . Ex ante the researcher may pick employment under π , but ex post the data may favor consumption. Because \mathcal{S} was insufficiently rich to accommodate switching, the ex ante choice remains binding even when the posterior π_X would prefer the other outcome—hence the dynamic inconsistency.

3 Inference

This section proposes conditional inference procedures for post estimates. We begin in Section 3.1 by introducing the problem of statistical inference in a model with normally distributed estimates. With the relevant notation and concepts established, we next motivate in Section 3.2 the use of conditional inference for post estimates and provide examples in Section 3.3. We then develop the conditional inference procedures in Sections 3.4–3.6 and provide an explicit example in Section 3.7. The normality assumption is motivated by standard large sample approximations, such as the central limit theorem and the law of large numbers. In Section 4, we describe how to implement analogues of the proposed conditional inference procedures

⁹In behavioral terms, present-bias and self-control generate the same logic (Laibson, 1997; O’Donoghue and Rabin, 1999; Frederick et al., 2002). Related commitment perspectives appear in rules-vs-discretion and menu/temptation models (Barro and Gordon, 1983; Schelling, 1960; Kreps, 1979; Gul and Pesendorfer, 2001).

using sample estimates and establish their uniform asymptotic validity.

3.1 Inference Setup

Let $\hat{\beta}$ be a vector that contains (i) the PAP estimates $\hat{\beta}_{pre} = S_{pre}(X)$ for each PAP specification $S_{pre} \in \mathcal{S}_{pre}$ and (ii) the post estimates $\hat{\beta}_{post} = S_{post}(X)$ for each deviation $S_{post} \in \mathcal{S}_{post}^X$ made under X . We assume that $\hat{\beta}$ is normally distributed with unknown mean β and known positive definite covariance matrix Σ :

$$\hat{\beta} \sim N(\beta, \Sigma). \quad (4)$$

Let $\mathbb{P}_\beta\{\cdot\}$ and $\mathbb{E}_\beta[\cdot]$ denote probabilities and expectations under this distribution.¹⁰ Each $\hat{\beta}_{pre}$ corresponds to a PAP estimand $\beta_{pre} = \mathbb{E}_\beta[\hat{\beta}_{pre}]$, while each $\hat{\beta}_{post}$ corresponds to a post estimand $\beta_{post} = \mathbb{E}_\beta[\hat{\beta}_{post}]$. By construction, β contains all the PAP and post estimands. We consider the problem of conducting inference on parameters that can be expressed as linear combinations of these estimands: $v'\beta_{pre}$ and $l'\beta_{post}$, where $v, l \neq 0$.¹¹ For instance, in Example 1, the effects of the various treatments correspond to different contrasts of the regression coefficients.

Inference for PAP Estimates. Given PAP estimates $\hat{\beta}_{pre}$, let Σ_{pre} denote the covariance matrix induced by Σ . By equation (4), any linear combination of PAP estimates is normally distributed:

$$v'\hat{\beta}_{pre} \sim N(v'\beta_{pre}, \sigma_{pre}^2), \quad \sigma_{pre}^2 = v'\Sigma_{pre}v.$$

Given a desired significance level $\alpha \in (0, 1)$, a confidence interval $CI_\alpha(X)$ has correct coverage for $v'\beta_{pre}$ across data realizations $X \in \mathcal{X}$ if it satisfies

$$\mathbb{P}_\beta\{v'\beta_{pre} \in CI_\alpha(X)\} = 1 - \alpha, \quad \forall \beta. \quad (5)$$

In words, a confidence interval $CI_\alpha(X)$ that satisfies criteria (5) contains the parameter of interest $v'\beta_{pre}$ with probability $1 - \alpha$ for any possible value of the unknown β . This is a standard criteria for valid inference (Lehmann and Romano, 2024). One example is the conventional two-sided interval $[v'\hat{\beta}_{pre} \pm z_{1-\alpha/2}\sigma_{pre}]$, where z_α denotes the α -quantile of the $N(0, 1)$ distribution. This interval satisfies

$$\mathbb{P}_\beta\{v'\beta_{pre} \in [v'\hat{\beta}_{pre} \pm z_{1-\alpha/2}\sigma_{pre}]\} = 1 - \alpha, \quad \forall \beta,$$

¹⁰We implicitly assume that estimators and confidence intervals of interest depend on X through $\hat{\beta}$.

¹¹This setup is quite general, since differentiable nonlinear functions of asymptotically normal estimates will also be asymptotically normal by the delta method. These nonlinear functions can be elements of $\hat{\beta}$.

and therefore yields valid inference for PAP estimates.

Inference for Post Estimates. Unlike the preregistered specifications \mathcal{S}_{pre} , the researcher’s deviations \mathcal{S}_{post}^X may depend on the observed data X . We say that a data realization X *induces* the deviation S_{post} if $S_{post} \in \mathcal{S}_{post}^X$ in that data realization. For a given S_{post} , we may then collect the set of all data realizations that induce the deviation:

$$\mathcal{X}_{post} = \{X \in \mathcal{X} : S_{post} \in \mathcal{S}_{post}^X\}.$$

We refer to \mathcal{X}_{post} as the *deviation set* for S_{post} . By definition, the researcher is only interested in post estimates $\hat{\beta}_{post}$ when $X \in \mathcal{X}_{post}$. In view of this, we propose confidence intervals that satisfy the following criteria: Given significance level $\alpha \in (0,1)$, a confidence interval $CI_\alpha(X)$ has correct *conditional coverage* for $l'\beta_{post}$ given $X \in \mathcal{X}_{post}$ if it satisfies

$$\mathbb{P}_\beta\{l'\beta_{post} \in CI_\alpha(X) \mid X \in \mathcal{X}_{post}\} = 1 - \alpha, \quad \forall \beta. \quad (6)$$

Intuitively, such intervals ensure correct coverage for $l'\beta_{post}$ across the set of data realizations $\mathcal{X}_{post} \subseteq \mathcal{X}$ where the researcher is interested in $\hat{\beta}_{post}$. This conditional coverage criteria generalizes the *unconditional* criteria in (5). Indeed, since PAP estimates are always reported, the analogous “deviation set” for a PAP estimate yields $\{X \in \mathcal{X} : S_{pre} \in \mathcal{S}_{pre}\} = \mathcal{X}$, in which case conditional and unconditional coverage coincide.

Known Deviation Set. To obtain $CI_\alpha(X)$ with correct conditional coverage, we assume the researcher can (and does) articulate the deviation sets \mathcal{X}_{post} for each $S_{post} \in \mathcal{S}_{post}^X$ at hand. This is a natural first step for deriving conditional inference procedures. We relax this assumption in Section 6, where we discuss robustness of conditional inference procedures to various forms of misspecification. To make the above concepts concrete, we return to the example deviation in Bessone et al. (2021). In the following example, let $se(\cdot)$ denote the standard deviation of an estimator under normal distribution (4).

Example 1 (Continued). Bessone et al. (2021) initially preregistered the fully interacted specification $\hat{\beta}_{pre} = (\hat{\tau}_N, \hat{\tau}_{NE}, \hat{\tau}_{NI}, \hat{\tau}_E, \hat{\tau}_I)'$, where

- $\hat{\tau}_N$: effect of naps only
- $\hat{\tau}_{NE}$: effect of naps \times encouragement
- $\hat{\tau}_{NI}$: effect of naps \times incentives
- $\hat{\tau}_E$: effect of encouragement only
- $\hat{\tau}_I$: effect of incentives only

After seeing the data and estimating $\hat{\beta}_{pre}$, the authors report three additional post estimates, $\hat{\beta}_{post} = (\hat{\tau}_{E+I}, \hat{\tau}_{NE+NI}, \hat{\tau}_{N+NE+NI})'$, corresponding to the estimates which (i) pool incentives and encouragement into a single “night sleep” treatment ($\hat{\tau}_{E+I}$ and $\hat{\tau}_{NE+NI}$) and (ii) collapse interaction effects ($\hat{\tau}_{N+NE+NI}$).¹² The authors find that, “these results ($\hat{\beta}_{pre}$) provide evidence that naps have an overall positive effect on outcomes, while increases in night sleep do not. However... this analysis has limited statistical power.” The authors justify turning to a “simplified but higher-powered version of this analysis” with the following three reasons:

(A) “Each of the night sleep treatments alone had no significant effect on this overall index,” i.e.,

$$\frac{|\hat{\tau}_E|}{\text{se}(\hat{\tau}_E)} \leq z_{1-\eta/2} \text{ and } \frac{|\hat{\tau}_I|}{\text{se}(\hat{\tau}_I)} \leq z_{1-\eta/2}.$$

(B) “In contrast, naps alone had a positive, marginally significant effect,” i.e.,

$$\frac{|\hat{\tau}_N|}{\text{se}(\hat{\tau}_N)} \geq z_{1-\eta/2}.$$

(C) “Those who received a night sleep treatment in addition to naps had very similar effects to those with naps only,” i.e.,

$$\frac{|\hat{\tau}_N - \hat{\tau}_{NE}|}{\text{se}(\hat{\tau}_N - \hat{\tau}_{NE})} \leq z_{1-\eta/2} \text{ and } \frac{|\hat{\tau}_N - \hat{\tau}_{NI}|}{\text{se}(\hat{\tau}_N - \hat{\tau}_{NI})} \leq z_{1-\eta/2}.$$

Taking these points at face value, a plausible characterization of the authors’ deviation set for $\hat{\beta}_{post} = (\hat{\tau}_{E+I}, \hat{\tau}_{NE+NI}, \hat{\tau}_{N+NE+NI})'$ is the set of data realizations satisfying the intersection of these events, i.e.,

$$\mathcal{X}_{post} = \{X \in \mathcal{X} : (\text{A}) \text{ and } (\text{B}) \text{ and } (\text{C})\}.$$

3.2 Motivating Conditional Inference

We argue that *conditional* coverage of $l'\beta_{post}$ should be the preferred inferential objective—not unconditional coverage, as is convention—since post estimates are only reported or acted upon in select states of the world. We see two main reasons for this. First, selective reporting can lead to distortions in the density of observed reports of $\hat{\beta}_{post}$, as we will demonstrate. Second, conditional validity is, in a certain sense, necessary and sufficient for confidence sets to be valid on average across studies, which is the very property underlying existing frequentist coverage objectives.

¹²The estimate $\hat{\tau}_{NE+NI}$ is not reported in the main paper, but rather in Online Appendix Table A.VIII (where the authors pool the two night sleep treatments but include a separate indicator for individuals who received a combination of either night sleep treatment along with the nap treatment).




Figure 3: Unconditional Density of $(\hat{\beta}_{pre}, \hat{\beta}_{post})$ vs. Conditional $(\hat{\beta}_{pre}, \hat{\beta}_{post}) | \hat{\beta}_{pre} \geq 0$.

$X = (\hat{\beta}_{pre}, \hat{\beta}_{post}) \sim N(0, \Sigma)$, $\sigma_{pre}^2 = \sigma_{post}^2 = 1$, $\rho = 0.5$. *Left:* Unconditional joint density of $(\hat{\beta}_{pre}, \hat{\beta}_{post})$. *Center:* Joint density of $(\hat{\beta}_{pre}, \hat{\beta}_{post})$ conditional on $\hat{\beta}_{pre} \geq 0$. *Right:* Marginal density of $\hat{\beta}_{post}$ under selective reporting (blue) compared to unconditional density (red). Conventional 5% significance cutoffs shown with red dashed lines.

Conditional Coverage for Given Study. Selective reporting can distort the density of observed reports $\hat{\beta}_{post}$ relative to the unconditional density assumed by conventional inference procedures. This is demonstrated in the following toy example.

Example 2. Let

$$\begin{pmatrix} \hat{\beta}_{pre} \\ \hat{\beta}_{post} \end{pmatrix} \sim N(0, \Sigma), \quad \Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix},$$

where $\rho = 0.5$ is the correlation between $\hat{\beta}_{pre}$ and $\hat{\beta}_{post}$. The leftmost panel in Figure 3 shows the joint unconditional density of $(\hat{\beta}_{pre}, \hat{\beta}_{post})$, i.e., when $\mathcal{X}_{post} = \mathbb{R}^2$. The center panel shows the joint density of $(\hat{\beta}_{pre}, \hat{\beta}_{post})$ conditional on $\hat{\beta}_{pre} \geq 0$, i.e., when $\hat{\beta}_{post}$ is selectively reported based on $\mathcal{X}_{post} = \{X \in \mathbb{R}^2 : \hat{\beta}_{pre} \geq 0\}$. The rightmost panel superimposes the associated marginal densities of $\hat{\beta}_{post}$ under selective (blue) and unconditional (red) reporting regimes. To give a reference of magnitude, the red dashed lines show the conventional 5% significance test cutoffs. As can be seen, the missing mass in the post-selection joint density causes the density of *observed reports* of $\hat{\beta}_{post}$ to look quite different from the unconditional density of $\hat{\beta}_{post}$ assumed by conventional inference procedures. By ignoring the correlation between the conditioning event and estimate $\hat{\beta}_{post}$, conventional procedures yield inflated type I error, undercoverage, and miscalibrated conclusions.

Average Coverage Across Studies. The conditional coverage criterion requires that $CI_\alpha(X)$ contain $l'\beta_{post}$ with high probability across data realizations $X \in \mathcal{X}_{post}$, where each data realization can be viewed as a particular draw of the same experiment in a given study. Thus,

conditional coverage provides a natural inference criterion for a *given* study. However, one might ponder the relevance of conditional coverage when there are *multiple* studies, each with their own experiments and potential PAP deviations. In Appendix B, we establish a general sense in which conditional coverage for *each* study is necessary and sufficient for controlling average coverage across *all* studies.

Formally, we consider a sampling model where deviation sets are drawn from a rich class of distributions reflecting the unanticipated nature of PAP deviations. Under this structure, we show that to ensure average coverage across studies, it is both *necessary and sufficient* to ensure conditional coverage in each study. This equivalence result provides further motivation for the conditional inference approach to PAP deviations.

3.3 Types of Deviation Sets

We now discuss leading examples of deviation sets. Section 3.3.1 outlines deviations based on the significance of PAP estimates, which we discuss extensively in future sections. Section 3.3.2 highlights classes of deviation sets that yield no inference distortions in our framework.

3.3.1 Deviations Based on Significance

Researchers often consider post estimates $\hat{\beta}_{post}$ for which reasoning about \mathcal{X}_{post} is predicated on values of PAP estimates $\hat{\beta}_{pre}$. In particular, $\hat{\beta}_{post}$ is often of interest when linear combinations of PAP estimates $v'\hat{\beta}_{pre}$ cross significance cutoffs $\kappa \geq 0$. Here we focus on $\kappa = z_{1-\eta/2}\sigma_{pre}$, where $\sigma_{pre} = \sqrt{v'\Sigma_{pre}v}$, which corresponds to a conventional two-sided *statistical* significance test with significance level $\eta \in (0,1)$. However, κ can also be based on economic significance.

Statistical Significance. The researcher deviates when $v'\hat{\beta}_{pre}$ is significantly large:

$$\mathcal{X}_{post} = \left\{ X \in \mathcal{X} : |v'\hat{\beta}_{pre}| \geq z_{1-\eta/2}\sigma_{pre} \right\}.$$

For example, consider Dube et al. (2025), who study the effects of a cognitive-skills training program (Sit-D) for police officers on discretionary arrests. In their Section IV.B, they say

“In this section we check the robustness of our main results to alternate specifications and variable definitions. Our discretionary arrests category constitutes a small subset of arrests . . . that we thought were most likely to fall in response to the training. To gauge if Sit-D also reduces other types of arrests over which officers might have discretion, we use a different measure of discretionary arrests . . . which consists of arrests for “nonindex” crimes under the FBI’s classification. We did not specify analyzing this outcome in our PAP but examine it for robustness.”

One can interpret the above quote as suggesting a deviation set based on statistical significance: $\hat{\beta}_{pre}$ is the vector of coefficients on the regression with the preregistered measure, v is the unit vector that picks the coefficient on SitD, and $\hat{\beta}_{post}$ is the vector of coefficients for the regression with the non-preregistered measure, with $v'\hat{\beta}_{pre}$ significant at level $\eta=0.05$ (Dube et al., 2025, Table IV).

Statistical Insignificance. The researcher deviates when $v'\hat{\beta}_{pre}$ is significantly small:

$$\mathcal{X}_{post} = \left\{ X \in \mathcal{X} : |v'\hat{\beta}_{pre}| \leq z_{1-\eta/2}\sigma_{pre} \right\}.$$

For example, in the quote from Bessone et al. (2021) presented in Example 1, two of the reasons the authors give for pooling, (A) and (B), are based on statistical insignificance. We consider various representations of \mathcal{X}_{post} for Bessone et al. (2021) in the empirical application in Section 5.

3.3.2 Non-Distortionary Deviations

An important and immediate implication of our framework is that when the deviation event $\{X \in \mathcal{X}_{post}\}$ is *independent* of the post estimates $\hat{\beta}_{post}$, one does *not* have to correct conventional inferences. In particular, since the conventional interval $[l'\hat{\beta}_{post} \pm z_{1-\alpha/2}\sigma_{post}]$ is a function of $\hat{\beta}_{post}$, independence yields

$$\mathbb{P}_{\beta} \left\{ l'\beta_{post} \in [l'\hat{\beta}_{post} \pm z_{1-\alpha/2}\sigma_{post}] | X \in \mathcal{X}_{post} \right\} = \mathbb{P}_{\beta} \left\{ l'\beta_{post} \in [l'\hat{\beta}_{post} \pm z_{1-\alpha/2}\sigma_{post}] \right\} = 1 - \alpha, \quad \forall \beta.$$

Such *non-distortionary* deviations can broadly be grouped into two cases, depicted in Figure 4 and discussed below.

Unconditional Reporting. The first case (left panel of Figure 4) is *unconditional reporting*, where even though $\hat{\beta}_{post}$ was not registered in the PAP, it would have been reported *no matter the data observed*. In other words, $\mathcal{X}_{post} = \mathcal{X}$. While seemingly a trivial result, there are nonetheless many examples of such deviations that make researchers uneasy in the status quo. Instances of $\mathcal{X}_{post} = \mathcal{X}$ might include: (i) honest mistakes, oversights in PAP specification, or adjustments for feasibility; (ii) unanticipated events (e.g., change in policy context, global pandemic); or (iii) newly acquired knowledge about the setting/context, econometric methods, or theoretical insights in the literature.¹³ For concreteness, we now consider an explicit example of (iii).

¹³Examples of (i) include Rafkin et al. (2021, Appendix D.5); Bhat et al. (2022, pages 9-10); Alsan et al. (2024, footnote 24); Jacobson and Powell (2024, footnotes 18 and 24); Agte et al. (2022, Table A5); Finkelstein and Notowidigdo (2019, footnote 9); Kaur et al. (2024, footnote 9); Kelley et al. (2024, page 3140). Examples of (ii) include Kremer et al. (2009, Section III.B); Evsyukova et al. (2025, footnote 7). Examples of (iii) include Bandiera et al. (2021, footnote 26); Field et al. (2021, page 2345); Bessone et al. (2021, pages 1914-15); Giacobino et al. (2024, footnote 18).




Figure 4: Non-Distortionary Deviations

$X = (\hat{\beta}_{pre}, \hat{\beta}_{post}) \sim N(0, \Sigma)$, $\sigma_{pre}^2 = \sigma_{post}^2 = 1$, $\rho \in \{0.5, 0\}$. *Left*: $\hat{\beta}_{pre}$ and $\hat{\beta}_{post}$ are correlated ($\rho = 0.5$), but $\hat{\beta}_{post}$ is reported unconditionally ($\mathcal{X}_{post} = \mathbb{R}^2$). *Center*: $\hat{\beta}_{post}$ is reported conditional on $\hat{\beta}_{pre} \geq 0$ ($\mathcal{X}_{post} = \{X \in \mathbb{R}^2 : \hat{\beta}_{pre} \geq 0\}$), but $\hat{\beta}_{post}$ and $\hat{\beta}_{pre}$ are uncorrelated ($\rho = 0$). *Right*: Overlapped marginal density of $\hat{\beta}_{post}$ in either case. Conventional 5% significance cutoffs shown with red dashed lines.

Example 3. [Bandiera et al. \(2021\)](#) study how the allocation of authority between procurement offices (POs) and their supervising accountant generals (AGs) affects performance, as measured by the price spent on procured goods. The primary result—based on PAP estimates—is that greater PO autonomy reduces prices. The secondary result—based on post estimates—is that the benefits of autonomy are higher when AG incentives are misaligned, as measured by the extent to which the AG delays approvals for PO purchases. The exact language surrounding the deviation is as follows.

“In our preanalysis plan we did not prespecify that we would study heterogeneity by the AG’s type. As the experiment rolled out and we discussed its effects with our study participants we came to realize the importance of the type of the AG in determining how the treatments, particularly the autonomy treatment, affected the way POs were able to change the way they carried out procurement.”

This above language suggests a deviation set equal to the entire sample space: $\mathcal{X}_{post} = \mathcal{X}$.

Independent Estimates. The second case (center panel of Figure 4) is conditional reporting of $\hat{\beta}_{post}$ based on estimates $\hat{\beta}_{pre}$ that are *independent* of $\hat{\beta}_{post}$. There are mechanical ways the researcher can ensure this: A key example is pilot studies where researchers (i) register a PAP and conduct a pilot of their study on one population and (ii) make adjustments based on results from the pilot before conducting the main study on a larger population.¹⁴ When the population

¹⁴An example that broadly falls into this class is [Dean \(2024, Appendix D\)](#).

of the pilot is sampled independently from that of the main study, any sort of deviation from the PAP based on findings in the pilot yields no inference distortions. Another example in this vein is sample splitting, where one computes estimates on one half of the data, uses those results to determine their deviations, and computes the corresponding $\hat{\beta}_{post}$ using the other half of the data—however, such procedures come at the cost of statistical precision.

Of course, independence need not be a mechanical feature of the experimental design. Consider an intervention disrupted by a lightning strike which changes the underlying sample, as in [Kremer et al. \(2009\)](#).¹⁵ This changes the estimand’s interpretation, as the conditioning event may index a different state/population s . However, conventional inference will be valid for this estimand. That is, we now have $\hat{\beta}|\{S=s\} \stackrel{d}{=} \hat{\beta}(s)|\{S=s\} \stackrel{d}{=} \hat{\beta}(s) \sim N(\beta(s), \Sigma(s))$, where the second equality follows from independence of selection variable S (e.g., random lightning strike) and estimates $\hat{\beta}(s)$.

3.4 Quantile Conditionally Unbiased Estimation

We now present general conditional inference procedures. In Section 3.4.1, we state the formal inference objectives for confidence intervals and point estimators. In Sections 3.4.2–3.4.3, we detail the construction of optimal procedures that meet these objectives.

3.4.1 Objective

To perform valid conditional inference, it suffices to derive estimators $\hat{\mu}_\alpha$ that are α -quantile conditionally unbiased in the sense that their overestimation probability for $l'\beta_{post}$ conditional on $X \in \mathcal{X}_{post}$ is equal to α :

$$\mathbb{P}_\beta\{l'\beta_{post} \leq \hat{\mu}_\alpha \mid X \in \mathcal{X}_{post}\} = \alpha, \quad \forall \beta. \quad (7)$$

Given such $\hat{\mu}_\alpha$, the confidence interval $CI_\alpha(X) = [\hat{\mu}_{\alpha/2}, \hat{\mu}_{1-\alpha/2}]$ provides correct conditional coverage:

$$\mathbb{P}_\beta\{\hat{\mu}_{\alpha/2} \leq l'\beta_{post} \leq \hat{\mu}_{1-\alpha/2} \mid X \in \mathcal{X}_{post}\} = 1 - \alpha, \quad \forall \beta. \quad (8)$$

Alternatively, if we want a point estimator of $l'\beta_{post}$, we can use $\hat{\mu}_{1/2}$ as a median conditionally unbiased estimator:

$$\mathbb{P}_\beta\{l'\beta_{post} \leq \hat{\mu}_{1/2} \mid X \in \mathcal{X}_{post}\} = \frac{1}{2}, \quad \forall \beta. \quad (9)$$

In words, the conditional median of $\hat{\mu}_{1/2}$ is equal to $l'\beta_{post}$.

¹⁵See also [Olken \(2015, footnote 2\)](#) for a discussion.

3.4.2 Derivation

To derive $\hat{\mu}_\alpha$, we follow arguments from [Fithian et al. \(2014\)](#). Assume there exists a set \mathcal{B}_{post} of positive measure such that

$$\mathcal{X}_{post} = \{X \in \mathcal{X} : \hat{\beta} \in \mathcal{B}_{post}\}, \quad (10)$$

where the set \mathcal{B}_{post} is implicitly allowed to depend on Σ . In other words, we focus on deviations that arise due to values of $\hat{\beta}$. This broadly accommodates the deviations that occur in practice, such as those based solely on PAP estimates $\hat{\beta}_{pre}$. For deviation sets satisfying (10), we have

$$\hat{\beta}_{post} | X \in \mathcal{X}_{post} \stackrel{d}{=} \hat{\beta}_{post} | \hat{\beta} \in \mathcal{B}_{post}.$$

Thus, our goal is to obtain valid inference for $l'\beta_{post}$ conditional on $\hat{\beta} \in \mathcal{B}_{post}$.

Notation. Let l_{post} denote the vector induced by (i) matrix multiplication to select $\hat{\beta}_{post}$ from $\hat{\beta}$ and (ii) vector multiplication of $\hat{\beta}_{post}$ by l , so that $l'\hat{\beta}_{post} = l'_{post}\hat{\beta}$. Let Σ_{post} denote the covariance matrix for $\hat{\beta}_{post}$, so that the variance of $l'\hat{\beta}_{post}$ is $\sigma_{post}^2 = l'_{post}\Sigma l_{post} = l'\Sigma_{post}l$.

The main challenge is that $l'\hat{\beta}_{post}$ is no longer normally distributed after conditioning on $\hat{\beta} \in \mathcal{B}_{post}$, owing to correlation between $l'\hat{\beta}_{post}$ and $\hat{\beta}$:

$$\begin{pmatrix} l'\hat{\beta}_{post} \\ \hat{\beta} \end{pmatrix} \sim N\left(\begin{pmatrix} l'\beta_{post} \\ \beta \end{pmatrix}, \begin{pmatrix} \sigma_{post}^2 & l'_{post}\Sigma \\ \Sigma l_{post} & \Sigma \end{pmatrix}\right). \quad (11)$$

To account for this correlation, consider the residual \hat{r} from the regression of $\hat{\beta}$ on $l'\hat{\beta}_{post}$ under their joint unconditional distribution (11):

$$\hat{r} = \hat{\beta} - \gamma l'\hat{\beta}_{post}, \quad \gamma = \frac{\text{Cov}_\beta(\hat{\beta}, l'\hat{\beta}_{post})}{\text{Var}_\beta(l'\hat{\beta}_{post})} = \frac{\Sigma l_{post}}{\sigma_{post}^2}.$$

The conditional distribution of $l'\hat{\beta}_{post}$ given $\{\hat{\beta} \in \mathcal{B}_{post}, \hat{r} = r\}$ is a $N(l'\beta_{post}, \sigma_{post}^2)$ distribution truncated to the set $\mathcal{Z}_{post}(r) = \{z \in \mathbb{R} : r + \gamma z \in \mathcal{B}_{post}\}$:

$$l'\hat{\beta}_{post} | \{\hat{\beta} \in \mathcal{B}_{post}, \hat{r} = r\} \stackrel{d}{=} l'\hat{\beta}_{post} | \{l'\hat{\beta}_{post} \in \mathcal{Z}_{post}(r), \hat{r} = r\} \stackrel{d}{=} l'\hat{\beta}_{post} | \{l'\hat{\beta}_{post} \in \mathcal{Z}_{post}(r)\}, \quad (12)$$

where the second equality follows from independence of \hat{r} and $l'\hat{\beta}_{post}$. To proceed, let

$$F_{TN}(z; \mu, \sigma^2, \mathcal{Z}) = \frac{\int_{-\infty}^z \phi\left(\frac{t-\mu}{\sigma}\right) \mathbf{1}\{t \in \mathcal{Z}\} dt}{\int_{-\infty}^{\infty} \phi\left(\frac{t-\mu}{\sigma}\right) \mathbf{1}\{t \in \mathcal{Z}\} dt} \quad (13)$$

denote the cumulative distribution function (CDF) of the $N(\mu, \sigma^2)$ distribution truncated to a set \mathcal{Z} , where $\phi(z)$ denotes the probability density function (PDF) of the $N(0, 1)$ distribution. Letting $\hat{\mathcal{Z}}_{post} = \mathcal{Z}_{post}(\hat{r})$, equation (12) and the probability integral transform yields

$$F_{TN}(l'\hat{\beta}_{post}; l'\beta_{post}, \sigma_{post}^2, \hat{\mathcal{Z}}_{post}) | \{\hat{\beta} \in \mathcal{B}_{post}, \hat{r} = r\} \sim U(0, 1), \quad \forall r, \quad \forall \beta.$$

The above holds for all r , so we can integrate out the residual \hat{r} to obtain

$$F_{TN}(l'\hat{\beta}_{post}; l'\beta_{post}, \sigma_{post}^2, \hat{\mathcal{Z}}_{post}) | \hat{\beta} \in \mathcal{B}_{post} \sim U(0, 1), \quad \forall \beta. \quad (14)$$

Thus, conditional on $\hat{\beta} \in \mathcal{B}_{post}$, the probability of observing $F_{TN}(l'\hat{\beta}_{post}; l'\beta_{post}, \sigma_{post}^2, \hat{\mathcal{Z}}_{post}) \geq 1 - \alpha$ is equal to α . The truncated normal distribution has strict monotone likelihood ratio in μ , and hence its CDF is strictly decreasing in the potential values $\mu \in \mathbb{R}$ of $l'\beta_{post}$. Thus, given $\hat{\beta} \in \mathcal{B}_{post}$, there exists unique $\hat{\mu}_\alpha^*$ for which

$$F_{TN}(l'\hat{\beta}_{post}; \hat{\mu}_\alpha^*, \sigma_{post}^2, \hat{\mathcal{Z}}_{post}) = 1 - \alpha. \quad (15)$$

The estimator $\hat{\mu}_\alpha^*$ is quantile conditionally unbiased in the sense of criteria (7):

$$\mathbb{P}_\beta \{l'\beta_{post} \leq \hat{\mu}_\alpha^* | \hat{\beta} \in \mathcal{B}_{post}\} = \mathbb{P}_\beta \{F_{TN}(l'\hat{\beta}_{post}; l'\beta_{post}, \sigma_{post}^2, \hat{\mathcal{Z}}_{post}) \geq 1 - \alpha | \hat{\beta} \in \mathcal{B}_{post}\} = \alpha, \quad \forall \beta.$$

3.4.3 Proposed Procedures

Based on estimator $\hat{\mu}_\alpha^*$, the proposed conditional inference procedures are as follows.

Confidence Intervals. We propose $CI_\alpha^*(X) = [\hat{\mu}_{\alpha/2}^*, \hat{\mu}_{1-\alpha/2}^*]$ as confidence intervals for $l'\beta_{post}$, which yields

$$\mathbb{P}_\beta \{l'\beta_{post} \in CI_\alpha^*(X) | \hat{\beta} \in \mathcal{B}_{post}\} = 1 - \alpha, \quad \forall \beta.$$

That is, $CI_\alpha^*(X)$ has correct conditional coverage for $l'\beta_{post}$.

Point Estimators. We propose $\hat{\mu}_{1/2}^*$ as point estimators for $l'\beta_{post}$, which yields

$$\mathbb{P}_\beta \{l'\beta_{post} \leq \hat{\mu}_{1/2}^* | \hat{\beta} \in \mathcal{B}_{post}\} = \frac{1}{2}, \quad \forall \beta.$$

That is, the conditional median of $\hat{\mu}_{1/2}^*$ is equal to $l'\beta_{post}$.

3.5 Statistical Properties

We now highlight two statistical properties of $\hat{\mu}_\alpha^*$. First, the estimator $\hat{\mu}_\alpha^*$ is optimal in the class of quantile conditionally unbiased estimators, in a broad sense formalized by Pfanzagl (1994). A statement of this result in our setup and notation is as follows.

Proposition 1 (Pfanzagl (1994)). *Given any $\hat{\mu}_\alpha$ that satisfies criteria (7), and any quasiconvex loss function $d \mapsto L(d, l'\beta_{post})$ that attains its minimum at $d = l'\beta_{post}$, we have*

$$\mathbb{E}_\beta[L(\hat{\mu}_\alpha^*, l'\beta_{post}) | \hat{\beta} \in \mathcal{B}_{post}] \leq \mathbb{E}_\beta[L(\hat{\mu}_\alpha, l'\beta_{post}) | \hat{\beta} \in \mathcal{B}_{post}], \quad \forall \beta.$$

Proof. See Appendix D.1. ■

Thus, for a broad class of loss functions, the estimator $\hat{\mu}_\alpha^*$ yields lower expected loss than any other quantile conditionally unbiased estimator $\hat{\mu}_\alpha$. This means there is limited scope for improving upon $\hat{\mu}_\alpha^*$ for conditional inference: $\hat{\mu}_{1/2}^*$ is an optimal point estimator for $l'\beta_{post}$ and $CI_\alpha^*(X) = [\hat{\mu}_{\alpha/2}^*, \hat{\mu}_{1-\alpha/2}^*]$ is an optimal (equal-tailed) confidence interval for $l'\beta_{post}$.

The second property is that conditional inferences with $\hat{\mu}_\alpha^*$ will agree with conventional *unconditional* inferences when the conditioning event $\{\hat{\beta} \in \mathcal{B}_{post}\}$ occurs with high probability. That is, for deviations where the use of conventional point estimators $l'\hat{\beta}_{post}$ and confidence intervals $[l'\hat{\beta}_{post} \pm z_{1-\alpha/2}\sigma_{post}]$ would yield minimal distortions, the researcher does not pay a price when using the conditional analogues $(\hat{\mu}_{1/2}^*, CI_\alpha^*(X))$ instead.

To gain intuition for this property, note $\{\hat{\beta} \in \mathcal{B}_{post}\} = \{l'\hat{\beta}_{post} \in \hat{\mathcal{Z}}_{post}\}$ and consider the case where $\mathcal{B}_{post} = \mathbb{R}^{\dim(\hat{\beta})}$, so that $\hat{\mathcal{Z}}_{post} = \mathbb{R}$. In this case, formula (13) yields

$$1 - \alpha = F_{TN}(l'\hat{\beta}_{post}; \hat{\mu}_\alpha^*, \sigma_{post}^2, \hat{\mathcal{Z}}_{post}) = \Phi\left(\frac{l'\hat{\beta}_{post} - \hat{\mu}_\alpha^*}{\sigma_{post}}\right), \quad (16)$$

where $\Phi(z)$ is the CDF of the $N(0,1)$ distribution. This is solved by $\hat{\mu}_\alpha^* = l'\hat{\beta}_{post} + z_\alpha\sigma_{post}$. Thus, when $\mathbb{P}_\beta\{\hat{\beta} \in \mathcal{B}_{post}\} \approx 1$ so that $\mathcal{B}_{post} \approx \mathbb{R}^{\dim(\hat{\beta})}$, we expect $\hat{\mu}_\alpha^* \approx l'\hat{\beta}_{post} + z_\alpha\sigma_{post}$. We formalize this intuition in the following result, based on Andrews et al. (2024, Appendix C.3).

Proposition 2 (Andrews et al. (2024)). *Hold Σ fixed and consider a sequence of sets \mathcal{B}_{post}^m and estimands β_m . If $\mathbb{P}_{\beta_m}\{\hat{\beta} \in \mathcal{B}_{post}^m\} \rightarrow 1$, then $\hat{\mu}_{\alpha,m}^*$ converges in probability to $l'\hat{\beta}_{post} + z_\alpha\sigma_{post}$ conditional on $\hat{\beta} \in \mathcal{B}_{post}^m$. That is, for each $\varepsilon > 0$, we have*

$$\mathbb{P}_{\beta_m}\left\{|\hat{\mu}_{\alpha,m}^* - (l'\hat{\beta}_{post} + z_\alpha\sigma_{post})| > \varepsilon \mid \hat{\beta} \in \mathcal{B}_{post}^m\right\} \rightarrow 0.$$

Thus, conditional on $\hat{\beta} \in \mathcal{B}_{post}^m$, $\hat{\mu}_{1/2,m}^*$ converges in probability to $l'\hat{\beta}_{post}$, and the endpoints of $CI_{\alpha,m}^*(X) = [\hat{\mu}_{\alpha/2,m}^*, \hat{\mu}_{1-\alpha/2,m}^*]$ converge in probability to the endpoints of $[l'\hat{\beta}_{post} \pm z_{1-\alpha/2}\sigma_{post}]$.

Proof. See Appendix D.2. ■

3.6 Polyhedral Conditioning

To compute $\hat{\mu}_\alpha^*$ we focus on conditioning events $\hat{\beta} \in \mathcal{B}_{post}$ such that, for a finite set of matrices $A_{post,k}$ and cutoff vectors $c_{post,k}$ yielding polyhedra $\{\hat{\beta} : A_{post,k}\hat{\beta} \leq c_{post,k}\}$, $k=1, \dots, K$, we have

$$\mathcal{B}_{post} = \bigcup_{k=1}^K \{\hat{\beta} : A_{post,k}\hat{\beta} \leq c_{post,k}\}. \quad (17)$$

That is, we restrict attention to deviations that arise from the estimates $\hat{\beta}$ falling into a union of polyhedra, where the cutoff vectors $c_{post,k}$ are implicitly allowed to depend on Σ . This setup yields computationally tractable $\hat{\mu}_\alpha^*$, and broadly accommodates the types of deviations that occur in practice, such as those based on statistical significance cutoffs.

Statistical Significance Cutoffs. Suppose $\hat{\beta}_{post}$ is of interest when the linear combination $v'\hat{\beta}_{pre}$ of PAP estimates passes some significance threshold governed by η . Let $v'\hat{\beta}_{pre} = v'_{pre}\hat{\beta}$, where v_{pre} denotes the vector induced by (i) matrix multiplication to select $\hat{\beta}_{pre}$ from $\hat{\beta}$ and (ii) vector multiplication of $\hat{\beta}_{pre}$ by v . A two-sided statistical significance cutoff yields

$$\mathcal{B}_{post} = \{\hat{\beta} : |v'\hat{\beta}_{pre}| \geq z_{1-\eta/2}\sigma_{pre}\} = \{\hat{\beta} : A_{post,1}\hat{\beta} \leq c_{post,1}\} \cup \{\hat{\beta} : A_{post,2}\hat{\beta} \leq c_{post,2}\}, \quad (18)$$

where $A_{post,1} = v'_{pre}$, $A_{post,2} = -v'_{pre}$, and $c_{post,1} = c_{post,2} = -z_{1-\eta/2}\sigma_{pre}$. On the other hand, a two-sided statistical insignificance cutoff yields

$$\mathcal{B}_{post} = \{\hat{\beta} : |v'\hat{\beta}_{pre}| \leq z_{1-\eta/2}\sigma_{pre}\} = \{\hat{\beta} : A_{post}\hat{\beta} \leq c_{post}\},$$

where $A_{post} = (-v_{pre}, v_{pre})'$ and $c_{post} = (z_{1-\eta/2}\sigma_{pre}, z_{1-\eta/2}\sigma_{pre})'$.

Under polyhedral conditioning, we obtain a convenient representation for the truncation set $\hat{\mathcal{Z}}_{post}$, based on Lee et al. (2016, Lemma 5.1). In what follows, we define the maximum over the empty set as $-\infty$ and the minimum over the empty set as $+\infty$.

Proposition 3 (Lee et al. (2016)). *Define*

$$\begin{aligned} \hat{Z}_{post,k}^- &= \max_{j \in J_k^-} \frac{(c_{post,k})_j - (A_{post,k}\hat{r})_j}{(A_{post,k}\gamma)_j}, & J_k^- &= \{j : (A_{post,k}\gamma)_j < 0\}, \\ \hat{Z}_{post,k}^+ &= \min_{j \in J_k^+} \frac{(c_{post,k})_j - (A_{post,k}\hat{r})_j}{(A_{post,k}\gamma)_j}, & J_k^+ &= \{j : (A_{post,k}\gamma)_j > 0\}, \\ \hat{Z}_{post,k}^0 &= \min_{j \in J_k^0} (c_{post,k})_j - (A_{post,k}\hat{r})_j, & J_k^0 &= \{j : (A_{post,k}\gamma)_j = 0\}. \end{aligned}$$

If the set \mathcal{B}_{post} takes the form (17), then

$$\hat{\mathcal{Z}}_{post} = \bigcup_{k=1}^K \hat{\mathcal{Z}}_{post,k}, \quad \hat{\mathcal{Z}}_{post,k} = \begin{cases} [\hat{Z}_{post,k}^-, \hat{Z}_{post,k}^+], & \hat{Z}_{post,k}^0 \geq 0, \\ \emptyset, & \hat{Z}_{post,k}^0 < 0. \end{cases}$$

Moreover, for each k , the event $\{A_{post,k} \hat{\beta} \leq c_{post,k}\}$ is equivalent to $\{l' \hat{\beta}_{post} \in \hat{\mathcal{Z}}_{post,k}\}$.

Proof. See Appendix D.3. ■

In addition to yielding computationally tractable estimators $\hat{\mu}_\alpha^*$, the polyhedral structure in (17) allow us to better interpret the truncation set $\hat{\mathcal{Z}}_{post}$, and hence the behavior of $\hat{\mu}_\alpha^*$ along relevant dimensions of the inference problem. The following section illustrates this behavior when the deviation event is characterized by a one-sided significance test.

3.7 Example of Explicit Derivation

Consider inference on post estimand $\beta_{post} \in \mathbb{R}$ conditional on PAP estimate $\hat{\beta}_{pre} \in \mathbb{R}$ crossing a one-sided $\eta = 0.05$ significance cutoff:

$$\hat{\beta}_{post} | \hat{\beta}_{pre} \geq 1.96, \quad \Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}, \quad \rho \geq 0.$$

Here ρ is the correlation between $\hat{\beta}_{pre}$ and $\hat{\beta}_{post}$, and we have assumed $\sigma_{pre} = \sigma_{post} = 1$. Let

$$CI_{0.95}^* \left(\hat{\beta}_{post} | \{ \hat{\beta}_{pre} \geq 1.96 \} \right) = [\hat{\mu}_{0.025}^*, \hat{\mu}_{0.975}^*]$$

denote the corrected (i.e., conditional) 95% interval. Figure 5 depicts how this interval (in green) varies as the realized $\hat{\beta}_{pre}$ moves further from the reporting threshold of 1.96. To give a sense of magnitudes, we (i) mark significance cutoffs $z_{1-\eta}$ on the horizontal axis for different values of η and (ii) depict the conventional 95% interval (in orange). Each panel depicts this behavior over different values of the correlation ρ between $\hat{\beta}_{pre}$ and $\hat{\beta}_{post}$.

There are two major takeaways from these plots of what drives distortions. (i) *Distance*: The closer $\hat{\beta}_{pre}$ is to the cutoff of the conditioning event (e.g., the distance between $\hat{\beta}_{pre}$ and 1.96), the greater the distortion from the conventional intervals. By contrast, when the realized $\hat{\beta}_{pre}$ is far above the cutoff, accounting for conditional reporting has a negligible effect. (ii) *Strength of Correlation*: When the dependence (governed by ρ) between $\hat{\beta}_{post}$ and the conditioning event $\{\hat{\beta}_{pre} \geq 1.96\}$ grows, the distortion becomes more pronounced. As seen in the first panel, the corrected and conventional intervals are equal when $\hat{\beta}_{post}$ and $\hat{\beta}_{pre}$ are independent ($\rho = 0$).




Figure 5: Corrected $CI_{0.95}^*(\hat{\beta}_{post} | \{\hat{\beta}_{pre} \geq 1.96\})$ vs. Conventional $CI_{0.95}(\hat{\beta}_{post})$

Remark 1. The above example highlights that the corrections produced by $\hat{\mu}_\alpha^*$ depend on the covariance between the estimator $\hat{\beta}_{post}$ and the conditioning event $\{\hat{\beta}_{pre} \geq 1.96\}$. In Appendix C, we show that this holds more broadly for the general case of estimator $l'\hat{\beta}_{post}$ and conditioning event $\cup_k \{A_{post,k}, \hat{\beta} \leq c_{post,k}\}$.

4 Feasible Inference

We have thus far assumed that the vector $\hat{\beta}$ of PAP and post estimates is normally distributed with mean β and a known covariance matrix Σ . This assumption is motivated by large-sample asymptotic results that yield asymptotic normality of estimates (e.g., under the central limit theorem) and consistency of covariance matrix estimators (e.g., under the law of large numbers). These asymptotic results underlie conventional procedures used for unconditional inference, which replace $(\hat{\beta}, \Sigma)$ with analogues $(\hat{\beta}_n, \hat{\Sigma}_n)$ constructed from a sample of size n . Following this same logic, Section 4.1 shows how to implement plug-in versions of the conditional inference procedures proposed in Section 3. In Section 4.2, we establish the uniform asymptotic validity of these plug-in procedures over a broad class of probability distributions as $n \rightarrow \infty$. We prove

this asymptotic validity in Appendix E.

4.1 Implementation

Given a sample of size n from some unknown distribution P_n , the researcher constructs PAP and post estimates, which we collect into a vector $\hat{\beta}_n$. The researcher also constructs a corresponding covariance matrix estimator $\hat{\Sigma}_n$. The researcher is interested in a linear combination $l' \hat{\beta}_{post,n}$ of some vector of post estimates $\hat{\beta}_{post,n}$. As before, we let l_{post} denote the vector induced by (i) matrix multiplication to select $\hat{\beta}_{post,n}$ from $\hat{\beta}_n$ and (ii) vector multiplication of $\hat{\beta}_{post,n}$ by l , so that $l' \hat{\beta}_{post,n} = l'_{post} \hat{\beta}_n$. The corresponding variance estimator is $\hat{\sigma}_{post,n}^2 = l'_{post} \hat{\Sigma}_n l_{post} = l' \hat{\Sigma}_{post,n} l$. As we formalize below in Section 4.2, we can think of $\hat{\beta}_n$ as corresponding to some underlying estimand $\beta(P_n)$. Thus, the goal is to conduct inference on parameter $l' \beta_{post}(P_n) = l'_{post} \beta(P_n)$.

The feasible conditional inference procedures are plug-in versions of the procedures from Section 3. That is, the feasible plug-in procedures replace all expressions that depend on $(\hat{\beta}, \hat{\Sigma})$ in Section 3 with analogous expressions that depend on $(\hat{\beta}_n, \hat{\Sigma}_n)$. We explicitly describe these plug-in procedures in Section 4.1.1, and provide a concrete example in Section 4.1.2.

4.1.1 Explicit Procedures

Step 1. Derive matrices and vectors $(A_{post,k}, \hat{c}_{post,k,n})$, $k = 1, \dots, K$.

Let $X_{1:n} \in \mathcal{X}_n$ represent the sample data. The researcher deviates to post estimates $\hat{\beta}_{post,n}$ when $X_{1:n} \in \mathcal{X}_{post,n}$, where

$$\mathcal{X}_{post,n} = \{X_{1:n} \in \mathcal{X}_n : \hat{\beta}_n \in \hat{\mathcal{B}}_{post,n}\}, \quad \hat{\mathcal{B}}_{post,n} = \bigcup_{k=1}^K \{\hat{\beta}_n : A_{post,k} \hat{\beta}_n \leq \hat{c}_{post,k,n}\}.$$

The polyhedra $\{\hat{\beta}_n : A_{post,k} \hat{\beta}_n \leq \hat{c}_{post,k,n}\}$, $k = 1, \dots, K$, are based on nonrandom matrices $A_{post,k}$ and random cutoff vectors $\hat{c}_{post,k,n}$ that may depend on $\hat{\Sigma}_n$.

Step 2. Compute truncation set $\hat{\mathcal{Z}}_{post,n}$.

To account for conditioning, compute the residual

$$\hat{r}_n = \hat{\beta}_n - \hat{\gamma}_n l' \hat{\beta}_{post,n}, \quad \hat{\gamma}_n = \frac{\hat{\Sigma}_n l_{post}}{\hat{\sigma}_{post,n}^2},$$

and compute the truncation quantities

$$\hat{Z}_{post,k,n}^- = \max_{j \in \hat{J}_{k,n}^-} \frac{(\hat{c}_{post,k,n})_j - (A_{post,k} \hat{r}_n)_j}{(A_{post,k} \hat{\gamma}_n)_j}, \quad \hat{J}_{k,n}^- = \{j : (A_{post,k} \hat{\gamma}_n)_j < 0\},$$

$$\begin{aligned}\hat{Z}_{post,k,n}^+ &= \min_{j \in \hat{J}_{k,n}^+} \frac{(\hat{c}_{post,k,n})_j - (A_{post,k} \hat{r}_n)_j}{(A_{post,k} \hat{\gamma}_n)_j}, & \hat{J}_{k,n}^+ &= \{j : (A_{post,k} \hat{\gamma}_n)_j > 0\}, \\ \hat{Z}_{post,k,n}^0 &= \min_{j \in \hat{J}_{k,n}^0} (\hat{c}_{post,k,n})_j - (A_{post,k} \hat{r}_n)_j, & \hat{J}_{k,n}^0 &= \{j : (A_{post,k} \hat{\gamma}_n)_j = 0\}.\end{aligned}$$

As in Proposition 3, the event $\{\hat{\beta}_n \in \hat{\mathcal{B}}_{post,n}\}$ is equivalent to $\{l' \hat{\beta}_{post,n} \in \hat{\mathcal{Z}}_{post,n}\}$, where

$$\hat{\mathcal{Z}}_{post,n} = \bigcup_{k=1}^K \hat{\mathcal{Z}}_{post,k,n}, \quad \hat{\mathcal{Z}}_{post,k,n} = \begin{cases} [\hat{Z}_{post,k,n}^-, \hat{Z}_{post,k,n}^+], & \hat{Z}_{post,k,n}^0 \geq 0, \\ \emptyset, & \hat{Z}_{post,k,n}^0 < 0. \end{cases}$$

Step 3. Solve for the plug-in quantile unbiased estimator $\hat{\mu}_{\alpha,n}^*$.

Given $\hat{\beta}_n \in \hat{\mathcal{B}}_{post,n}$, the plug-in quantile unbiased estimator is the unique $\hat{\mu}_{\alpha,n}^*$ that solves

$$F_{TN}(l' \hat{\beta}_{post,n}; \hat{\mu}_{\alpha,n}^*, \hat{\sigma}_{post,n}^2, \hat{\mathcal{Z}}_{post,n}) = 1 - \alpha.$$

Computation is fast, since this amounts to finding the root of strictly monotone function. The corresponding conditional inference procedures are

- (i) confidence intervals $CI_{\alpha,n}^*(X_{1:n}) = [\hat{\mu}_{\alpha/2,n}^*, \hat{\mu}_{1-\alpha/2,n}^*]$; and
- (ii) point estimators $\hat{\mu}_{1/2,n}^*$.

Section 4.2 establishes the asymptotic validity of these feasible plug-in procedures for inference on $l' \beta_{post}(P_n)$ as $n \rightarrow \infty$.

4.1.2 Concrete Example

Consider $\hat{\beta}_n = (\hat{\beta}_{pre,n}, \hat{\beta}_{post,n})' \in \mathbb{R}^2$ and $\beta(P_n) = (\beta_{pre}(P_n), \beta_{post}(P_n))' \in \mathbb{R}^2$. We want inference on post estimand $\beta_{post}(P_n)$ conditional on the PAP estimate $\hat{\beta}_{pre,n}$ crossing a two-sided statistical significance cutoff governed by $\eta \in (0,1)$:

$$\hat{\mathcal{B}}_{post,n} = \left\{ \begin{pmatrix} \hat{\beta}_{pre,n} \\ \hat{\beta}_{post,n} \end{pmatrix} : |\hat{\beta}_{pre,n}| \geq z_{1-\eta/2} \hat{\sigma}_{pre,n} \right\}, \quad \hat{\Sigma}_n = \begin{pmatrix} \hat{\sigma}_{pre,n}^2 & \hat{\sigma}_{cov,n} \\ \hat{\sigma}_{cov,n} & \hat{\sigma}_{post,n}^2 \end{pmatrix}, \quad \hat{\sigma}_{cov,n} > 0,$$

where $\hat{\sigma}_{cov,n}$ is the covariance estimator for $\hat{\beta}_{pre,n}$ and $\hat{\beta}_{post,n}$. In this case, $l = l_{post} = (0,1)'$.

Step 1. Derive matrices and vectors $(A_{post,k}, \hat{c}_{post,k,n})$, $k = 1, \dots, K$.

In this case, the set $\hat{\mathcal{B}}_{post,n}$ takes the form of (18):

$$\hat{\mathcal{B}}_{post,n} = \left\{ \begin{pmatrix} \hat{\beta}_{pre,n} \\ \hat{\beta}_{post,n} \end{pmatrix} : \hat{\beta}_{pre,n} \leq -z_{1-\eta/2} \hat{\sigma}_{pre,n} \right\} \cup \left\{ \begin{pmatrix} \hat{\beta}_{pre,n} \\ \hat{\beta}_{post,n} \end{pmatrix} : -\hat{\beta}_{pre,n} \leq -z_{1-\eta/2} \hat{\sigma}_{pre,n} \right\}.$$

Thus, $K=2$ and

$$A_{post,1} = (1,0), \quad A_{post,2} = (-1,0), \quad \hat{c}_{post,1,n} = \hat{c}_{post,2,n} = -z_{1-\eta/2} \hat{\sigma}_{pre,n}.$$

Step 2. Compute truncation set $\hat{\mathcal{Z}}_{post,n}$.

The regression step yields

$$\hat{\gamma}_n = \frac{\hat{\Sigma}_n l_{post}}{\hat{\sigma}_{post,n}^2} = \begin{pmatrix} \frac{\hat{\sigma}_{cov,n}}{\hat{\sigma}_{post,n}^2} \\ 1 \end{pmatrix}, \quad \hat{r}_n = \hat{\beta}_n - \hat{\gamma}_n \hat{\beta}_{post,n} = \begin{pmatrix} \hat{\beta}_{pre,n} - \frac{\hat{\sigma}_{cov,n}}{\hat{\sigma}_{post,n}^2} \hat{\beta}_{post,n} \\ 0 \end{pmatrix}.$$

For $k=1$, we obtain $\hat{Z}_{post,1,n}^0 = 0$ and

$$\hat{\mathcal{Z}}_{post,1,n} = [-\infty, \hat{Z}_{post,1,n}^+], \quad \hat{Z}_{post,1,n}^+ = \frac{-z_{1-\eta/2} \hat{\sigma}_{pre,n} - \left(\hat{\beta}_{pre,n} - \frac{\hat{\sigma}_{cov,n}}{\hat{\sigma}_{post,n}^2} \hat{\beta}_{post,n} \right)}{\frac{\hat{\sigma}_{cov,n}}{\hat{\sigma}_{post,n}^2}}.$$

For $k=2$ we obtain $\hat{Z}_{post,2,n}^0 = 0$ and

$$\hat{\mathcal{Z}}_{post,2,n} = [\hat{Z}_{post,2,n}^-, +\infty], \quad \hat{Z}_{post,2,n}^- = \frac{z_{1-\eta/2} \hat{\sigma}_{pre,n} - \left(\hat{\beta}_{pre,n} - \frac{\hat{\sigma}_{cov,n}}{\hat{\sigma}_{post,n}^2} \hat{\beta}_{post,n} \right)}{\frac{\hat{\sigma}_{cov,n}}{\hat{\sigma}_{post,n}^2}}.$$

Thus, $\{|\hat{\beta}_{pre,n}| \geq z_{1-\eta/2} \hat{\sigma}_{pre,n}\} = \{\hat{\beta}_{post,n} \in \hat{\mathcal{Z}}_{post,n}\}$, where

$$\hat{\mathcal{Z}}_{post,n} = [-\infty, \hat{Z}_{post,1,n}^+] \cup [\hat{Z}_{post,2,n}^-, +\infty].$$

Step 3. Solve for the plug-in quantile unbiased estimator $\hat{\mu}_{\alpha,n}^*$.

Given $|\hat{\beta}_{pre,n}| \geq z_{1-\eta/2} \hat{\sigma}_{pre,n}$, formula (13) yields

$$F_{TN}(z; \mu, \hat{\sigma}_{post,n}^2, \hat{Z}_{post,n}) = \begin{cases} \frac{\Phi\left(\frac{z-\mu}{\hat{\sigma}_{post,n}}\right)}{\Phi\left(\frac{\hat{Z}_{post,1,n}^+ - \mu}{\hat{\sigma}_{post,n}}\right) + 1 - \Phi\left(\frac{\hat{Z}_{post,2,n}^- - \mu}{\hat{\sigma}_{post,n}}\right)}, & z < \hat{Z}_{post,1,n}^+, \\ \frac{\Phi\left(\frac{z-\mu}{\hat{\sigma}_{post,n}}\right) - \Phi\left(\frac{\hat{Z}_{post,2,n}^- - \mu}{\hat{\sigma}_{post,n}}\right) + \Phi\left(\frac{\hat{Z}_{post,1,n}^+ - \mu}{\hat{\sigma}_{post,n}}\right)}{\Phi\left(\frac{\hat{Z}_{post,1,n}^+ - \mu}{\hat{\sigma}_{post,n}}\right) + 1 - \Phi\left(\frac{\hat{Z}_{post,2,n}^- - \mu}{\hat{\sigma}_{post,n}}\right)}, & z \geq \hat{Z}_{post,2,n}^-. \end{cases}$$

The estimator $\hat{\mu}_{\alpha,n}^*$ is the unique solution to $F_{TN}(\hat{\beta}_{post,n}; \hat{\mu}_{\alpha,n}^*, \hat{\sigma}_{post,n}^2, \hat{Z}_{post,n}) = 1 - \alpha$.

4.2 Uniform Asymptotic Validity

We now show that inference based on $\hat{\mu}_{\alpha,n}^*$ is uniformly asymptotically valid, formally stated in equation (19) below. Existing results by [Andrews et al. \(2024\)](#) and [McCloskey \(2024\)](#) can be used to prove uniformity the case of $K = 1$. Here, we prove the general case for $K \geq 1$. In the proof we assume the rows of $A_{post,k}$ are nonzero, i.e., $(A_{post,k})_j \neq 0$ for all (j,k) , which rules out, for example, a deviation event of the form $\{0 \leq \hat{\sigma}_{pre,n} - \hat{\sigma}_{post,n}\}$.

Environment. We suppose that a sample of size n is drawn from some unknown distribution $P_n \in \mathcal{P}_n$, where \mathcal{P}_n is a class of probability distributions corresponding to a sample of size n . For example, given a class of distributions \mathcal{P}_0 with bounded moments, if we have i.i.d. draws of size n from some fixed $P_0 \in \mathcal{P}_0$, then $P_n = (P_0)^n$ is the product distribution and \mathcal{P}_n is the set of all products of a fixed distribution with bounded moments. As we take $n \rightarrow \infty$, there is a corresponding sequence of unknown distributions $\{P_n\} \in \times_{n=1}^{\infty} \mathcal{P}_n$, where $\times_{n=1}^{\infty} \mathcal{P}_n$ is the sequence of distribution classes, and the notation $\{P_n\} \in \times_{n=1}^{\infty} \mathcal{P}_n$ means that $P_n \in \mathcal{P}_n$ for all n . Below we let $P \in \cup_{n=1}^{\infty} \mathcal{P}_n$ index elements belonging to \mathcal{P}_n for some n .

Asymptotic Normality. We first assume that the scaled estimates $\tilde{\beta}_n = \sqrt{n} \hat{\beta}_n$ are uniformly asymptotically normal. Let BL_1 denote the set of real-valued functions that are bounded above in absolute value by one and have Lipschitz constant bounded above by one ([van der Vaart and Wellner, 1996](#), Section 1.12). Furthermore, given a candidate covariance matrix Σ , let $\lambda_{min}(\Sigma)$ and $\lambda_{max}(\Sigma)$ denote its minimum and maximum eigenvalues.

Assumption 1. There exist $(\beta(P), \Sigma(P))$ such that for $\tilde{\beta}_n(P) = \sqrt{n} \beta(P)$ and $\xi_P \sim N(0, \Sigma(P))$,

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \sup_{f \in BL_1} \left| \mathbb{E}_P[f(\tilde{\beta}_n - \tilde{\beta}_n(P))] - \mathbb{E}[f(\xi_P)] \right|.$$

Moreover, there exists finite $\bar{\lambda} > 0$ such that $1/\bar{\lambda} \leq \lambda_{\min}(\Sigma(P)) \leq \lambda_{\max}(\Sigma(P)) \leq \bar{\lambda}$ for all $P \in \mathcal{P}_n$.

Assumption 1 requires that $\sqrt{n}(\hat{\beta}_n - \beta(P))$ converges to $N(0, \Sigma(P))$ in bounded Lipschitz metric, uniformly in P . This is a standard way to define uniform convergence in distribution.¹⁶ For example, if the components of $\hat{\beta}_n$ are sample averages of unit-level observations, then uniform convergence follows from bounds on the moments of the observations and bounds on dependence across observations. Assumption 1 also requires the eigenvalues of $\Sigma(P)$ to be uniformly bounded above and away from zero, which ensures that $\sqrt{n}(\hat{\beta}_n - \beta(P))$ is stochastically bounded with nonzero asymptotic variance.

Consistent Covariance Matrix Estimation. We next assume that the scaled covariance matrix estimator $\tilde{\Sigma}_n = n\hat{\Sigma}_n$ is consistent for $\Sigma(P)$, uniformly over P .

Assumption 2. For each $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \mathbb{P}_P \left\{ \|\tilde{\Sigma}_n - \Sigma(P)\| > \varepsilon \right\} = 0.$$

If $\hat{\Sigma}_n$ is appropriate for the setting at hand (e.g., sample covariance for iid data, long-run covariance for time series data), then Assumption 2 follows from the same kind of sufficient conditions that justify Assumption 1.

Consistent Cutoff Vector Estimation. Finally, we assume that the scaled cutoff vectors $\tilde{c}_{post,k,n} = \sqrt{n}\hat{c}_{post,k,n}$ are uniformly consistent, in similar fashion to $\tilde{\Sigma}_n$.

Assumption 3. For each $k = 1, \dots, K$, there exist $c_{post,k}(P)$ such that for each $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \mathbb{P}_P \left\{ \|\tilde{c}_{post,k,n} - c_{post,k}(P)\| > \varepsilon \right\} = 0.$$

Moreover, there exists finite $\bar{\lambda}_c > 0$ such that $\|c_{post,k}(P)\| \leq \bar{\lambda}_c$ for all $P \in \mathcal{P}_n$ and k .¹⁷

For example, often the cutoff vectors take the form $\hat{c}_{post,k,n} = q_{post,k} \sqrt{C'_{post,k} \hat{\Sigma}_n C_{post,k}}$, where $C_{post,k} \neq 0$ and $q_{post,k} \neq 0$ are nonrandom vectors. Assumption 2 then implies

$$\begin{aligned} & \sup_{P \in \mathcal{P}_n} \mathbb{P}_P \left\{ \left| \sqrt{C'_{post,k} \tilde{\Sigma}_n C_{post,k}} - \sqrt{C'_{post,k} \Sigma(P) C_{post,k}} \right| > \varepsilon \right\} \\ & \leq \sup_{P \in \mathcal{P}_n} \mathbb{P}_P \left\{ \sqrt{\left| C'_{post,k} \tilde{\Sigma}_n C_{post,k} - C'_{post,k} \Sigma(P) C_{post,k} \right|} > \varepsilon \right\} \\ & \leq \sup_{P \in \mathcal{P}_n} \mathbb{P}_P \left\{ \left| C'_{post,k} (\tilde{\Sigma}_n - \Sigma(P)) C_{post,k} \right| > \varepsilon^2 \right\} \end{aligned}$$

¹⁶Examples include Andrews et al. (2024, Assumption 6) and McCloskey (2024, Assumption 2).

¹⁷This assumption is analogous to McCloskey (2024, Assumption 1).

$$\leq \sup_{P \in \mathcal{P}_n} \mathbb{P}_P \left\{ \left\| \tilde{\Sigma}_n - \Sigma(P) \right\| > \frac{\varepsilon^2}{\|C_{post,k}\|^2} \right\} \rightarrow 0,$$

where the first inequality follows from $|\sqrt{x} - \sqrt{y}| \leq \sqrt{|x - y|}$ for $x, y \geq 0$, and the last equality follows from the matrix operator norm bound. Thus, under Assumption 2, the above $\tilde{c}_{post,k,n} = q_{post,k} \sqrt{C'_{post,k} \tilde{\Sigma}_n C_{post,k}}$ satisfies Assumption 3 with $c_{post,k}(P) = q_{post,k} \sqrt{C'_{post,k} \Sigma(P) C_{post,k}}$.

Uniform Asymptotic Validity. Let $\beta_{post}(P)$ denote the components of the centering vector $\beta(P)$ corresponding to $\hat{\beta}_{post,n}$. Under Assumptions 1-3, we show in Appendix E that inference based on $\hat{\mu}_{\alpha,n}^*$ is uniformly asymptotically valid in the sense that

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \left| \mathbb{P}_P \left\{ l' \beta_{post}(P) \leq \hat{\mu}_{\alpha,n}^* \mid \hat{\beta}_n \in \hat{\mathcal{B}}_{post,n} \right\} - \alpha \right| \mathbb{P}_P \left\{ \hat{\beta}_n \in \hat{\mathcal{B}}_{post,n} \right\} = 0. \quad (19)$$

That is, $\hat{\mu}_{\alpha,n}^*$ satisfies an asymptotic analogue of the quantile conditional unbiasedness criteria (7) for the normal model $\hat{\beta} \sim N(\beta, \Sigma)$.¹⁸ The required uniformity in the class of distributions \mathcal{P}_n is the asymptotic analogue of requiring that procedures in the normal model be valid regardless of the unknown mean β .¹⁹

Confidence Intervals. The above convergence implies uniformly valid conditional coverage for our proposed confidence intervals:

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \left| \mathbb{P}_P \left\{ l' \beta_{post}(P) \in CI_{\alpha,n}^*(X_{1:n}) \mid \hat{\beta}_n \in \hat{\mathcal{B}}_{post,n} \right\} - (1 - \alpha) \right| \mathbb{P}_P \left\{ \hat{\beta}_n \in \hat{\mathcal{B}}_{post,n} \right\} = 0,$$

which is the asymptotic analogue of criteria (8) from the normal model.

Point Estimators. The above convergence also implies uniformly valid conditional median-unbiasedness for our proposed point estimators:

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \left| \mathbb{P}_P \left\{ l' \beta_{post}(P) \leq \hat{\mu}_{1/2,n}^* \mid \hat{\beta}_n \in \hat{\mathcal{B}}_{post,n} \right\} - \frac{1}{2} \right| \mathbb{P}_P \left\{ \hat{\beta}_n \in \hat{\mathcal{B}}_{post,n} \right\} = 0,$$

which is the asymptotic analogue of criteria (9) from the normal model.

¹⁸The multiplication by $\mathbb{P}_P \left\{ \hat{\beta}_n \in \hat{\mathcal{B}}_{post,n} \right\}$ accounts for sequences where the probability of the conditioning event converges to zero. This is a standard way to account for such cases (Andrews et al., 2024; McCloskey, 2024).

¹⁹The impossibility results of Leeb and Pötscher (2006) do not apply here, since the above approach does not attempt to consistently estimate the conditional distribution of $l' \hat{\beta}_{post,n}$ given deviation. Rather, it accounts for conditioning by using the plug-in analogue of pivotal quantity (14) for $l' \beta_{post}$ from the normal model.

5 Application

We now return to Example 1 and apply our approach to the empirical results of [Bessone et al. \(2021\)](#), using the authors’ stated reasons for deviating to consider different possible deviation sets. Recall that the authors cite three reasons for reporting the non-preregistered pooled coefficient $\hat{\tau}_{N+NE+NI}$:

- (A) **Statistical Power:** “Those who received a night sleep treatment in addition to naps had very similar effects to those with naps only,” i.e.,

$$\mathcal{X}_{post}^A = \left\{ X \in \mathcal{X} : \frac{|\hat{\tau}_N - \hat{\tau}_{NE}|}{\text{se}(\hat{\tau}_N - \hat{\tau}_{NE})} \leq z_{1-\eta/2} \text{ and } \frac{|\hat{\tau}_N - \hat{\tau}_{NI}|}{\text{se}(\hat{\tau}_N - \hat{\tau}_{NI})} \leq z_{1-\eta/2} \right\}.$$

- (B) **Interpretability:** “Each of the night sleep treatments alone had no significant effect on this overall index,” i.e.,

$$\mathcal{X}_{post}^B = \left\{ X \in \mathcal{X} : \frac{|\hat{\tau}_E|}{\text{se}(\hat{\tau}_E)} \leq z_{1-\eta/2} \text{ and } \frac{|\hat{\tau}_I|}{\text{se}(\hat{\tau}_I)} \leq z_{1-\eta/2} \right\}.$$

- (C) **Economic Significance:** “In contrast, naps alone had a positive, marginally significant effect,” i.e.,

$$\mathcal{X}_{post}^C = \left\{ X \in \mathcal{X} : \frac{|\hat{\tau}_N|}{\text{se}(\hat{\tau}_N)} \geq z_{1-\eta/2} \right\}.$$

With these three plausible deviation sets in mind, we can consider the empirical consequences of correcting for the conditional reporting of $\hat{\tau}_{N+NE+NI}$ by conditioning on sequential intersections of these events. Table 1 shows the corrected 95% confidence intervals for the three discussed deviation sets alongside the original CI reported in the paper. The same significance cutoff of $\eta=0.05$ is used for all constraints in (A), (B), and (C). We show in subsequent tables how the choice of significance cutoff $z_{1-\eta/2}$ used for (A), (B), and (C) impacts these CIs.

Table 1: Corrected Confidence Intervals Under Selective Reporting

Conditioning Event(s) in \mathcal{X}_{post}	$CI_{0.95}^*$
No conditioning (original)	[0.0313, 0.1820]
(A)	[0.0313, 0.1820]
(A) and (B)	[-0.0001, 0.1837]
(A) and (B) and (C)	[-0.0001, 0.1837]

Conditioning on (A) vs. No Conditioning. We first observe that conditioning on (A) alone yields a $CI_{0.95}^*$ almost exactly equal to the conventional $CI_{0.95}$. Observe that (A) is made up of four inequalities:

$$\begin{aligned} -z_{1-\eta/2} &\leq \frac{\hat{\tau}_N - \hat{\tau}_{NE}}{\text{se}(\hat{\tau}_N - \hat{\tau}_{NE})} \leq z_{1-\eta/2}, \\ -z_{1-\eta/2} &\leq \frac{\hat{\tau}_N - \hat{\tau}_{NI}}{\text{se}(\hat{\tau}_N - \hat{\tau}_{NI})} \leq z_{1-\eta/2}. \end{aligned}$$

Recall from the toy example of a single constraint from Figure 5 that the corrected confidence intervals $CI_{0.95}^*$ converge to conventional confidence intervals at a sufficient distance from the cutoff for reporting. To contextualize this distance, consider the constraint closest to binding in the data,

$$\hat{\tau}_N - \hat{\tau}_{NI} \leq 1.96 \cdot \text{se}(\hat{\tau}_N - \hat{\tau}_{NI}).$$

Loosely speaking, the unchanged confidence intervals in this case reflect that the realized estimates in Bessone et al. (2021) are sufficiently “far” from the boundary of \mathcal{X}_{post}^A .²⁰ To provide a sense of the *scale* of this distance, we can consider an exercise of holding estimates $\hat{\tau}_N, \hat{\tau}_{NI}$ fixed and varying $\text{se}(\hat{\tau}_N - \hat{\tau}_{NI})$ through $\hat{\sigma}_{NI}^2$ alone. To see any appreciable difference between $CI_{0.95}^*$ and the conventional $CI_{0.95}$, one would need $\hat{\sigma}_{NI}^2$ to be around 5 times smaller than its observed value.

Conditioning on (A) and (B). Adding the condition that both $\hat{\tau}_E$ and $\hat{\tau}_I$ are near-zero imposes a constraint which is much closer to binding, and has an appreciable impact on the result of conditioning. For further intuition, Table 2 shows how the confidence intervals change when one varies the threshold for significance in (B), i.e., $z_{1-\eta/2}$.

Table 2: Conditional Confidence Intervals for \mathcal{X}_{post}^{AB}

Threshold for significance $z_{1-\eta/2}$ in (B)	$CI_{0.95}^*$
No conditioning (original)	[0.0313, 0.1820]
$\eta=0.01$	[0.02492, 0.18216]
$\eta=0.05$	[-0.00012, 0.18365]
$\eta=0.1$	[-0.04088, 0.18578]

In the data, this change is driven by the insignificance constraint on $\hat{\tau}_E$, which grows closer to binding as $\eta=0.1$ increases, i.e., the standard for significance becomes more lax.

²⁰The corrected intervals $CI_{0.95}^*$ also do not change for different choices of significance cutoffs, $z_{1-\eta}$, for $\alpha \in \{0.01, 0.05, 0.1\}$ (not shown given redundancy).

Conditioning on (A), (B), and (C). As with the addition of constraint (A) relative to no conditioning, the addition of constraint (C) to {(A) and (B)} does not have a meaningful impact. This is because $\hat{\tau}_N$ is sufficiently far from the boundary of significance for conditioning to have a meaningful impact—the same holds true for conditioning either on {(A) and (C)} or on (C) alone. Overall, the above empirical results demonstrate that, depending on the reason for deviating, the adjustments from our procedures can range from having no difference to an economically significant difference relative to conventional practice.

6 Robustness to Misspecification of Deviation Set

We have so far assumed that for any $\hat{\beta}_{post} = S_{post}(X)$, the reported deviation set $\mathcal{X}_{post} = \{X \in \mathcal{X} : S_{post} \in \mathcal{S}_{post}^X\}$ is correct. This assumption is justified when the researcher is honest and capable of articulating \mathcal{X}_{post} . In some settings, however, either of these assumptions may seem implausible. For example, an earnest researcher may face cognitive or communication costs when reporting \mathcal{X}_{post} (much resembling the costs faced during initial PAP specification). Alternately, a nefarious researcher may know their true \mathcal{X}_{post} , but strategically report a different deviation set (e.g., to obtain shorter confidence intervals).

In this section, we consider the possibility that the researcher reports some deviation set different from the truth, $\tilde{\mathcal{X}}_{post} \neq \mathcal{X}_{post}$, leading to potentially biased point estimators $\hat{\mu}_{1/2}^*(X, \tilde{\mathcal{X}}_{post})$ and confidence intervals $CI_\alpha^*(X, \tilde{\mathcal{X}}_{post})$ with incorrect coverage. We refer to $\tilde{\mathcal{X}}_{post}$ as the *reported deviation set*. In cases where the reported deviation set varies with the realized data $x \in \mathcal{X}$, $\tilde{\mathcal{X}}_{post}(x) \subseteq \mathcal{X}$ is defined more generally as a correspondence $\tilde{\mathcal{X}}_{post} : \mathcal{X} \rightrightarrows \mathcal{X}$.²¹

6.1 Impossibility Result

To frame upcoming discussion, we open with an impossibility result: Proposition 4 states that no non-trivial conditional inference procedure can guarantee valid conditional coverage when the researcher is fully unrestricted in reporting $\tilde{\mathcal{X}}_{post}$. For instance, no non-trivial procedure can insure against a nefarious researcher strategically choosing a particular $\tilde{\mathcal{X}}_{post}$ to exclude some β_0 .

Proposition 4. *For $X \in \mathbb{R}^{\dim(X)}$ continuous with full support, let $CS_\alpha(X, \tilde{\mathcal{X}})$ be any confidence set procedure for which there exists $(x_0, \tilde{\mathcal{X}}_0)$ with $x_0 \in \tilde{\mathcal{X}}_0$ such that $\beta_0 \notin CS_\alpha(x_0, \tilde{\mathcal{X}}_0)$ (i.e., the procedure is “capable of rejecting” β_0) and β_0 is not contained in $CS_\alpha(x, \tilde{\mathcal{X}}_0)$ for all x in some open*

²¹Note that defining results in terms of this correspondence abstracts from underlying assumptions about the researcher (e.g., incentives, utility, accuracy) generating the reporting behavior. Therefore, while we will point out possible assumptions about the researcher that *generate* certain reporting behavior (i.e., properties of correspondence $\tilde{\mathcal{X}}_{post}(x)$), our results are agnostic to these assumptions.

neighborhood $U \subseteq \tilde{\mathcal{X}}_0$ of x_0 .²² Then there exists $(\mathcal{X}_{post}, \tilde{\mathcal{X}}_{post})$ such that (i) the researcher deviates with positive probability for $X \in \mathcal{X}_{post}$ and (ii) conditional coverage of $CS_\alpha(X, \tilde{\mathcal{X}}_{post})$ is zero.

Proof. See Appendix D.4. ■

Reframed, Proposition 4 implies that *some* form of additional structure on reporting behavior must be imposed to guarantee valid coverage. Therefore, while the assumptions in previous sections of the researcher being honest and accurate in specifying their deviation sets may feel unpalatably strong, *some* type of assumption on $\tilde{\mathcal{X}}_{post}(x)$ must be made. Observe also that a reinterpretation of this result can be understood as implying that an arbitrarily skeptical reviewer can always find an $\tilde{\mathcal{X}}'_{post}$ which invalidates the researcher’s reported results, should they wish to. In this sense, structure on allowable deviation sets also allows the *researcher* to defend their results against arbitrarily unfavorable counter-assertions of their “true” deviation set, which we demonstrate more concretely in Section 6.3.

We now discuss two possible sources of such structure: Section 6.2 considers *partial* reports $\tilde{\mathcal{X}}_{post} \subseteq \mathcal{X}_{post}$, and Section 6.3 presents sensitivity analysis for assessing the robustness of results to $\tilde{\mathcal{X}}_{post} \neq \mathcal{X}_{post}$, more generally.

6.2 Local Reporting

It may happen that an honest researcher can describe their deviation behavior *local* to the realized $x \in \mathcal{X}$, but has difficulty articulating deviation behavior across all of \mathcal{X}_{post} . For instance, when \mathcal{X}_{post} consists of many disjoint components, the researcher may have a clear understanding of components “local” to the realized data draw $x \in \mathcal{X}$, but not for those “distant” from x . This disjoint structure is particularly likely when there are *many potential motives* for reporting $\hat{\beta}_{post}$, but only certain motives are *relevant* at any given draw $X \in \mathcal{X}$. In such a setting, articulating the full \mathcal{X}_{post} requires the researcher to enumerate all hypothetical motives for reporting $\hat{\beta}_{post}$, otherwise reporting an incomplete description of $\tilde{\mathcal{X}}_{post} \subseteq \mathcal{X}_{post}$.

We can show that under certain “local consistency” conditions in reporting behavior for some partial component $\tilde{\mathcal{X}}_{post} \subseteq \mathcal{X}_{post}$ containing the data realization (i.e., $X \in \tilde{\mathcal{X}}_{post}$), estimators $\hat{\mu}_\alpha^*(X, \tilde{\mathcal{X}}_{post})$ and confidence intervals $CI_\alpha^*(X, \tilde{\mathcal{X}}_{post})$ that condition on said partial component yield valid inferences. Note, however, that these inferences will be less precise than those based on procedures that condition on the full \mathcal{X}_{post} . For concreteness, we discuss this result in terms of a stylized example.

Example 4 (Same Specification, Different Motive). A researcher is studying the effect of peer quality on student performance. Depending on preliminary findings $\hat{\beta}_{pre}$ from the PAP, the

²²We restrict to procedures “capable of rejecting” β_0 so as to disregard trivial procedures for coverage (such as the constant set $CS_\alpha(\cdot, \cdot) \equiv \mathbb{R}^{\dim(\beta)}$).

researcher may be prompted to consider the following ex-post analysis S_{post} :

$$\text{TestScore}_i = \alpha + \beta_{post} \cdot \text{PeerScore}_i + \gamma \cdot \text{Controls}_i + \varepsilon_i, \quad (20)$$

where β_{post} measures the relationship between a student's outcome and the average test scores of their classroom peers. Depending on PAP outcomes, the researcher may pursue specification S_{post} under different interpretive lenses, leading to distinct *motives* for reporting $\hat{\beta}_{post}$. For instance, suppose the researcher employs a research design that implies quasi-random variation in peer scores across classrooms (e.g., due to institutional features). After collecting the data, the researcher performs a series of balance checks to verify successful randomization. If such checks pass, the researcher interprets (20) as a strategy for estimating a *causal peer effect* (m_1). Should the balance checks fail, however, (e.g., peer scores are found to be tightly correlated with students' prior achievement or classroom assignment), (20) is still of interest, but might now serve as a way to diagnose *sorting mechanisms* in administrative placement (m_2).

These two motives for reporting (m_1 : causal spillover; m_2 : diagnostic for sorting) correspond to mutually exclusive interpretations of the data: m_1 relies on peer scores being as-good-as-random, while m_2 assumes peer scores are non-random and structured. These interpretations then also correspond to disjoint deviation sets $\mathcal{X}_1 \cap \mathcal{X}_2 = \emptyset$ with $\mathcal{X}_1, \mathcal{X}_2 \subseteq \mathcal{X}_{post}$, defined informally as:

$$\mathcal{X}_1 = \{X \in \mathcal{X} \mid \text{peer-score variance large, peer scores plausibly exogenous}\}$$

$$\mathcal{X}_2 = \{X \in \mathcal{X} \mid \text{peer-score assignment correlated with prior performance}\}$$

When the realized data $x \in \mathcal{X}_1$ causes the researcher to view β_{post} as a “causal spillover” ($m=1$), they might not be primed to consider the case $x \in \mathcal{X}_2$ where β_{post} is a “diagnostic for sorting” ($m=2$). As a result, the researcher may report an incomplete deviation set, $\tilde{\mathcal{X}}_{post} = \mathcal{X}_1 \subseteq \mathcal{X}_{post}$.

The following proposition gives us that under certain reporting conditions, it is valid to condition on just \mathcal{X}_1 when $X \in \mathcal{X}_1$ and \mathcal{X}_2 when $X \in \mathcal{X}_2$.

Proposition 5. *Let $\mathcal{X}_{post} = \{X \in \mathcal{X} : S_{post} \in \mathcal{S}_{post}^X\}$ be the deviation set for $\hat{\beta}_{post} = S_{post}(X)$, and let $CI_\alpha(X, \tilde{\mathcal{X}})$ be a confidence interval procedure which delivers valid conditional coverage as in (6) for $X \in \tilde{\mathcal{X}}$, i.e.,*

$$\mathbb{P}_\beta \left\{ l' \beta_{post} \in CI_\alpha(X, \tilde{\mathcal{X}}) \mid X \in \tilde{\mathcal{X}} \right\} \geq 1 - \alpha, \quad \forall \beta.$$

For $\{\mathcal{X}_m\}_{m=1}^M$ a partition of \mathcal{X}_{post} , i.e., $\bigcup_m \mathcal{X}_m = \mathcal{X}_{post}$ and $\forall m \neq m', \mathcal{X}_m \cap \mathcal{X}_{m'} = \emptyset$, define “local

report function" $\tilde{\mathcal{X}}_{post}(X): \mathcal{X} \rightarrow \{\mathcal{X}_1, \dots, \mathcal{X}_M\}$ as

$$\tilde{\mathcal{X}}_{post}(X) = \mathcal{X}_m \quad \text{if } X \in \mathcal{X}_m.$$

Then $CI_\alpha(X, \tilde{\mathcal{X}}_{post}(X))$ using local report function $\tilde{\mathcal{X}}_{post}(X)$ delivers valid conditional coverage $\forall X \in \mathcal{X}_{post}$, i.e.,

$$\mathbb{P}_\beta \left\{ l' \beta_{post} \in CI_\alpha(X, \tilde{\mathcal{X}}_{post}(X)) \mid X \in \mathcal{X}_{post} \right\} \geq 1 - \alpha, \quad \forall \beta.$$

Proof. See Appendix D.5. ■

The crucial requirement here is that the researcher behaves “locally coherently” in reporting the incomplete deviation set \mathcal{X}_m for all $X \in \mathcal{X}_m$. To be concrete, consider again the partial deviation sets from Example 4. Define $\mathcal{X}_1 \subseteq \mathcal{X}_{post}$ by the following two sample-based conditions:

- (i) **Sufficient dispersion of peer scores across classrooms:** for a prespecified threshold $\delta > 0$,

$$\widehat{\text{Var}}_j(\overline{\text{PeerScore}}_j) > \delta,$$

where $\widehat{\text{Var}}_j$ denotes the *sample* variance across classrooms of the classroom mean peer score $\overline{\text{PeerScore}}_j$.

- (ii) **No baseline correlation with prior achievement:** estimating

$$\text{PreScore}_i = \eta + \rho \text{PeerScore}_i + \nu_i,$$

we fail to reject $H_0: \rho = 0$ at the 5% level (i.e., $p(\hat{\rho}) > 0.05$).

Together, these define

$$\mathcal{X}_1 = \left\{ X \in \mathcal{X} : \widehat{\text{Var}}_j(\overline{\text{PeerScore}}_j) > \delta \text{ and } p(\hat{\rho}) > 0.05 \right\}.$$

This form of “coherence” can conceptually be broken into two conditions. The first condition is that for every data draw $X \in \mathcal{X}_1$, the researcher reports the same \mathcal{X}_1 as their deviation set. For this example, the first condition would be violated if there existed some counterfactual draw $X \in \mathcal{X}_1$ for which the researcher reported a different threshold $\delta' \neq \delta$ or different significance cutoff $\eta' \neq 0.05$ for $\hat{\rho}$.

The second condition requires that the partial reports $\{\mathcal{X}_m\}_m$ be *disjoint* subsets of the sample space \mathcal{X} , i.e., $\mathcal{X}_m \cap \mathcal{X}_{m'} = \emptyset$ for all $m \neq m'$. Disjointness fits settings where the researcher has multiple, distinct interpretive frames for reporting $\hat{\beta}_{post}$, each tied to a separate region of the data, so that at any realization X only one interpretation applies. This requirement is naturally satisfied when interpretations are mutually exclusive, as in Example 4.

6.3 Sensitivity Analysis

In settings where the deviation set takes the form of a cutoff rule, the researcher may not be able to discern the exact value of the cutoff. In this sense, the true cutoff is fuzzy. For example, suppose $\hat{\beta}_{post}$ is of interest because a PAP estimate $\hat{\beta}_{pre}$ was observed to be small. That is, there exists some cutoff $\kappa_0 > 0$ for which

$$\mathcal{X}_{post} = \mathcal{X}_0 = \{X \in \mathcal{X} : |\hat{\beta}_{pre}| \leq \kappa_0\}.$$

However, the exact value of κ_0 may not be obvious to the researcher. After some introspection, the researcher reports $\tilde{\kappa}$ as an approximation to κ_0 . This yields reported deviation set

$$\tilde{\mathcal{X}}_{post} = \{X \in \mathcal{X} : |\hat{\beta}_{pre}| \leq \tilde{\kappa}\}.$$

In such settings, a natural robustness exercise is to plot conditionally valid intervals for a range of κ above and below the reported $\tilde{\kappa}$. Formally, given $\varepsilon > 0$, let $\tilde{\mathcal{K}}_\varepsilon$ be a set of κ such that $\tilde{\kappa} - \varepsilon \leq \kappa \leq \tilde{\kappa} + \varepsilon$ for each $\kappa \in \tilde{\mathcal{K}}_\varepsilon$. One can plot

$$CI_\alpha^*(X, \mathcal{X}_\kappa), \quad \mathcal{X}_\kappa = \{X \in \mathcal{X} : |\hat{\beta}_{pre}| \leq \kappa\}, \quad \forall \kappa \in \tilde{\mathcal{K}}_\varepsilon. \quad (21)$$

This allows one to assess the sensitivity of conclusions to different potential values of κ_0 .

This robustness exercise is formally justified under the condition that $|\tilde{\kappa} - \kappa_0| \leq \varepsilon$ for known ε . Under this condition, if the conclusions that the researcher reaches with $CI_\alpha^*(X, \tilde{\mathcal{X}}_{post})$ can also be reached with $CI_\alpha^*(X, \mathcal{X}_\kappa)$ for each $\kappa \in \tilde{\mathcal{K}}_\varepsilon$, then we know the researcher would reach those conclusions with the true $CI_\alpha^*(X, \mathcal{X}_0)$. While knowledge of ε may be a strong assumption, one can in practice vary ε to determine the largest ε for which the researcher's conclusions persist. Intuitively, if we believe that $\tilde{\kappa} \approx \kappa_0$, then we expect κ_0 to fall into $\tilde{\mathcal{K}}_\varepsilon$ for some ε that is not too large. Plausible departures from the reported $\tilde{\kappa}$ ought to not lead to massive changes in results.

As intuition for this exercise, see Figure 6, which considers $(\hat{\beta}_{pre}, \hat{\beta}_{post})$ from Example 2, where $\rho = 0.25$. Recall the plots from Figure 5. We consider the exercise of fixing a particular data realization $(\hat{\beta}_{pre}, \hat{\beta}_{post})$ and plotting the estimates and confidence intervals as a function of cutoff $\tilde{\kappa}$. The green lines show the corrected $CI_{0.95}^*$ for $\hat{\beta}_{post}$, with the first panel corresponding to $\mathcal{X}_{post} = \{\hat{\beta}_{pre} \geq \tilde{\kappa}\}$. As one can see, the reported $\tilde{\kappa}$ can yield very different intervals: as one reports $\tilde{\kappa}$ further and further from $\hat{\beta}_{pre}$, the confidence intervals move from including to excluding zero.

Implications for Reporting Conventions. For deviations based on statistical significance cutoffs $\kappa_\eta = z_{1-\eta/2}\sigma_{pre}$, a conventional choice of significance level is $\eta = 0.05$. By sticking to convention, there is less need for the scrutiny in (21). But for cutoffs with potentially no obvious conventions, such as those based on economic significance, the robustness exercise in (21) is




Figure 6: Effect of moving cutoff $\tilde{\kappa}$ on $CI_{0.95}^*$

useful. To avoid such scrutiny, one can preregister definitions of economic significance for their primary outcomes. Note that the above analysis is also valid for any deviation set \mathcal{X}_0 known up to a finite set of fuzzy cutoffs κ_0 . Moreover, while we focused on confidence intervals $CI_{\alpha}^*(X, \mathcal{X}_{\kappa})$, the same sensitivity analysis can be applied to point estimators $\hat{\mu}_{1/2}^*(\mathcal{X}_{\kappa})$.

7 Conclusion

This paper considers the statistical consequences of deviating from prespecified analysis. We first develop a general model of preregistration in the research process, which we use to demonstrate that PAP deviations can be viewed as a form of conditional reporting that, if left unacknowledged, yields invalid inference. Our framework yields two recommendations. First, researchers should adopt conditional inference as the relevant criteria for non-prespecified analysis. Second, given this conditional criteria, researchers should articulate their corresponding reasons for deviating, then leverage \mathcal{X}_{post} to report corrected (i.e., conditional) inferences alongside conventional (i.e., unconditional) inferences. To this end, we provide general and tractable inference procedures for obtaining confidence intervals with correct conditional coverage and point estimators that are conditionally unbiased. We formalize our conditional inference objectives in a model with normally distributed estimates, and present general procedures for constructing confidence intervals and point estimators that are conditionally valid. We then show how to implement these procedures in practice, providing uniformity guarantees for non-normal estimates. Using data from [Bessone et al. \(2021\)](#), we demonstrate that, depending on the deviation set and data realization, accounting for conditional reporting can range from having no difference to an economically significant difference relative to conventional practice. In particular, when the data puts one close to the boundary of reporting, the impact of conditioning can be large, whereas when the data realization is sufficiently far from the boundary, there is no change from

unconditional inference. This framework has direct implications for considering the validity of past non-preregistered results reported in papers.

We conclude with a discussion of the robustness of our procedures to misspecification of the reported deviation set. Our results suggest possible directions for future work. In particular, this framework may have implications for the optimal length and detail of PAP specification, yielding possible rules of thumb for researchers and journals. This framework may also suggest new paradigms for adaptive experiments with data-driven selection.

References

- Agte, P., Bernhardt, A., Field, E. M., Pande, R., and Rigol, N. (2022). Investing in the next generation: The long-run impacts of a liquidity shock. *American Economic Review*.
- Alsan, M., Durvasula, M., Gupta, H., Schwartzstein, J., and Williams, H. (2024). Representation and extrapolation: evidence from clinical trials. *The quarterly journal of economics*, 139(1):575–635.
- Anderson, M. L. and Magruder, J. (2022). Highly powered analysis plans. Technical report, National Bureau of Economic Research.
- Andrews, I. and Kasy, M. (2019). Identification of and correction for publication bias. *American Economic Review*, 109(8):2766–2794.
- Andrews, I., Kitagawa, T., and McCloskey, A. (2024). Inference on winners. *The Quarterly Journal of Economics*, 139(1):305–358.
- Baicker, K., Finkelstein, A., et al. (2025). Oregon health insurance experiment: Documents. <https://www.nber.org/programs-projects/projects-and-centers/oregon-health-insurance-experiment/oregon-health-insurance-experiment-documents>. Accessed: 2025-10-01.
- Bandiera, O., Best, M. C., Khan, A. Q., and Prat, A. (2021). The allocation of authority in organizations: A field experiment with bureaucrats. *The Quarterly Journal of Economics*, 136(4):2195–2242.
- Banerjee, A., Duflo, E., Finkelstein, A., Katz, L. F., Olken, B. A., and Sautmann, A. (2020a). In praise of moderation: Suggestions for the scope and use of pre-analysis plans for rcts in economics. Technical report, National Bureau of Economic Research.
- Banerjee, A. V., Chassang, S., Montero, S., and Snowberg, E. (2020b). A theory of experimenters: Robustness, randomization, and balance. *American Economic Review*, 110(4):1206–1230.
- Barro, R. J. and Gordon, D. B. (1983). Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12(1):101–121.
- Berger, J. O. (2013). *Statistical decision theory and Bayesian analysis*. Springer Science & Business Media.
- Bessone, P., Rao, G., Schilbach, F., Schofield, H., and Toma, M. (2021). The economic consequences of increasing sleep among the urban poor. *The Quarterly Journal of Economics*, 136(3):1887–1941.

- Bhat, B., De Quidt, J., Haushofer, J., Patel, V. H., Rao, G., Schilbach, F., and Vautrey, P.-L. P. (2022). The long-run effects of psychotherapy on depression, beliefs, and economic outcomes. Technical report, National Bureau of Economic Research.
- Brodeur, A., Carrell, S. E., Figlio, D. N., and Lusher, L. (2023). Unpacking p-hacking and publication bias. IZA Discussion Papers 16369, Institute of Labor Economics (IZA).
- Brodeur, A., Cook, N., Hartley, J. S., and Heyes, A. (2024). Do preregistration and preanalysis plans reduce p-hacking and publication bias? evidence from 15,992 test statistics and suggestions for improvement. *Journal of Political Economy Microeconomics*, 2(3):527–561.
- Brodeur, A., Cook, N., and Heyes, A. (2020). Methods matter: p-hacking and publication bias in causal analysis in economics. *American Economic Review*, 110(11):3634–3660.
- Burlig, F. (2018). Improving transparency in observational social science research: A pre-analysis plan approach. *Economics Letters*, 168:56–60.
- Christensen, G., Freese, J., and Miguel, E. (2019). *Transparent and Reproducible Social Science Research: How to Do Open Science*. University of California Press, Oakland, CA.
- Christensen, G. and Miguel, E. (2018). Transparency, reproducibility, and the credibility of economics research. *Journal of Economic Literature*, 56(3):920–80.
- Coffman, L. and Niederle, M. (2015). Pre-analysis plans have limited upside, especially where replications are feasible. *Journal of Economic Perspectives*, 29(3):81–98.
- De Angelis, C., Drazen, J. M., Frizelle, F. A., Haug, C., Hoey, J., Horton, R., Kotzin, S., Laine, C., Marušić, A., Overbeke, A. J. P. M., Schroeder, T. V., Sox, H. C., and Van Der Weyden, M. B. (2004). Clinical trial registration: A statement from the international committee of medical journal editors. *New England Journal of Medicine*, 351(12):1250–1251.
- Dean, J. T. (2024). Noise, cognitive function, and worker productivity. *American Economic Journal: Applied Economics*.
- Dube, O., MacArthur, S. J., and Shah, A. K. (2025). A cognitive view of policing. *The Quarterly Journal of Economics*, 140(1):745–791.
- Evsyukova, Y., Rusche, F., and Mill, W. (2025). Linkedout? a field experiment on discrimination in job network formation. *The Quarterly Journal of Economics*, 140(1):283–334.
- Field, E., Pande, R., Rigol, N., Schaner, S., and Troyer Moore, C. (2021). On her own account: How strengthening women’s financial control impacts labor supply and gender norms. *American Economic Review*, 111(7):2342–2375.

- Finkelstein, A. and Notowidigdo, M. J. (2019). Take-up and targeting: Experimental evidence from snap. *The Quarterly Journal of Economics*, 134(3):1505–1556.
- Finkelstein, A., Taubman, S., Wright, B., Bernstein, M., Gruber, J., Newhouse, J. P., Allen, H., Baicker, K., and Group, O. H. S. (2012). The Oregon Health Insurance Experiment: Evidence from the First Year. *The Quarterly Journal of Economics*, 127(3):1057–1106.
- Fithian, W., Sun, D., and Taylor, J. (2014). Optimal inference after model selection. *arXiv preprint arXiv:1410.2597*.
- Frederick, S., Loewenstein, G., and O’Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2):351–401.
- Giacobino, H., Huillery, E., Michel, B., and Sage, M. (2024). Schoolgirls, not brides: Education as a shield against child marriage. *American Economic Journal: Applied Economics*.
- Glennerster, R. (2017). The practicalities of running randomized evaluations: Partnerships, measurement, ethics, and transparency. In Banerjee, A. and Duflo, E., editors, *Handbook of Economic Field Experiments*, volume 1, pages 175–243. North-Holland.
- Gul, F. and Pesendorfer, W. (2001). Temptation and self-control. *Econometrica*, 69(6):1403–1435.
- Humphreys, M., de la Sierra, R. S., and van der Windt, P. (2013). Fishing, commitment, and communication: A proposal for comprehensive nonbinding research registration. *Political Analysis*, 21:1–20.
- Imai, T., Toussaert, S., Baillon, A., Dreber, A., Ertac, S., Johannesson, M., Neyse, L., and Villeval, M. C. (2025). Pre-registration and pre-analysis plans in experimental economics. IZA Discussion Paper 17821, Institute of Labor Economics (IZA). Working paper.
- Jacobson, M. and Powell, D. (2024). Price sensitivity and information barriers to the take-up of naloxone. *American Economic Review*.
- Kasy, M. and Spiess, J. (2023). Optimal pre-analysis plans: Statistical decisions subject to implementability. *arXiv preprint arXiv:2208.09638*.
- Kaur, S., Mullainathan, S., Oh, S., and Schilbach, F. (2024). Do financial concerns make workers less productive? *The Quarterly Journal of Economics*, 140(1):635–689.
- Kelley, E. M., Lane, G., and Schönholzer, D. (2024). Monitoring in small firms: Experimental evidence from kenyan public transit. *American Economic Review*.

- Kremer, M., Miguel, E., and Thornton, R. (2009). Incentives to learn. *The Review of Economics and statistics*, 91(3):437–456.
- Kreps, D. M. (1979). A representation theorem for “Preference for Flexibility”. *Econometrica*, 47(3):565–577.
- Kydland, F. E. and Prescott, E. C. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–491.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–477.
- Lee, J. D., Sun, D. L., Sun, Y., and Taylor, J. E. (2016). Exact post-selection inference, with application to the lasso. *The Annals of Statistics*.
- Leeb, H. and Pötscher, B. M. (2006). Can one estimate the conditional distribution of post-model-selection estimators? *The Annals of Statistics*, 34(5):2554–2591.
- Lehmann, E. and Romano, J. P. (2024). *Testing statistical hypotheses*. Springer.
- Lehmann, E. L. and Casella, G. (2006). *Theory of point estimation*. Springer Science & Business Media.
- Ludwig, J., Mullainathan, S., and Spiess, J. (2019). Augmenting pre-analysis plans with machine learning. In *Aea papers and proceedings*, volume 109, pages 71–76. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.
- Markovic, J., Xia, L., and Taylor, J. (2017). Unifying approach to selective inference with applications to cross-validation. *arXiv preprint arXiv:1703.06559*.
- McCloskey, A. (2024). Hybrid confidence intervals for informative uniform asymptotic inference after model selection. *Biometrika*, 111(1):109–127.
- McKenzie, D. (2012). A pre-analysis plan checklist. Accessed: March 31, 2025.
- Miguel, E. (2021). Evidence on research transparency in economics. *Journal of Economic Perspectives*, 35(3):193–214.
- O’Donoghue, T. and Rabin, M. (1999). Doing it now or later. *American Economic Review*, 89(1):103–124.
- Ofose, G. K. and Posner, D. N. (2020). Do pre-analysis plans hamper publication? *AEA Papers and Proceedings*, 110(May):70–74.

- Ofosu, G. K. and Posner, D. N. (2023). Pre-analysis plans: an early stocktaking. *Perspectives on Politics*, 21(1):174–190.
- Olken, B. A. (2015). Promises and perils of pre-analysis plans. *Journal of Economic Perspectives*, 29(3):61–80.
- Pfanzagl, J. (1979). On optimal median unbiased estimators in the presence of nuisance parameters. *The Annals of Statistics*, pages 187–193.
- Pfanzagl, J. (1994). *Parametric Statistical Theory*. De Gruyter, Berlin.
- Rafkin, C., Shree Kumar, A., and Vautrey, P.-L. (2021). When guidance changes: Government stances and public beliefs. *Journal of Public Economics*, 196:104319.
- Rao, C. R., Rao, C. R., Statistiker, M., Rao, C. R., and Rao, C. R. (1973). *Linear statistical inference and its applications*, volume 2. Wiley New York.
- Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press, Cambridge, MA.
- Simmons, J. P., Nelson, L. D., and Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22(11):1359–1366.
- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23(3):165–180.
- Tian, X. and Taylor, J. (2017). Asymptotics of selective inference. *Scandinavian Journal of Statistics*, 44(2):480–499.
- Tibshirani, R. J., Rinaldo, A., Tibshirani, R., and Wasserman, L. (2018). Uniform asymptotic inference and the bootstrap after model selection. *The Annals of Statistics*, 46(3):1255–1287.
- Van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.
- van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer.
- Walker, M. (2017). *Why We Sleep: Unlocking the Power of Sleep and Dreams*. Scribner, New York.
- Wicherts, J. M., Veldkamp, C. L. S., Augusteijn, H. E. M., Bakker, M., van Aert, R. C. M., and van Assen, M. A. L. M. (2016). Degrees of freedom in planning, running, analyzing, and reporting psychological studies: A checklist to avoid p-hacking. *Frontiers in Psychology*, 7:1832.

Appendix

The outline of the Appendix is as follows.

- Appendix A: Bayesian decision problem based on counterfactual PAPs.
- Appendix B: Conditional coverage for average coverage across studies.
- Appendix C: Covariance structure of conditioning events based on multiple polyhedra.
- Appendix D: Proofs of results from the main text.
- Appendix E: Proof of uniform asymptotic validity.

A Deviations Based on Counterfactual PAP

In Section 2.2, we illustrated how deviations can arise in a standard Bayesian decision problem when \mathcal{S} is constrained. In this section, we provide a variant of the Bayesian decision problem where deviations can arise even when \mathcal{S} is *unconstrained*. We give the setup in Section A.1, and provide a stylized example in A.2 that yields a deviation condition analogous to the one from Bessone et al. (2021), discussed in Example 1.

A.1 Counterfactual PAP Approach

If the researcher counterfactually entered the experiment with their posterior beliefs π_X , they would have preregistered

$$S_{post} \in \operatorname{argmin}_{S \in \mathcal{S}} \bar{R}(S, \pi_X), \quad \bar{R}(S, \pi_X) = \int_{\Theta} R(S, \theta) d\pi_X(\theta), \quad (22)$$

where $\bar{R}(S, \pi_X)$ is the average risk of S under π_X . We can think of S_{post} as the counterfactual PAP from solving problem (2) with π_X in the place of π . If one views \mathcal{S} as the relevant action space, then the counterfactual PAP approach (i.e., solving (2) and (22)) aligns with the conditional Bayes principle, which says to choose an action that minimizes average loss under one's current beliefs (Berger, 2013, Section 1.5.1).

By construction, the counterfactual PAP minimizes average risk over *future* data draws that are independent of X , making it a reasonable point of departure for future researchers studying similar topics. Indeed, the process of solving problems (2) and (22) is analogous to conducting a pilot study with initial PAP S_{pre} , seeing results X_{pre} , then registering updated PAP S_{post} before implementing the experiment at scale on an independent sample X_{post} . The difference here is that

the same data is used for both steps, $X = X_{pre} = X_{post}$, which must be accounted for when conducting inferences; see Section 3. Notably, deviations can arise in (22) even when \mathcal{S} is unrestricted. In this sense, the dynamic inconsistency that arises in the counterfactual PAP approach is more attributable to the shift in beliefs from π to π_X , relative to the posterior average loss approach in (3).

A.2 Stylized Example

To build further intuition for the counterfactual PAP approach, we solve problems (2) and (22) in a stylized example patterned after the deviation from Bessone et al. (2021), discussed in Example 1.

- The data $X = (\hat{\tau}_1, \hat{\tau}_2, \sigma)$ consists of treatment effect estimates $\hat{\tau}_k$ with common standard deviation σ . Given treatment effect estimands τ_k , the estimates are normally distributed $\hat{\tau}_k \sim N(\tau_k, \sigma^2)$, independently across groups k .
- The parameter $\theta = (\tau_1, \tau_2, \sigma)$ contains the treatment effect estimands τ_k .
- Under the prior distribution π , estimands are drawn $\tau_k \stackrel{iid}{\sim} N(m, v^2)$, independently of σ . In particular, the researcher believes that the estimands τ_k capture some latent causal effect m , but maintains a level of uncertainty, captured by v^2 .

Mapping to Example 1, the estimands (τ_1, τ_2) represent treatment effects for (i) workers that receive the incentives treatment and (ii) workers that receive the encouragement treatment. In this case, m represents the latent effectiveness of night sleep treatments.

The researcher wishes to estimate the treatment effect τ_1 for the first group, and chooses between estimators $w\hat{\tau}_1 + (1-w)\hat{\tau}_2$, for weights $w \in \{1, 1/2\}$. Choosing $w = 1$ yields $\hat{\tau}_1$, which is the unbiased plug-in estimator for τ_1 . Choosing $w = 1/2$ yields $(\hat{\tau}_1 + \hat{\tau}_2)/2$, which is the pooled estimator that may be biased for τ_1 , but has smaller variance than the plug-in estimator. The loss function is squared error, as in equation (1): $L(w\hat{\tau}_1 + (1-w)\hat{\tau}_2, \theta) = (w\hat{\tau}_1 + (1-w)\hat{\tau}_2 - \tau_1)^2$. The risk is given by the mean squared error (MSE):

$$R(w, \theta) = \underbrace{(1-w)^2(\tau_2 - \tau_1)^2}_{\text{squared bias}} + \underbrace{\sigma^2(w^2 + (1-w)^2)}_{\text{variance}}.$$

MSE reflects the trade-off between bias and variance associated with the choice of weights w . Note that, in this setup, the set of specifications \mathcal{S} under consideration is one-to-one with the weighting schemes $w \in \{1, 1/2\}$. The first scheme puts all the weight on the first group, while the second scheme pools the estimates according to equal shares.

Let $\mathbb{E}_\pi[\cdot]$ denote expectations with respect to the prior distribution. To solve problem (2), the researcher compares the average risk of $w = 1$ and $w = 1/2$ under their prior.

$$\bar{R}(1, \pi) = \mathbb{E}_\pi[\sigma^2], \quad \bar{R}(1/2, \pi) = \frac{v^2 + \mathbb{E}_\pi[\sigma^2]}{2}.$$

We consider a researcher who expects the sampling variance of the estimates to be low relative to the dispersion of the treatment effect distribution: $\mathbb{E}_\pi[\sigma^2] < v^2$. In this case, it is optimal for the researcher to preregister the specification that yields an unbiased estimator: $S_{pre}(X) = \hat{\tau}_1$. Intuitively, since $\mathbb{E}_\pi[\sigma^2] < v^2$, the precision gains from pooling are perceived to be low, leading the researcher to choose $w = 1$.

Upon observing $X = (\hat{\tau}_1, \hat{\tau}_2, \sigma)$, the researcher learns the true σ , and updates their prior beliefs over τ_k to posterior beliefs where the treatment effect estimands are distributed

$$\tau_k | X \sim N\left(\frac{v^2}{v^2 + \sigma^2} \hat{\tau}_k + \left(1 - \frac{v^2}{v^2 + \sigma^2}\right) m, \left(1 - \frac{v^2}{v^2 + \sigma^2}\right) v^2\right),$$

independently across k . That is, observing X leads the researcher to update their prior mean m towards the treatment effect estimates $\hat{\tau}_k$, while reducing their prior uncertainty. The extent of belief updating is governed by $v^2/(v^2 + \sigma^2)$, which measures the researcher's prior uncertainty relative to the sampling variability of the treatment effect estimates.

To solve problem (22), the researcher compares the average risk of $w = 1$ and $w = 1/2$ under their posterior:

$$\bar{R}(1, \pi_X) = \sigma^2, \quad \bar{R}(1/2, \pi_X) = \left(1 - \frac{v^2}{v^2 + \sigma^2}\right) \frac{v^2}{2} + \left(\frac{v^2}{v^2 + \sigma^2}\right)^2 \frac{(\hat{\tau}_2 - \hat{\tau}_1)^2}{4} + \frac{\sigma^2}{2}.$$

Thus, it is optimal for the researcher to deviate from their PAP if and only if the t -statistic $\hat{T} = |\hat{\tau}_2 - \hat{\tau}_1|/\sqrt{2\sigma^2}$ satisfies

$$\hat{T} \leq \sqrt{\frac{\sigma^2}{v^2} \left(1 + \frac{\sigma^2}{v^2}\right)}.$$

In particular, the deviation $S_{post}(X) = (\hat{\tau}_1 + \hat{\tau}_2)/2$ occurs when the treatment effect estimates $\hat{\tau}_k$ are sufficiently homogeneous relative to the researcher's prior beliefs. In such data realizations, the scope for bias is lower than expected, so the researcher pools the estimates.

In summary, the researcher preregisters $S_{pre}(X) = \hat{\tau}_1$ and deviates to $S_{post}(X) = (\hat{\tau}_1 + \hat{\tau}_2)/2$ when

$$\mathbb{E}_\pi[\sigma^2] \leq v^2 \leq \frac{\sigma^2}{\hat{T}^2} \left(\frac{1}{2} + \sqrt{\frac{1}{4} + \hat{T}^2}\right).$$

These inequalities reflect the scenario (i) the researcher had decided not to pool from the start

and (ii) it became optimal to pool thereafter. If either of these conditions fail, then one would not observe the above switch from not pooling to pooling.

B Average Coverage

Consider a sequence of studies $t = 1, \dots, T$. For each study, there is a baseline vector of estimates $X_t \in \mathbb{R}^d$ with mean $\mu_t \in \mathbb{R}^d$ and positive definite covariance matrix $\Sigma_t \in \mathbb{R}_{++}^{d \times d}$. The researcher in study t reports $l'_t X_t$ if and only if $X_t \in \mathcal{X}_t$, where $l_t \in \mathbb{R}^d$ is a linear combination vector and $\mathcal{X}_t \in \mathcal{L}(\mathbb{R}^d)$ is a Lebesgue measurable deviation set. In data realizations $X_t \in \mathcal{X}_t$, researchers use intervals $CI_\alpha(X_t; \Sigma_t, \mathcal{X}_t, l_t)$ to conduct inference on $l'_t \mu_t$. We suppress the dependence on $(\Sigma_t, \mathcal{X}_t, l_t)$, and simply denote the confidence intervals by $CI_{\alpha,t}(X_t) = CI_\alpha(X_t; \Sigma_t, \mathcal{X}_t, l_t)$.

The average coverage rate (ACR) across studies is

$$ACR_T = \frac{\sum_{t=1}^T \mathbb{1}\{l'_t \mu_t \in CI_{\alpha,t}(X_t)\} \mathbb{1}\{X_t \in \mathcal{X}_t\}}{\sum_{t=1}^T \mathbb{1}\{X_t \in \mathcal{X}_t\}}.$$

We want an average coverage rate that stays above $1 - \alpha$ as $T \rightarrow \infty$, in analogy to criteria (6). Otherwise, the research discipline spanned by the above sequence of studies may systematically produce misleading results. Below we establish a general sense in which conditional coverage for each study is *necessary and sufficient* for controlling average coverage across all studies.

To place structure on the above problem, we suppose there exists an unknown probability distribution $F \in \mathcal{F}$ on the sample space $\mathcal{V} = \mathbb{R}^d \times \mathbb{R}_{++}^{d \times d} \times \mathcal{L}(\mathbb{R}^d) \times \mathbb{R}^d$ such that

$$V_t = (\mu_t, \Sigma_t, \mathcal{X}_t, l_t) \stackrel{iid}{\sim} F, \quad t = 1, \dots, T,$$

where \mathcal{F} is a class of probability distributions on \mathcal{V} . This assumption approximates a research environment where similarly qualified researchers pursue independent prespecified analysis, obtain their data, and then potentially deviate from their PAPs.²³ Denoting $X = (X'_1, \dots, X'_T)'$, $\mu = (\mu'_1, \dots, \mu'_T)'$, $\Sigma = \text{diag}(\Sigma_1, \dots, \Sigma_T)$, and $V = (V_1, \dots, V_T)$, we assume that

$$X|V \stackrel{d}{=} X|\mu, \Sigma \sim N(\mu, \Sigma).$$

In particular, conditional on V , the baseline estimates are independently normally distributed across t , where $X_t|V \stackrel{d}{=} X_t|V_t \sim N(\mu_t, \Sigma_t)$. Thus, since V_t is independent across t , we have that (X_t, V_t) is independent across t . Define $\mathbb{P}_v\{\cdot\} = \mathbb{P}_F\{\cdot|V_t = v\}$, which does not depend on F under our sampling assumptions. Denote $v = (\mu_v, \Sigma_v, \mathcal{X}_v, l_v)$ for a particular realization $V_t = v$, so that $CI_{\alpha,v}(X_t) = CI_\alpha(X_t; \Sigma_v, \mathcal{X}_v, l_v)$.

²³Similar sampling assumptions are made in models of publication bias (Andrews and Kasy, 2019).

Let $w(v) = \mathbb{P}_v\{X_t \in \mathcal{X}_v\}$ denote the probability of deviation under $N(\mu_t, \Sigma_t)|V_t = v$, which does not depend on t under our sampling assumptions. The law of iterated expectations yields

$$\mathbb{E}_F[\mathbb{1}\{l'_t \mu_t \in CI_{\alpha,t}(X_t)\} \mathbb{1}\{X_t \in \mathcal{X}_t\}] = \int_v w(v) \mathbb{P}_v\{l'_v \mu_v \in CI_{\alpha,v}(X_t)|X_t \in \mathcal{X}_v\} dF(v),$$

which does not depend on t . Thus, law of large numbers implies

$$\text{plim}_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}\{l'_t \mu_t \in CI_{\alpha,t}(X_t)\} \mathbb{1}\{X_t \in \mathcal{X}_t\} = \int_v w(v) \mathbb{P}_v\{l'_v \mu_v \in CI_{\alpha,v}(X_t)|X_t \in \mathcal{X}_v\} dF(v).$$

An analogous argument for the denominator, together with Slutsky's lemma, yields

$$ACR_F = \text{plim}_{T \rightarrow \infty} ACR_T = \frac{\int_v w(v) \mathbb{P}_v\{l'_t \mu_t \in CI_{\alpha,v}(X_t)|X_t \in \mathcal{X}_v\} dF(v)}{\int_v w(v) dF(v)},$$

provided that the denominator is positive.

We are now prepared to state the result. In what follows, $\mathcal{V}_+ = \{v \in \mathcal{V} : w(v) > 0\}$ denotes the set of v for which the probability of deviation $w(v) = \mathbb{P}_v\{X_t \in \mathcal{X}_v\}$ under $N(\mu_t, \Sigma_t)|V_t = v$ is positive, and \mathcal{F}_+ denotes the set of distributions on \mathcal{V}_+ .

Proposition 6. *Suppose that \mathcal{F} satisfies $\mathcal{F}_+ \subseteq \mathcal{F}$ and $\inf_{F \in \mathcal{F}} \int_v w(v) dF(v) > 0$. Then, to obtain correct average coverage across studies in the sense that*

$$ACR_F \geq 1 - \alpha, \quad \forall F \in \mathcal{F},$$

it is necessary and sufficient to ensure conditional coverage for each study in the sense that

$$\mathbb{P}_v\{l'_v \mu_v \in CI_{\alpha,v}(X_t)|X_t \in \mathcal{X}_v\} \geq 1 - \alpha, \quad \forall v \in \mathcal{V}_+.$$

Proof of Proposition 6.

For the sufficiency direction, observe that for all $F \in \mathcal{F}$,

$$\begin{aligned} ACR_F &= \frac{\int_{v \in \mathcal{V}_+} w(v) \mathbb{P}_v\{l'_v \mu_v \in CI_{\alpha,v}(X_t)|X_t \in \mathcal{X}_v\} dF(v)}{\int_{v \in \mathcal{V}_+} w(v) dF(v)} \\ &\geq \frac{\int_{v \in \mathcal{V}_+} w(v)(1 - \alpha) dF(v)}{\int_{v \in \mathcal{V}_+} w(v) dF(v)} = 1 - \alpha. \end{aligned}$$

For the necessity direction, suppose there exists $v \in \mathcal{V}_+$ for which

$$\mathbb{P}_v\{l'_v \mu_v \in CI_{\alpha,v}(X_t) | X_t \in \mathcal{X}_v\} < 1 - \alpha.$$

But since $\mathcal{F}_+ \subseteq \mathcal{F}$, the probability distribution $F_v \in \mathcal{F}_+$ that puts probability one on v must be contained in \mathcal{F} , which yields $ACR_{F_v} < 1 - \alpha$. Thus, by contraposition, the necessity direction holds. \blacksquare

Proposition 6 shows that, if the class of distributions \mathcal{F} that generates $V_t = (\mu_t, \Sigma_t, \mathcal{X}_t, l_t)$ is rich enough, then conditional coverage in each study is necessary and sufficient for obtaining correct average coverage across studies.²⁴ Intuitively, when \mathcal{F} is rich enough, there will exist unfavorable distributions $F \in \mathcal{F}$ that invalidate unconditional inference procedures. Allowing for rich classes of \mathcal{F} in this way seems warranted, given that deviations occur precisely because researchers cannot anticipate them.

C Covariance Structure of Conditioning Event

The truncation quantities $(\hat{Z}_{post,k}^-, \hat{Z}_{post,k}^+, \hat{Z}_{post,k}^0)$ depend on the sign and magnitude of the terms $(A_{post,k}\gamma)_j$. These terms are the coefficients from regressing the conditioning event variables $A_{post,k}\hat{\beta}$ on the estimator of interest $l'\hat{\beta}_{post}$:

$$(A_{post,k}\gamma)_j = \left(A_{post,k} \frac{\Sigma l_{post}}{\sigma_{post}^2} \right)_j = \left(A_{post,k} \frac{\text{Cov}_\beta(\hat{\beta}, l'\hat{\beta}_{post})}{\text{Var}_\beta(l'\hat{\beta}_{post})} \right)_j = \left(\frac{\text{Cov}_\beta(A_{post,k}\hat{\beta}, l'\hat{\beta}_{post})}{\text{Var}_\beta(l'\hat{\beta}_{post})} \right)_j.$$

Thus, the truncated distribution used to construct $\hat{\mu}_\alpha^*$ depends on the covariance structure of the conditioning event:

$$\text{Cov}_\beta(A_{post,k}\hat{\beta}, l'\hat{\beta}_{post}), \quad k = 1, \dots, K.$$

Of course, $\hat{\mu}_\alpha^*$ also depends on the variance term $\sigma_{post}^2 = \text{Var}_\beta(l'\hat{\beta}_{post})$, but since this term does not depend on $(A_{post,k})_{k=1}^K$, we ignore it in this discussion.

To illustrate, suppose one's experiment consists of pilot data X_0 and primary data X_1 from independent samples; that is, $X = (X_0, X_1)$, where X_0 is independent of X_1 . If the deviation set for $\hat{\beta}_{post} = S_{post}(X_1)$ is based on estimates $\hat{\beta}_0 \subseteq \hat{\beta}$ computed on X_0 , then there exist matrices $(A_{post,k,0})_{k=1}^K$ such that $A_{post,k}\hat{\beta} = A_{post,k,0}\hat{\beta}_0$ for each k . In this case,

$$\text{Cov}_\beta(A_{post,k}\hat{\beta}, l'\hat{\beta}_{post}) = A_{post,k,0}\text{Cov}_\beta(\hat{\beta}_0, l'\hat{\beta}_{post}) = 0, \quad \forall k.$$

²⁴Fithian et al. (2014, Example 1) consider an analogous setup, but do not show necessity.

Thus, $[\hat{Z}_{post,k}^-, \hat{Z}_{post,k}^+] = [-\infty, +\infty]$ for all k so that $\hat{\mathcal{Z}}_{post} = [-\infty, +\infty]$. In this case, by equation (16), conditional inference based on $\hat{\mu}_\alpha^*$ reduces to conventional unconditional inference.

As highlighted above, the adjustments produced by $\hat{\mu}_\alpha^*$ depend on the covariance structure of the conditioning event. When there is zero covariance, our proposed conditional intervals and point estimators reduce to their unconditional analogues:

$$A_{post,k}\gamma = 0, \quad \forall k \quad \implies \quad \begin{pmatrix} CI_\alpha^*(X) \\ \hat{\mu}_{1/2}^* \end{pmatrix} = \begin{pmatrix} [l'\hat{\beta}_{post} \pm z_{1-\alpha/2}\sigma_{post}] \\ l'\hat{\beta}_{post} \end{pmatrix}.$$

Moreover, reductions to unconditional inference can occur even without full independence. For example, suppose the estimates $\hat{\beta}$ fall into a polyhedra k where $A_{post,k}\gamma = 0$. Then, we have $\hat{\mathcal{Z}}_{post,k} = [-\infty, +\infty]$, which yields $\hat{\mathcal{Z}}_{post} = [-\infty, +\infty]$. Notably, this can occur even if there are other polyhedra $k' \neq k$ with $A_{post,k'}\gamma \neq 0$. Thus, even without independent samples, it is possible for inferences based on $\hat{\mu}_\alpha^*$ to align with conventional ones, depending on the particular data realization.

D Proofs of Results from the Main Text

D.1 Proof of Proposition 1

It suffices to verify that our setup satisfies the conditions of Pfanzagl (1994, Theorem 5.5.13). To this end, first note that observing $\hat{\beta}$ is equivalent to observing $(l'\hat{\beta}_{post}, \hat{r})$, since $\hat{r} = \hat{\beta} - \gamma l'\hat{\beta}_{post}$ for known γ . Moreover,

$$\begin{pmatrix} l'\hat{\beta}_{post} \\ \hat{r} \end{pmatrix} \sim N \left(\begin{pmatrix} l'\beta_{post} \\ \beta - \gamma l'\beta_{post} \end{pmatrix}, \begin{pmatrix} \sigma_{post}^2 & 0 \\ 0 & \Sigma_r \end{pmatrix} \right), \quad \Sigma_r = \Sigma - \frac{\Sigma l_{post} l_{post}' \Sigma}{\sigma_{post}^2}, \quad \forall \beta \in \mathbb{R}^{\dim(\beta)}.$$

Following Rao et al. (1973, Section 8a.4), there exists a dominating measure for this class of normal distributions, with corresponding density $f(v; l'\beta_{post})f(r; \beta - \gamma l'\beta_{post})$, where

$$f(v; \theta) = \frac{1}{\sqrt{2\pi\sigma_{post}^2}} \exp\left(-\frac{(v-\theta)^2}{2\sigma_{post}^2}\right), \quad \theta \in \Theta = \{l'_{post}\beta : \beta \in \mathbb{R}^{\dim(\beta)}\} = \mathbb{R},$$

and

$$f(r; \eta) = \frac{1}{\sqrt{|2\pi\Sigma_r|_+}} \exp\left(-\frac{(r-\eta)'\Sigma_r^+(r-\eta)}{2}\right), \quad \eta \in H = \{(I - \gamma l'_{post})\beta : \beta \in \mathbb{R}^{\dim(\beta)}\},$$

where $(\Sigma_r^+, |\Sigma_r|_+)$ denote the Moore-Penrose inverse and pseudo-determinant of Σ_r , respectively. Note that $\beta \mapsto (l'\beta_{post}, \beta - \gamma l'\beta_{post})$ is an invertible transformation. Thus, it suffices to show that

$\hat{\mu}_\alpha^*$ dominates quantile conditionally unbiased estimators of the form $\hat{\mu}_\alpha = \mu_\alpha(l' \hat{\beta}_{post}, \hat{r})$ over the class of conditional densities given by

$$f(v; \theta) f(r; \eta) \frac{\mathbb{1}\{r + \gamma v \in \mathcal{B}_{post}\}}{\mathbb{P}_{\theta, \eta}\{\hat{r} + \gamma l' \hat{\beta}_{post} \in \mathcal{B}_{post}\}}, \quad (\theta, \eta) \in \Theta \times H,$$

where the aforementioned dominating measure is left implicit.

The above densities and the estimator $\hat{\mu}_\alpha^*$ satisfy the conditions of [Pfanzagl \(1994, Theorem 5.5.13\)](#). The conclusion of that theorem together with [Pfanzagl \(1994, Proposition 2.5.3\)](#) yields

$$\mathbb{E}_{\theta, \eta}[L(\hat{\mu}_\alpha^*, \theta) | \hat{\beta} \in \mathcal{B}_{post}] \leq \mathbb{E}_{\theta, \eta}[L(\hat{\mu}_\alpha, \theta) | \hat{\beta} \in \mathcal{B}_{post}], \quad \forall (\theta, \eta) \in \Theta \times H,$$

for any quasiconvex loss function $d \mapsto L(d, \theta)$ that attains its minimum at $d = \theta$. By definition of Θ and H and invertibility of $\beta \mapsto (l' \beta_{post}, \beta - \gamma l' \beta_{post})$, the conclusion follows.

D.2 Proof of Proposition 2

Let $\hat{\mu}_\alpha^\Phi = l' \hat{\beta}_{post} + z_\alpha \sigma_{post}$. To show $\mathbb{P}_{\beta_m}\{|\hat{\mu}_{\alpha, m}^* - \hat{\mu}_\alpha^\Phi| > \varepsilon | \hat{\beta} \in \mathcal{B}_{post}^m\} \rightarrow 0$, it suffices to show

$$\mathbb{P}_{\beta_m}\{|\hat{\mu}_{\alpha, m}^* - \hat{\mu}_\alpha^\Phi| > \varepsilon | \hat{\beta} \in \mathcal{B}_{post}^m, \hat{r}_m\} \xrightarrow{p} 0$$

Since $\mathbb{P}_{\beta_m}\{\hat{\beta} \in \mathcal{B}_{post}^m\} \rightarrow 1$ implies $\mathbb{P}_{\beta_m}\{\hat{\beta} \in \mathcal{B}_{post}^m | \hat{r}_m\} \xrightarrow{p} 1$, to show that the above converges in probability to zero when $\mathbb{P}_{\beta_m}\{\hat{\beta} \in \mathcal{B}_{post}^m\} \rightarrow 1$, it suffices to show that

$$\mathbb{P}_{\beta_m}\{\hat{\beta} \in \mathcal{B}_{post}^m | \hat{r}_m\} \xrightarrow{p} 1 \implies \mathbb{P}_{\beta_m}\{|\hat{\mu}_{\alpha, m}^* - \hat{\mu}_\alpha^\Phi| > \varepsilon | \hat{\beta} \in \mathcal{B}_{post}^m, \hat{r}_m\} \xrightarrow{p} 0. \quad (23)$$

Following similar notation to [Andrews et al. \(2024, Proof of Proposition 9\)](#), denote

$$g(\mathcal{B}_{post}, \beta, r) = \mathbb{P}_\beta\{\hat{\beta} \in \mathcal{B}_{post} | \hat{r} = r\}, \quad \mathcal{G}(\kappa) = \{(\mathcal{B}_{post}, \beta, r) : 1 - g(\mathcal{B}_{post}, \beta, r) \leq \kappa\},$$

$$h_\varepsilon(\mathcal{B}_{post}, \beta, r) = \mathbb{P}_\beta\{|\hat{\mu}_\alpha^* - \hat{\mu}_\alpha^\Phi| > \varepsilon | \hat{\beta} \in \mathcal{B}_{post}, \hat{r} = r\}, \quad \mathcal{H}(\varepsilon, \delta) = \{(\mathcal{B}_{post}, \beta, r) : h_\varepsilon(\mathcal{B}_{post}, \beta, r) \leq \delta\}.$$

Suppose we can show that for any sequence $(\mathcal{B}_{post}^m, \beta_m, r_m)$ where $g(\mathcal{B}_{post}^m, \beta_m, r_m) \rightarrow 1$, we have $h_\varepsilon(\mathcal{B}_{post}^m, \beta_m, r_m) \rightarrow 0$. Then, for each (ε, δ) there exists $\kappa(\varepsilon, \delta)$ such that $\mathcal{G}(\kappa(\varepsilon, \delta)) \subseteq \mathcal{H}(\varepsilon, \delta)$. In this case, we have

$$\begin{aligned} \mathbb{P}_{\beta_m}\{1 - \mathbb{P}_{\beta_m}\{\hat{\beta} \in \mathcal{B}_{post}^m | \hat{r}_m\} \leq \kappa(\varepsilon, \delta)\} &= \mathbb{P}_{\beta_m}\{(\mathcal{B}_{post}^m, \beta_m, \hat{r}_m) \in \mathcal{G}(\kappa(\varepsilon, \delta))\} \\ &\leq \mathbb{P}_{\beta_m}\{(\mathcal{B}_{post}^m, \beta_m, \hat{r}_m) \in \mathcal{H}(\varepsilon, \delta)\} \\ &= \mathbb{P}_{\beta_m}\{\mathbb{P}_{\beta_m}\{|\hat{\mu}_{\alpha, m}^* - \hat{\mu}_\alpha^\Phi| > \varepsilon | \hat{\beta} \in \mathcal{B}_{post}^m, \hat{r}_m\} \leq \delta\}, \end{aligned}$$

By $\mathbb{P}_{\beta_m} \left\{ \hat{\beta} \in \mathcal{B}_{post}^m | \hat{r}_m \right\} \xrightarrow{p} 1$, the LHS converges to one for any κ , so that the RHS converges to one for any (ε, δ) , which implies (23). Thus, to conclude the proof, it suffices to show that for any sequence $(\mathcal{B}_{post}^m, \beta_m, r_m)$ where $g(\mathcal{B}_{post}^m, \beta_m, r_m) \rightarrow 1$, we have $h_\varepsilon(\mathcal{B}_{post}^m, \beta_m, r_m) \rightarrow 0$.

Let $(\mathcal{B}_{post}^m, \beta_m, r_m)$ be such a sequence. Recall $\mathcal{Z}_{post}^m(r) = \{z \in \mathbb{R} : r + \gamma z \in \mathcal{B}_{post}^m\}$ so that

$$g(\mathcal{B}_{post}^m, \beta_m, r_m) = \mathbb{P}_{\beta_m} \left\{ l' \hat{\beta}_{post} \in \mathcal{Z}_{post}^m(r_m) | \hat{r}_m = r_m \right\} = \mathbb{P}_{\beta_m} \left\{ l' \hat{\beta}_{post} \in \mathcal{Z}_{post}^m(r_m) \right\},$$

where the last equality follows from independence of \hat{r}_m and $l' \hat{\beta}_{post}$. Likewise, letting $\hat{\mu}_{\alpha, m}^*(r)$ denote the unique solution to $F_{TN}(l' \hat{\beta}_{post}; \hat{\mu}_{\alpha, m}^*(r), \sigma_{post}^2, \mathcal{Z}_{post}^m(r)) = 1 - \alpha$ given $l' \hat{\beta}_{post} \in \mathcal{Z}_{post}^m(r)$, we have

$$\begin{aligned} h_\varepsilon(\mathcal{B}_{post}^m, \beta_m, r_m) &= \mathbb{P}_{\beta_m} \left\{ |\hat{\mu}_{\alpha, m}^*(r_m) - \hat{\mu}_\alpha^\Phi| > \varepsilon | l' \hat{\beta}_{post} \in \mathcal{Z}_{post}^m(r_m), \hat{r}_m = r_m \right\} \\ &= \mathbb{P}_{\beta_m} \left\{ |\hat{\mu}_{\alpha, m}^*(r_m) - \hat{\mu}_\alpha^\Phi| > \varepsilon | l' \hat{\beta}_{post} \in \mathcal{Z}_{post}^m(r_m) \right\}. \end{aligned}$$

Thus, to show (23), it suffices to show

$$\mathbb{P}_{\beta_m} \left\{ l' \hat{\beta}_{post} \in \mathcal{Z}_{post}^m(r_m) \right\} \rightarrow 1 \quad \implies \quad \mathbb{P}_{\beta_m} \left\{ |\hat{\mu}_{\alpha, m}^*(r_m) - \hat{\mu}_\alpha^\Phi| > \varepsilon | l' \hat{\beta}_{post} \in \mathcal{Z}_{post}^m(r_m) \right\} \rightarrow 0.$$

As we argue below, this setup is amenable to Van der Vaart (2000, Lemma 5.10). To set things up, denote $\bar{\mu}_{\alpha, m}^* = (\hat{\mu}_{\alpha, m}^*(r_m) - l' \hat{\beta}_{post}) / \sigma_{post}$ and $\bar{\mathcal{Z}}_{post}^m = (\mathcal{Z}_{post}^m(r_m) - l' \hat{\beta}_{post}) / \sigma_{post}$, and note that given $0 \in \bar{\mathcal{Z}}_{post}^m$,

$$\begin{aligned} F_{TN}(l' \hat{\beta}_{post}; \hat{\mu}_{\alpha, m}^*(r_m), \sigma_{post}^2, \mathcal{Z}_{post}^m(r_m)) &= \frac{\int_{-\infty}^{l' \hat{\beta}_{post}} \phi \left(\frac{t - \hat{\mu}_{\alpha, m}^*(r_m)}{\sigma_{post}} \right) \mathbb{1} \{ t \in \mathcal{Z}_{post}^m(r_m) \} dt}{\int_{-\infty}^{\infty} \phi \left(\frac{t - \hat{\mu}_{\alpha, m}^*(r_m)}{\sigma_{post}} \right) \mathbb{1} \{ t \in \mathcal{Z}_{post}^m(r_m) \} dt} \\ &= \frac{\int_{-\infty}^0 \phi(u - \bar{\mu}_{\alpha, m}^*) \mathbb{1} \{ u \in \bar{\mathcal{Z}}_{post}^m \} du}{\int_{-\infty}^{\infty} \phi(u - \bar{\mu}_{\alpha, m}^*) \mathbb{1} \{ u \in \bar{\mathcal{Z}}_{post}^m \} du} \\ &= F_{TN}(0; \bar{\mu}_{\alpha, m}^*, 1, \bar{\mathcal{Z}}_{post}^m), \end{aligned}$$

which follows from the substitution $u = (t - l' \hat{\beta}_{post}) / \sigma_{post}$. Since $\hat{\mu}_{\alpha, m}^*(r_m)$ is the unique solution for which the LHS equals $1 - \alpha$, then $\bar{\mu}_{\alpha, m}^*$ is the unique zero to

$$\Psi_m(\bar{\mu}_{\alpha, m}^*) = (1 - \alpha) - F_{TN}(0; \bar{\mu}_{\alpha, m}^*, 1, \bar{\mathcal{Z}}_{post}^m) = 0, \quad \forall m.$$

Given $0 \in \bar{\mathcal{Z}}_{post}^m$, the maps $z \mapsto \Psi_m(z)$ are strictly increasing over $z \in \mathbb{R}$ and have unique zeros $\bar{\mu}_{\alpha, m}^*$. Thus, we appeal to the arguments of Van der Vaart (2000, Lemma 5.10). In particular,

if there exists a function $z \mapsto \Psi(z)$ such that (i) $\Psi_m(z) | \{0 \in \bar{\mathcal{Z}}_{post}^m\} \xrightarrow{p} \Psi(z)$ for all $z \in \mathbb{R}$ and (ii) $\Psi(z_\alpha - \varepsilon) < 0 < \Psi(z_\alpha + \varepsilon)$ for every $\varepsilon > 0$, then we obtain $\bar{\mu}_{\alpha,m}^* | \{0 \in \bar{\mathcal{Z}}_{post}^m\} \xrightarrow{p} z_\alpha$. Indeed, under the monotonicity property, we have

$$\mathbb{P}_{\beta_m} \left\{ \Psi_m(z_\alpha - \varepsilon) < 0 < \Psi_m(z_\alpha + \varepsilon) | 0 \in \bar{\mathcal{Z}}_{post}^m \right\} \leq \mathbb{P}_{\beta_m} \left\{ z_\alpha - \varepsilon < \bar{\mu}_{\alpha,m}^* < z_\alpha + \varepsilon | 0 \in \bar{\mathcal{Z}}_{post}^m \right\},$$

The LHS converges to one, since under conditions (i) and (ii) we have $(\Psi_m(z_\alpha - \varepsilon), \Psi_m(z_\alpha + \varepsilon)) | \{0 \in \bar{\mathcal{Z}}_{post}^m\} \xrightarrow{p} (\Psi(z_\alpha - \varepsilon), \Psi(z_\alpha + \varepsilon))$ and $\Psi(z_\alpha - \varepsilon) < 0 < \Psi(z_\alpha + \varepsilon)$. Thus, the conclusion of the lemma yields

$$\mathbb{P}_{\beta_m} \left\{ |\hat{\mu}_{\alpha,m}^*(r_m) - \hat{\mu}_\alpha^\Phi| > \varepsilon | l' \hat{\beta}_{post} \in \mathcal{Z}_{post}^m(r_m) \right\} = \mathbb{P}_{\beta_m} \left\{ |\bar{\mu}_{\alpha,m}^* - z_\alpha| > \varepsilon | 0 \in \bar{\mathcal{Z}}_{post}^m \right\} \rightarrow 0,$$

which is the desired convergence. We now show that $\mathbb{P}_{\beta_m} \left\{ l' \hat{\beta}_{post} \in \mathcal{Z}_{post}^m(r_m) \right\} \rightarrow 1$ implies the existence of $\Psi(z)$ satisfying conditions (i) and (ii), which will conclude the proof.

Consider $\Psi(z) = (1 - \alpha) - \Phi(-z)$, which has unique zero at z_α , and hence satisfies condition (ii). Now, in considering $\mathbb{P}_{\beta_m} \left\{ |\Psi_m(z) - \Psi(z)| > \varepsilon | 0 \in \bar{\mathcal{Z}}_{post}^m \right\}$, we bound the inner absolute value as follows. Letting $\xi \stackrel{d}{=} l' \hat{\beta}_{post} - l'_{post} \beta_m \stackrel{\beta_m}{\sim} N(0, \sigma_{post}^2)$, we have

$$\begin{aligned} |\Psi_m(z) - \Psi(z)| &= \left| F_{TN}(0; z, 1, \bar{\mathcal{Z}}_{post}^m) - \Phi(-z) \right| \\ &= \left| F_{TN}(\xi; \xi + z\sigma_{post}, \sigma_{post}^2, \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m) - \Phi(-z) \right| \\ &\leq \frac{\left| \int_{-\infty}^{\xi} \phi\left(\frac{t - (\xi + z\sigma_{post})}{\sigma_{post}}\right) (\mathbb{1}\{t \in \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m\} - 1) dt \right|}{\int_{-\infty}^{\infty} \phi\left(\frac{t - (\xi + z\sigma_{post})}{\sigma_{post}}\right) \mathbb{1}\{t \in \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m\} dt} \\ &\quad + \left| \frac{1}{\int_{-\infty}^{\infty} \frac{1}{\sigma_{post}} \phi\left(\frac{t - (\xi + z\sigma_{post})}{\sigma_{post}}\right) \mathbb{1}\{t \in \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m\} dt} - 1 \right| \Phi(-z) \\ &\leq \frac{\int_{-\infty}^{\infty} \phi\left(\frac{t - (\xi + z\sigma_{post})}{\sigma_{post}}\right) \mathbb{1}\{t \notin \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m\} dt}{\int_{-\infty}^{\infty} \phi\left(\frac{t - (\xi + z\sigma_{post})}{\sigma_{post}}\right) \mathbb{1}\{t \in \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m\} dt} \\ &\quad + \left| \frac{1}{\int_{-\infty}^{\infty} \frac{1}{\sigma_{post}} \phi\left(\frac{t - (\xi + z\sigma_{post})}{\sigma_{post}}\right) \mathbb{1}\{t \in \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m\} dt} - 1 \right| \Phi(-z) \end{aligned}$$

$$= f_m(\xi, z).$$

For a given realization $\xi = v$, and letting $\xi_{z,v} \sim N(v + z\sigma_{post}, \sigma_{post}^2)$, the above is equal to

$$f_m(v, z) = \frac{\mathbb{P}\{\xi_{z,v} \notin \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m\}}{\mathbb{P}\{\xi_{z,v} \in \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m\}} + \left| \frac{1}{\mathbb{P}\{\xi_{z,v} \in \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m\}} - 1 \right| \Phi(-z)$$

Since $\mathbb{P}\{\xi \in \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m\} \rightarrow 1$ and $\xi_{z,v}$ has full support for each (z, v) , then $f_m(v, z) \rightarrow 0$ for each (z, v) . In particular, $f_m(\xi, z)$ converges to zero almost surely: $\mathbb{P}\{\lim_{m \rightarrow \infty} f_m(\xi, z) = 0\} = 1$. This implies convergence in probability $\mathbb{P}\{f_m(\xi, z) > \varepsilon\} \rightarrow 0$, and hence

$$\begin{aligned} \mathbb{P}_{\beta_m}\{|\Psi_m(z) - \Psi(z)| > \varepsilon | 0 \in \bar{\mathcal{Z}}_{post}^m\} &\leq \mathbb{P}\{f_m(\xi, z) > \varepsilon | \xi \in \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m\} \\ &\leq \frac{\mathbb{P}\{f_m(\xi, z) > \varepsilon\}}{\mathbb{P}\{\xi \in \mathcal{Z}_{post}^m(r_m) - l'_{post}\beta_m\}} \\ &\rightarrow 0. \end{aligned}$$

Thus, $\Psi(z) = (1 - \alpha) - \Phi(-z)$ satisfies condition (i).

D.3 Proof of Proposition 3

For each k , [Lee et al. \(2016, Lemma 5.1\)](#) implies

$$\begin{aligned} \{A_{post,k}\hat{\beta} \leq c_{post,k}\} &= \{A_{post,k}(\gamma l' \hat{\beta}_{post} + \hat{r}) \leq c_{post,k}\} \\ &= \{A_{post,k} \gamma l' \hat{\beta}_{post} \leq c_{post,k} - A_{post,k} \hat{r}\} \\ &= \{(A_{post,k} \gamma)_j l' \hat{\beta}_{post} \leq (c_{post,k})_j - (A_{post,k} \hat{r})_j, \forall j\} \\ &= \left\{ l' \hat{\beta}_{post} \geq \frac{(c_{post,k})_j - (A_{post,k} \hat{r})_j}{(A_{post,k} \gamma)_j}, \quad \forall j \in J_k^- \right\} \\ &\cap \left\{ l' \hat{\beta}_{post} \leq \frac{(c_{post,k})_j - (A_{post,k} \hat{r})_j}{(A_{post,k} \gamma)_j}, \quad \forall j \in J_k^+ \right\} \\ &\cap \{(c_{post,k})_j - (A_{post,k} \hat{r})_j \geq 0, \quad \forall j \in J_k^0\} \end{aligned}$$

Thus, for each k , the event $\{A_{post,k}\hat{\beta} \leq c_{post,k}\}$ is equivalent to $\{l' \hat{\beta}_{post} \in \hat{\mathcal{Z}}_{post,k}\}$. Following the same above steps, we also have

$$\hat{\mathcal{Z}}_{post} = \{z \in \mathbb{R} : \hat{r} + \gamma z \in \mathcal{B}_{post}\} = \bigcup_{k=1}^K \{z \in \mathbb{R} : A_{post,k}(\gamma z + \hat{r}) \leq c_{post,k}\}$$

$$\begin{aligned}
&= \bigcup_{k=1}^K \left\{ z \in \mathbb{R} : z \in \widehat{\mathcal{Z}}_{post,k} \right\} \\
&= \bigcup_{k=1}^K \widehat{\mathcal{Z}}_{post,k}.
\end{aligned}$$

D.4 Proof of Proposition 4

Let $(x_0, \tilde{\mathcal{X}}_0)$ be as in the conditions, so we can take open neighborhood $U \subseteq \tilde{\mathcal{X}}_0$ of x_0 such that $\beta_0 \notin CS_\alpha(x, \tilde{\mathcal{X}}_0)$ for all $x \in U$. By definition of the set

$$A = \left\{ x \in \mathbb{R}^{\dim(X)} : \exists \tilde{\mathcal{X}}_{post} \subseteq \mathbb{R}^{\dim(X)} \text{ such that } x \in \tilde{\mathcal{X}}_{post} \text{ and } \beta_0 \notin CS_\alpha(x, \tilde{\mathcal{X}}_{post}) \right\},$$

we have that every $x \in U$ belongs to A by taking $\tilde{\mathcal{X}}_{post} = \tilde{\mathcal{X}}_0$. Since U is an open neighborhood, it has positive Lebesgue measure. Thus, A contains a set of positive measure, and so A itself has positive Lebesgue measure. By continuity and full support yielding density $p(x) > 0$ for all $x \in \mathbb{R}^{\dim(X)}$, we therefore obtain

$$\mathbb{P}\{X \in A\} \geq \mathbb{P}\{X \in U\} = \int_U p(x) dx > 0,$$

meaning that A has strictly positive probability under the distribution of X . Thus, a researcher who deviates for $\mathcal{X}_{post} = A$ will do so with positive probability. Moreover, by definition of A , one can report $\tilde{\mathcal{X}}_{post}$ as a function of $x \in A$ to ensure that β_0 will never be in $CS_\alpha(x, \tilde{\mathcal{X}}_{post})$, which yields conditional coverage equal to zero.

D.5 Proof of Proposition 5

By the law of total probability over the partition $\mathcal{X}_{post} = \bigcup_m \mathcal{X}_m$, for all β ,

$$\begin{aligned}
&\mathbb{P}_\beta \left\{ l' \beta_{post} \in CI_\alpha(X, \tilde{\mathcal{X}}_{post}(X)) \mid X \in \mathcal{X}_{post} \right\} \\
&= \sum_{m=1}^M \mathbb{P}_\beta \{ X \in \mathcal{X}_m \mid X \in \mathcal{X}_{post} \} \cdot \mathbb{P}_\beta \left\{ l' \beta_{post} \in CI_\alpha(X, \tilde{\mathcal{X}}_{post}(X)) \mid X \in \mathcal{X}_m \right\} \\
&= \sum_{m=1}^M \mathbb{P}_\beta \{ X \in \mathcal{X}_m \mid X \in \mathcal{X}_{post} \} \cdot \underbrace{\mathbb{P}_\beta \{ l' \beta_{post} \in CI_\alpha(X, \mathcal{X}_m) \mid X \in \mathcal{X}_m \}}_{\geq 1-\alpha} \\
&\geq (1-\alpha) \cdot \underbrace{\sum_{m=1}^M \mathbb{P}_\beta \{ X \in \mathcal{X}_m \mid X \in \mathcal{X}_{post} \}}_{=1} \\
&= 1-\alpha.
\end{aligned}$$

E Uniform Asymptotic Validity

We establish the uniform asymptotic validity of the feasible conditional inference procedures discussed Section 4 of the main text. The following analyses and proofs follow the approaches of Andrews et al. (2024) and McCloskey (2024). To make this section self-contained, we once again summarize the environment and procedure—though with some additional discussion, notation, and remarks.

E.1 Setup

Environment. We suppose that a sample of size n is drawn from some unknown distribution $P_n \in \mathcal{P}_n$, where \mathcal{P}_n is a class of probability distributions corresponding to a sample of size n . For example, given a class of distributions \mathcal{P}_0 with bounded moments, if we have i.i.d. draws of size n from some fixed $P_0 \in \mathcal{P}_0$, then $P_n = (P_0)^n$ is the product distribution and \mathcal{P}_n is the set of all products of a fixed distribution with bounded moments. As we take $n \rightarrow \infty$, there is a corresponding sequence of unknown distributions $\{P_n\} \in \times_{n=1}^{\infty} \mathcal{P}_n$, where $\times_{n=1}^{\infty} \mathcal{P}_n$ is the sequence of distribution classes, and the notation $\{P_n\} \in \times_{n=1}^{\infty} \mathcal{P}_n$ means that $P_n \in \mathcal{P}_n$ for all n . We use $P \in \cup_{n=1}^{\infty} \mathcal{P}_n$ to index elements belonging to \mathcal{P}_n for some n .

Estimates. Given a sample of size n from the unknown P_n , the researcher constructs PAP and post estimates, which we collect into a vector $\hat{\beta}_n$, alongside a corresponding covariance matrix estimator $\hat{\Sigma}_n$. The researcher is interested in a linear combination of a particular vector of post estimates: $l' \hat{\beta}_{post,n}$. Let l_{post} denote the vector induced by (i) matrix multiplication to select $\hat{\beta}_{post,n}$ from $\hat{\beta}_n$ and (ii) multiplication of $\hat{\beta}_{post,n}$ by l . Thus, $l' \hat{\beta}_{post,n} = l'_{post} \hat{\beta}_n$. We denote the corresponding variance estimator by $\hat{\sigma}_{post,n}^2 = l'_{post} \hat{\Sigma}_n l_{post}$.

Conditioning Event. The conditioning event corresponding to $\hat{\beta}_{post,n}$ is

$$\hat{\beta}_n \in \hat{\mathcal{B}}_{post,n}, \quad \hat{\mathcal{B}}_{post,n} = \bigcup_{k=1}^K \{\hat{\beta}_n : A_{post,k} \hat{\beta}_n \leq \hat{c}_{post,k,n}\},$$

where $\{\hat{\beta}_n : A_{post,k} \hat{\beta}_n \leq \hat{c}_{post,k,n}\}$ are disjoint polyhedra based on nonrandom matrices $A_{post,k} \in \mathbb{R}^{\dim(c_k) \times \dim(\beta)}$ and random vectors $\hat{c}_{post,k,n} \in \mathbb{R}^{\dim(c_k)}$. We assume the rows of $A_{post,k}$ are nonzero, i.e., $(A_{post,k})_j \neq 0$ for all (j,k) . When convenient, we use

$$A = \begin{pmatrix} A_1 \\ \vdots \\ A_K \end{pmatrix} \in \mathbb{R}^{\dim(c) \times \dim(\beta)}, \quad c = \begin{pmatrix} c_1 \\ \vdots \\ c_K \end{pmatrix} \in \mathbb{R}^{\dim(c)}.$$

to denote stacked matrices and vectors. To account for conditioning, we consider

$$\hat{r}_n = \hat{\beta}_n - \hat{\gamma}_n l'_{post} \hat{\beta}_n, \quad \hat{\gamma}_n = \frac{\hat{\Sigma}_n l_{post}}{\hat{\sigma}_{post,n}^2}.$$

Throughout, we define the maximum over the empty set as $-\infty$ and the minimum over the empty set as $+\infty$. With that in mind, define the sets

$$\begin{aligned} \hat{Z}_{post,k,n}^- &= \max_{j \in \hat{J}_{k,n}^-} \frac{(\hat{c}_{post,k,n})_j - (A_{post,k} \hat{r}_n)_j}{(A_{post,k} \hat{\gamma}_n)_j}, & \hat{J}_{k,n}^- &= \{j : (A_{post,k} \hat{\gamma}_n)_j < 0\}, \\ \hat{Z}_{post,k,n}^+ &= \min_{j \in \hat{J}_{k,n}^+} \frac{(\hat{c}_{post,k,n})_j - (A_{post,k} \hat{r}_n)_j}{(A_{post,k} \hat{\gamma}_n)_j}, & \hat{J}_{k,n}^+ &= \{j : (A_{post,k} \hat{\gamma}_n)_j > 0\}, \\ \hat{Z}_{post,k,n}^0 &= \min_{j \in \hat{J}_{k,n}^0} (\hat{c}_{post,k,n})_j - (A_{post,k} \hat{r}_n)_j, & \hat{J}_{k,n}^0 &= \{j : (A_{post,k} \hat{\gamma}_n)_j = 0\}. \end{aligned}$$

As in the main text, the event $\{A_{post,k} \hat{\beta}_n \leq \hat{c}_{post,k,n}\}$ is equivalent to event $\{l'_{post} \hat{\beta}_n \in \hat{\mathcal{Z}}_{post,k,n}\}$, and $\{\hat{\beta}_n \in \hat{\mathcal{B}}_{post,n}\}$ is equivalent to event $\{l'_{post} \hat{\beta}_n \in \hat{\mathcal{Z}}_{post,n}\}$, where

$$\hat{\mathcal{Z}}_{post,n} = \bigcup_{k=1}^K \hat{\mathcal{Z}}_{post,k,n}, \quad \hat{\mathcal{Z}}_{post,k,n} = \begin{cases} [\hat{Z}_{post,k,n}^-, \hat{Z}_{post,k,n}^+], & \hat{Z}_{post,k,n}^0 \geq 0, \\ \emptyset, & \hat{Z}_{post,k,n}^0 < 0. \end{cases}$$

Conditional Inference. Let $z \mapsto F_{TN}(z; \mu, \sigma^2, \mathcal{Z})$ denote the CDF of the $N(\mu, \sigma^2)$ distribution truncated to a set \mathcal{Z} . This truncated CDF takes the form

$$F_{TN}(z; \mu, \sigma^2, \mathcal{Z}) = \frac{\int_{-\infty}^z \phi\left(\frac{t-\mu}{\sigma}\right) \mathbf{1}\{t \in \mathcal{Z}\} dt}{\int_{-\infty}^{\infty} \phi\left(\frac{t-\mu}{\sigma}\right) \mathbf{1}\{t \in \mathcal{Z}\} dt}.$$

The truncated normal distribution with known variance has strict monotone likelihood ratio in μ , and hence $\mu \mapsto F_{TN}(z; \mu, \sigma^2, \mathcal{Z})$ is strictly decreasing in μ . Given $\hat{\beta}_n \in \hat{\mathcal{B}}_{post,n}$, the conditionally quantile unbiased estimator (QUE) is the unique $\hat{\mu}_{\alpha,n}^*$ that solves

$$F_{TN}(l'_{post} \hat{\beta}_n; \hat{\mu}_{\alpha,n}^*, \hat{\sigma}_{post,n}^2, \hat{\mathcal{Z}}_{post,n}) = 1 - \alpha.$$

The goal is to show that inference based on $\hat{\mu}_{\alpha,n}^*$ is uniformly asymptotically valid in the sense of Proposition 7 below.

Scaled Quantities. In what follows, we consider scaled estimates $\tilde{\beta}_n = \sqrt{n}\hat{\beta}_n$ and $\tilde{\Sigma}_n = n\hat{\Sigma}_n$, which induce variance $\tilde{\sigma}_{post,n}^2 = n\hat{\sigma}_{post,n}^2$, deviation set $\tilde{\mathcal{B}}_{post,n} = \sqrt{n}\hat{\mathcal{B}}_{post,n}$ such that

$$\tilde{\beta}_n \in \tilde{\mathcal{B}}_{post,n}, \quad \tilde{\mathcal{B}}_{post,n} = \bigcup_{k=1}^K \{\tilde{\beta}_n : A_{post,k}\tilde{\beta}_n \leq \tilde{c}_{post,k,n}\}, \quad \tilde{c}_{post,k,n} = \sqrt{n}\hat{c}_{post,k,n},$$

residualization step

$$\tilde{r}_n = \tilde{\beta}_n - \tilde{\gamma}_n l'_{post} \tilde{\beta}_n = \sqrt{n}\hat{r}_n, \quad \tilde{\gamma}_n = \frac{\tilde{\Sigma}_n l_{post}}{\tilde{\sigma}_{post,n}^2} = \hat{\gamma}_n,$$

and sets

$$\begin{aligned} \tilde{Z}_{post,k,n}^- &= \max_{j \in \tilde{J}_{k,n}^-} \frac{(\tilde{c}_{post,k,n})_j - (A_{post,k}\tilde{r}_n)_j}{(A_{post,k}\tilde{\gamma}_n)_j} = \sqrt{n}\hat{Z}_{post,k,n}^-, & \tilde{J}_{k,n}^- &= \{j : (A_{post,k}\tilde{\gamma}_n)_j < 0\}, \\ \tilde{Z}_{post,k,n}^+ &= \min_{j \in \tilde{J}_{k,n}^+} \frac{(\tilde{c}_{post,k,n})_j - (A_{post,k}\tilde{r}_n)_j}{(A_{post,k}\tilde{\gamma}_n)_j} = \sqrt{n}\hat{Z}_{post,k,n}^+, & \tilde{J}_{k,n}^+ &= \{j : (A_{post,k}\tilde{\gamma}_n)_j > 0\}, \\ \tilde{Z}_{post,k,n}^0 &= \min_{j \in \tilde{J}_{k,n}^0} (\tilde{c}_{post,k,n})_j - (A_{post,k}\tilde{r}_n)_j = \sqrt{n}\hat{Z}_{post,k,n}^0, & \tilde{J}_{k,n}^0 &= \{j : (A_{post,k}\tilde{\gamma}_n)_j = 0\}, \end{aligned}$$

such that the event $\{\tilde{\beta}_n \in \tilde{\mathcal{B}}_{post,n}\}$ is equivalent to event $\{l'_{post}\tilde{\beta}_n \in \tilde{\mathcal{Z}}_{post,n}\}$, where

$$\tilde{\mathcal{Z}}_{post,n} = \bigcup_{k=1}^K \tilde{\mathcal{Z}}_{post,k,n}, \quad \tilde{\mathcal{Z}}_{post,k,n} = \begin{cases} [\tilde{Z}_{post,k,n}^-, \tilde{Z}_{post,k,n}^+], & \tilde{Z}_{post,k,n}^0 \geq 0, \\ \emptyset, & \tilde{Z}_{post,k,n}^0 < 0. \end{cases}$$

Note that $\tilde{\mathcal{Z}}_{post,k,n} = \sqrt{n}\hat{\mathcal{Z}}_{post,k,n}$ so that $\tilde{\mathcal{Z}}_{post,n} = \sqrt{n}\hat{\mathcal{Z}}_{post,n}$.

E.2 Assumptions

For ease of reference, we restate Assumptions 1-3.

Let BL_1 denote the set of real-valued functions that are bounded above in absolute value by one and have Lipschitz constant bounded above by one. Furthermore, given matrix Σ , let $\lambda_{min}(\Sigma)$ and $\lambda_{max}(\Sigma)$ denote its minimum and maximum eigenvalues.

Assumption 1. There exist $(\beta(P), \Sigma(P))$ such that for $\tilde{\beta}_n(P) = \sqrt{n}\beta(P)$ and $\xi_P \sim N(0, \Sigma(P))$,

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \sup_{f \in BL_1} |\mathbb{E}_P[f(\tilde{\beta}_n - \tilde{\beta}_n(P))] - \mathbb{E}[f(\xi_P)]|.$$

Moreover, there exists finite $\bar{\lambda} > 0$ such that $1/\bar{\lambda} \leq \lambda_{min}(\Sigma(P)) \leq \lambda_{max}(\Sigma(P)) \leq \bar{\lambda}$ for all P .

Consider the scaled covariance estimator $\tilde{\Sigma}_n = n\hat{\Sigma}_n$

Assumption 2. For each $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \mathbb{P}_P \left\{ \|\tilde{\Sigma}_n - \Sigma(P)\| > \varepsilon \right\} = 0.$$

Consider $\tilde{c}_{post,k,n} = \sqrt{n} \hat{c}_{post,k,n}$ for $\{\tilde{\beta}_n : A_{post,k} \tilde{\beta}_n \leq \tilde{c}_{post,k,n}\} = \sqrt{n} \{\hat{\beta}_n : A_{post,k} \hat{\beta}_n \leq \hat{c}_{post,k,n}\}$.

Assumption 3. For each $k = 1, \dots, K$, there exist $c_{post,k}(P)$ such that for each $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \mathbb{P}_P \left\{ \|\tilde{c}_{post,k,n} - c_{post,k}(P)\| > \varepsilon \right\} = 0.$$

Moreover, there exists finite $\bar{\lambda}_c > 0$ such that $\|c_{post,k}(P)\| \leq \bar{\lambda}_c$ for all P and k .

E.3 Results

Given $\hat{\beta}_n \in \hat{\mathcal{B}}_{post,n}$, the QUE $\hat{\mu}_{\alpha,n}^*$ uniquely solves

$$F_{TN}(l'_{post} \hat{\beta}_n; \hat{\mu}_{\alpha,n}^*, \hat{\sigma}_{post,n}^2, \hat{\mathcal{Z}}_{post,n}) = 1 - \alpha.$$

Our goal is to show that conditional inference based on $\hat{\mu}_{\alpha,n}^*$ is uniformly asymptotically valid, in the following sense.

Proposition 7. *Under Assumptions 1-3, we have*

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \left| \mathbb{P}_P \left\{ \hat{\mu}_{\alpha,n}^* \geq l'_{post} \beta(P) \mid \hat{\beta}_n \in \hat{\mathcal{B}}_{post,n} \right\} - \alpha \right| \mathbb{P}_P \left\{ \hat{\beta}_n \in \hat{\mathcal{B}}_{post,n} \right\} = 0.$$

To prove Proposition 7, we state and prove a series of lemmas.

Sufficient Conditions Based on Subsequences. The first lemma shows that, to verify the uniform asymptotic validity in Proposition 7, it suffices to verify a type of pointwise asymptotic validity for every subsequence that satisfies a set of conditions. Let $\tilde{\mu}_{\alpha,n}^* = \sqrt{n} \hat{\mu}_{\alpha,n}^*$.

Lemma 1. *Suppose there exists finite $\bar{\lambda}, \bar{\lambda}_c > 0$ such that $1/\bar{\lambda} \leq \lambda_{min}(\Sigma(P)) \leq \lambda_{max}(\Sigma(P)) \leq \bar{\lambda}$ and $\|c_{post,k}(P)\| \leq \bar{\lambda}_c$, for all P and $k = 1, \dots, K$. Then, to show that*

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \left| \mathbb{P}_P \left\{ \tilde{\mu}_{\alpha,n}^* \geq l'_{post} \tilde{\beta}_n(P) \mid \tilde{\beta}_n \in \tilde{\mathcal{B}}_{post,n} \right\} - \alpha \right| \mathbb{P}_P \left\{ \tilde{\beta}_n \in \tilde{\mathcal{B}}_{post,n} \right\} = 0,$$

it suffices to show that for all subsequences $\{n_s\} \subseteq \{n\}$ and $\{P_{n_s}\} \in \times_{s=1}^{\infty} \mathcal{P}_{n_s}$ with

- (i) $\mathbb{P}_{P_{n_s}} \left\{ \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post,n_s} \right\} \rightarrow p^* \in (0, 1];$
- (ii) $A_{post,k} \tilde{\beta}_{n_s}(P_{n_s}) \rightarrow a_{post,k}^* \in [-\infty, +\infty]^{\dim(c_k)}$ for all k ;

(iii) $\Sigma(P_{n_s}) \rightarrow \Sigma^* \in \mathcal{M} = \{\Sigma : 1/\bar{\lambda} \leq \lambda_{\min}(\Sigma) \leq \lambda_{\max}(\Sigma) \leq \bar{\lambda}\}$; and

(iv) $c_{post,k}(P_{n_s}) \rightarrow c_{post,k}^* \in [-\bar{\lambda}_c, \bar{\lambda}_c]^{\dim(c_k)}$ for all k ;

we have that

$$\lim_{s \rightarrow \infty} \mathbb{P}_{P_{n_s}} \left\{ \tilde{\mu}_{\alpha, n_s}^* \geq l'_{post} \tilde{\beta}_{n_s}(P_{n_s}) \mid \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\} = \alpha.$$

Proof of Lemma 1.

Note $|x| = \max\{-x, x\}$. Thus, to show that

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \left| \mathbb{P}_P \left\{ \tilde{\mu}_{\alpha, n}^* \geq l'_{post} \tilde{\beta}_n(P) \mid \tilde{\beta}_n \in \tilde{\mathcal{B}}_{post, n} \right\} - \alpha \right| \mathbb{P}_P \left\{ \tilde{\beta}_n \in \tilde{\mathcal{B}}_{post, n} \right\} = 0,$$

it suffices to show both that

$$\liminf_{n \rightarrow \infty} \inf_{P \in \mathcal{P}_n} \left(\mathbb{P}_P \left\{ \tilde{\mu}_{\alpha, n}^* \geq l'_{post} \tilde{\beta}_n(P) \mid \tilde{\beta}_n \in \tilde{\mathcal{B}}_{post, n} \right\} - \alpha \right) \mathbb{P}_P \left\{ \tilde{\beta}_n \in \tilde{\mathcal{B}}_{post, n} \right\} \geq 0 \quad (24)$$

and that

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \left(\mathbb{P}_P \left\{ \tilde{\mu}_{\alpha, n}^* \geq l'_{post} \tilde{\beta}_n(P) \mid \tilde{\beta}_n \in \tilde{\mathcal{B}}_{post, n} \right\} - \alpha \right) \mathbb{P}_P \left\{ \tilde{\beta}_n \in \tilde{\mathcal{B}}_{post, n} \right\} \leq 0. \quad (25)$$

We claim that (24) holds if for all subsequences satisfying conditions (i)-(iv) of this lemma,

$$\liminf_{s \rightarrow \infty} \mathbb{P}_{P_{n_s}} \left\{ \tilde{\mu}_{\alpha, n_s}^* \geq l'_{post} \tilde{\beta}_{n_s}(P_{n_s}) \mid \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\} \geq \alpha. \quad (26)$$

Likewise, we claim that (25) holds if for all subsequences satisfying conditions (i)-(iv) of this lemma,

$$\limsup_{s \rightarrow \infty} \mathbb{P}_{P_{n_s}} \left\{ \tilde{\mu}_{\alpha, n_s}^* \geq l'_{post} \tilde{\beta}_{n_s}(P_{n_s}) \mid \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\} \leq \alpha. \quad (27)$$

Suppose these two claims are true. Then, to show that both (24) and (25) hold, it suffices to show that for all subsequences satisfying conditions (i)-(iv) of this lemma,

$$\lim_{s \rightarrow \infty} \mathbb{P}_{P_{n_s}} \left\{ \tilde{\mu}_{\alpha, n_s}^* \geq l'_{post} \tilde{\beta}_{n_s}(P_{n_s}) \mid \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\} = \alpha.$$

In particular, the above equation implies (26) and (27). We provide the explicit argument for the first claim that (26) implies (24).

Towards a contradiction, suppose that (26) holds for all subsequences satisfying conditions

(i)-(iv) of this lemma, but that (24) fails. If (24) fails, then there exists $\varepsilon > 0$ such that

$$\liminf_{n \rightarrow \infty} \inf_{P \in \mathcal{P}_n} \left(\mathbb{P}_P \left\{ \tilde{\mu}_{\alpha,n}^* \geq l'_{post} \tilde{\beta}_n(P) \mid \tilde{\beta}_n \in \tilde{\mathcal{B}}_{post,n} \right\} - \alpha \right) \mathbb{P}_P \left\{ \tilde{\beta}_n \in \tilde{\mathcal{B}}_{post,n} \right\} = L < -3\varepsilon.$$

By properties of the limit inferior, given $\varepsilon > 0$, there exists subsequence $\{n_q\} \subseteq \{n\}$ such that

$$\inf_{P \in \mathcal{P}_{n_q}} \left(\mathbb{P}_P \left\{ \tilde{\mu}_{\alpha,n_q}^* \geq l'_{post} \tilde{\beta}_{n_q}(P) \mid \tilde{\beta}_{n_q} \in \tilde{\mathcal{B}}_{post,n_q} \right\} - \alpha \right) \mathbb{P}_P \left\{ \tilde{\beta}_{n_q} \in \tilde{\mathcal{B}}_{post,n_q} \right\} = L_q < L + \varepsilon, \quad \forall q.$$

By properties of the infimum, given $\varepsilon > 0$ and q , there exists $P_{n_q} \in \mathcal{P}_{n_q}$ such that

$$\left(\mathbb{P}_{P_{n_q}} \left\{ \tilde{\mu}_{\alpha,n_q}^* \geq l'_{post} \tilde{\beta}_{n_q}(P_{n_q}) \mid \tilde{\beta}_{n_q} \in \tilde{\mathcal{B}}_{post,n_q} \right\} - \alpha \right) \mathbb{P}_{P_{n_q}} \left\{ \tilde{\beta}_{n_q} \in \tilde{\mathcal{B}}_{post,n_q} \right\} < L_q + \varepsilon.$$

Since $L < -3\varepsilon$ by hypothesis and $L_q < L + \varepsilon$, the above implies

$$\limsup_{q \rightarrow \infty} \left(\mathbb{P}_{P_{n_q}} \left\{ \tilde{\mu}_{\alpha,n_q}^* \geq l'_{post} \tilde{\beta}_{n_q}(P_{n_q}) \mid \tilde{\beta}_{n_q} \in \tilde{\mathcal{B}}_{post,n_q} \right\} - \alpha \right) \mathbb{P}_{P_{n_q}} \left\{ \tilde{\beta}_{n_q} \in \tilde{\mathcal{B}}_{post,n_q} \right\} < -\varepsilon. \quad (28)$$

We will show there must exist a subsequence $\{n_s\} \subseteq \{n_q\}$ that satisfies conditions (i)-(iv) of this lemma. Note that, if we can find such a subsequence, then (26) would yield

$$0 \leq \liminf_{s \rightarrow \infty} \left(\mathbb{P}_{P_{n_s}} \left\{ \tilde{\mu}_{\alpha,n_s}^* \geq l'_{post} \tilde{\beta}_{n_s}(P_{n_s}) \mid \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post,n_s} \right\} - \alpha \right),$$

which implies

$$0 \leq \liminf_{s \rightarrow \infty} \left(\mathbb{P}_{P_{n_s}} \left\{ \tilde{\mu}_{\alpha,n_s}^* \geq l'_{post} \tilde{\beta}_{n_s}(P_{n_s}) \mid \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post,n_s} \right\} - \alpha \right) \mathbb{P}_{P_{n_s}} \left\{ \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post,n_s} \right\}, \quad (29)$$

which yields a contradiction, since equation (28) implies that the above sequence is eventually bounded above by a negative constant.

Condition (i). Equation (28) implies

$$\liminf_{q \rightarrow \infty} \left(1 - \frac{\mathbb{P}_{P_{n_q}} \left\{ \tilde{\mu}_{\alpha,n_q}^* \geq l'_{post} \tilde{\beta}_{n_q}(P_{n_q}) \mid \tilde{\beta}_{n_q} \in \tilde{\mathcal{B}}_{post,n_q} \right\}}{\alpha} \right) \mathbb{P}_{P_{n_q}} \left\{ \tilde{\beta}_{n_q} \in \tilde{\mathcal{B}}_{post,n_q} \right\} = L' > \frac{\varepsilon}{\alpha}.$$

By properties of the limit inferior, given $\varepsilon/(2\alpha) > 0$, there exists q_0 such that

$$\left(1 - \frac{\mathbb{P}_{P_{n_q}} \left\{ \tilde{\mu}_{\alpha,n_q}^* \geq l'_{post} \tilde{\beta}_{n_q}(P_{n_q}) \mid \tilde{\beta}_{n_q} \in \tilde{\mathcal{B}}_{post,n_q} \right\}}{\alpha} \right) \mathbb{P}_{P_{n_q}} \left\{ \tilde{\beta}_{n_q} \in \tilde{\mathcal{B}}_{post,n_q} \right\} > L' - \frac{\varepsilon}{2\alpha}, \quad \forall q \geq q_0.$$

For such q , the term in parentheses on the LHS must be positive, and hence in $(0,1]$. Thus,

$$\mathbb{P}_{P_{n_q}} \left\{ \tilde{\beta}_{n_q} \in \tilde{\mathcal{B}}_{post, n_q} \right\} > L' - \frac{\varepsilon}{2\alpha} > \frac{\varepsilon}{2\alpha}, \quad \forall q \geq q_0.$$

The set $[\varepsilon/(2\alpha), 1]$ is compact and contains $\mathbb{P}_{P_{n_q}} \left\{ \tilde{\beta}_{n_q} \in \tilde{\mathcal{B}}_{post, n_q} \right\}$ for all $q \geq q_0$. Thus, there must exist a subsequence $\{n_r\} \subseteq \{n_q : q \geq q_0\}$ such that

$$\mathbb{P}_{P_{n_r}} \left\{ \tilde{\beta}_{n_r} \in \tilde{\mathcal{B}}_{post, n_r} \right\} \rightarrow p^* \in [\varepsilon/(2\alpha), 1].$$

Condition (ii). Consider $a = (a'_1, \dots, a'_K)' \in [-\infty, +\infty]^{\dim(c)}$. Under the standard order topology on $[-\infty, +\infty]$, the induced product topology on $[-\infty, +\infty]^{\dim(c)}$ is generated by the metric

$$d(a^1, a^2) = \max_j |\Phi(a_j^1) - \Phi(a_j^2)|, \quad a^1, a^2 \in [-\infty, \infty]^{\dim(c)},$$

where $z \mapsto \Phi(z)$ denotes the standard normal CDF with $\Phi(-\infty) = 0$ and $\Phi(+\infty) = 1$. The set $[-\infty, +\infty]^{\dim(c)}$ is compact under d and contains $A_{post} \tilde{\beta}_{n_r}(P_{n_r})$ for all r . Thus, there exists a further subsequence $\{n_t\} \subseteq \{n_r\}$ and limit a_{post}^* such that

$$d(A_{post} \tilde{\beta}_{n_t}(P_{n_t}), a_{post}^*) \rightarrow 0, \quad a_{post}^* \in [-\infty, +\infty]^{\dim(c)}.$$

In particular, $A_{post, k} \tilde{\beta}_{n_t}(P_{n_t}) \rightarrow a_{post, k}^* \in [-\infty, +\infty]^{\dim(c_k)}$ for all k .

Condition (iii). By assumption, $\Sigma(P) \in \mathcal{M}$ for all P . Thus, $\Sigma(P_{n_t}) \in \mathcal{M}$ for all t . By compactness of \mathcal{M} in the spectral norm, there then must exist a subsequence $\{n_u\} \subseteq \{n_t\}$ such that $\Sigma(P_{n_u}) \rightarrow \Sigma^* \in \mathcal{M}$.

Condition (iv). By assumption, we have $\|c_{post, k}(P)\| \leq \bar{\lambda}_c$, for all P and k . Thus, $c_{post}(P_{n_u}) \in [-\bar{\lambda}_c, \bar{\lambda}_c]^{\dim(c)}$ for all u . The set $[-\bar{\lambda}_c, \bar{\lambda}_c]^{\dim(c)}$ is compact, so there must exist a subsequence $\{n_s\} \subseteq \{n_u\}$ and limit c_{post}^* such that

$$c_{post}(P_{n_s}) \rightarrow c_{post}^* \in [-\bar{\lambda}_c, \bar{\lambda}_c]^{\dim(c)}.$$

In particular, $c_{post, k}(P_{n_s}) \rightarrow c_{post, k}^* \in [-\bar{\lambda}_c, \bar{\lambda}_c]^{\dim(c_k)}$ for all k .

Contradiction. We have obtained a subsequence $\{n_s\} \subseteq \{n_q\}$ that satisfies conditions (i)-(iv) of this lemma. By hypothesis, equation (26) holds for this subsequence:

$$\liminf_{s \rightarrow \infty} \mathbb{P}_{P_{n_s}} \left\{ \tilde{\mu}_{\alpha, n_s}^* \geq l'_{post} \tilde{\beta}_{n_s}(P_{n_s}) \mid \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\} \geq \alpha.$$

But this leads to a contradiction of equation (28), as we described under equation (29). ■

Representing the Coverage Event. To use Lemma 1 to prove Proposition 7, we first derive a representation of the coverage event $\{\tilde{\mu}_{\alpha,n}^* \geq l'_{post} \tilde{\beta}_n(P)\}$ in terms of appropriately normalized (centered and scaled) estimators. In what follows, recall that $\tilde{\beta}_n^* = \tilde{\beta}_n - \tilde{\beta}_n(P) = \sqrt{n}(\hat{\beta}_n - \beta(P))$. Moreover, define the normalized residual

$$\tilde{r}_n^* = \tilde{r}_n + \tilde{\gamma}_n l'_{post} \tilde{\beta}_n(P),$$

the normalized truncation quantities

$$\begin{aligned} \tilde{Z}_{post,k,n}^{*-} &= \max_{j \in \tilde{J}_{k,n}^-} \frac{(\tilde{C}_{post,k,n})_j - (A_{post,k} \tilde{r}_n^*)_j}{(A_{post,k} \tilde{\gamma}_n)_j} = \tilde{Z}_{post,k,n}^- - l'_{post} \tilde{\beta}_n(P), \\ \tilde{Z}_{post,k,n}^{*+} &= \min_{j \in \tilde{J}_{k,n}^+} \frac{(\tilde{C}_{post,k,n})_j - (A_{post,k} \tilde{r}_n^*)_j}{(A_{post,k} \tilde{\gamma}_n)_j} = \tilde{Z}_{post,k,n}^+ - l'_{post} \tilde{\beta}_n(P), \\ \tilde{Z}_{post,k,n}^{*0} &= \min_{j \in \tilde{J}_{k,n}^0} (\tilde{C}_{post,k,n})_j - (A_{post,k} \tilde{r}_n^*)_j = \tilde{Z}_{post,k,n}^0, \end{aligned}$$

and the normalized truncation sets

$$\tilde{\mathcal{Z}}_{post,n}^* = \bigcup_{k=1}^K \tilde{\mathcal{Z}}_{post,k,n}^*, \quad \tilde{\mathcal{Z}}_{post,k,n}^* = \begin{cases} [\tilde{Z}_{post,k,n}^{*-}, \tilde{Z}_{post,k,n}^{*+}], & \tilde{Z}_{post,k,n}^{*0} \geq 0 \\ \emptyset, & \tilde{Z}_{post,k,n}^{*0} < 0 \end{cases}.$$

Note that $\tilde{\mathcal{Z}}_{post,n}^* = \tilde{\mathcal{Z}}_{post,n} - l'_{post} \tilde{\beta}_n(P)$.

Lemma 2. *For all $n \geq 1$, we have*

$$\tilde{\mu}_{\alpha,n}^* \geq l'_{post} \tilde{\beta}_n(P) \iff 1 - \alpha \leq F_{TN}(l'_{post} \tilde{\beta}_n^*; 0, \tilde{\sigma}_{post,n}^2, \tilde{\mathcal{Z}}_{post,n}^*).$$

Proof of Lemma 2.

Noting that $\tilde{\mu}_{\alpha,n}^* = \sqrt{n} \hat{\mu}_{\alpha,n}^*$, we obtain

$$\begin{aligned} F_{TN}(l'_{post} \tilde{\beta}_n; \tilde{\mu}_{\alpha,n}^*, \tilde{\sigma}_{post,n}^2, \tilde{\mathcal{Z}}_{post,n}) &= \frac{\int_{-\infty}^{l'_{post} \tilde{\beta}_n} \phi\left(\frac{t - \tilde{\mu}_{\alpha,n}^*}{\tilde{\sigma}_{post,n}}\right) \mathbb{1}\{t \in \tilde{\mathcal{Z}}_{post,n}\} dt}{\int_{-\infty}^{\infty} \phi\left(\frac{t - \tilde{\mu}_{\alpha,n}^*}{\tilde{\sigma}_{post,n}}\right) \mathbb{1}\{t \in \tilde{\mathcal{Z}}_{post,n}\} dt} \\ &= \frac{\int_{-\infty}^{l'_{post} \sqrt{n} \hat{\beta}_n} \phi\left(\frac{t/\sqrt{n} - \hat{\mu}_{\alpha,n}^*}{\hat{\sigma}_{post,n}}\right) \mathbb{1}\{t/\sqrt{n} \in \hat{\mathcal{Z}}_{post,n}\} dt}{\int_{-\infty}^{\infty} \phi\left(\frac{t/\sqrt{n} - \hat{\mu}_{\alpha,n}^*}{\hat{\sigma}_{post,n}}\right) \mathbb{1}\{t/\sqrt{n} \in \hat{\mathcal{Z}}_{post,n}\} dt} \end{aligned}$$

$$\begin{aligned}
&= \frac{\int_{-\infty}^{l'_{post}\hat{\beta}_n} \phi\left(\frac{u - \hat{\mu}_{\alpha,n}^*}{\hat{\sigma}_{post,n}}\right) \mathbb{1}\{u \in \hat{\mathcal{Z}}_{post,n}\} du}{\int_{-\infty}^{\infty} \phi\left(\frac{u - \hat{\mu}_{\alpha,n}^*}{\hat{\sigma}_{post,n}}\right) \mathbb{1}\{u \in \hat{\mathcal{Z}}_{post,n}\} du} \\
&= F_{TN}(l'_{post}\hat{\beta}_n; \hat{\mu}_{\alpha,n}^*, \hat{\sigma}_{post,n}^2, \hat{\mathcal{Z}}_{post,n}),
\end{aligned}$$

where the third equality follows from substituting $u = t/\sqrt{n}$ and $du = dt/\sqrt{n}$. Since $\hat{\mu}_{\alpha,n}^*$ uniquely solves $F_{TN}(l'_{post}\hat{\beta}_n; \hat{\mu}_{\alpha,n}^*, \hat{\sigma}_{post,n}^2, \hat{\mathcal{Z}}_{post,n}) = 1 - \alpha$ given $\hat{\beta}_n \in \hat{\mathcal{B}}_{post,n}$, then given $\tilde{\beta}_n \in \tilde{\mathcal{B}}_{post,n}$ we have that $\tilde{\mu}_{\alpha,n}^*$ uniquely solves

$$F_{TN}(l'_{post}\tilde{\beta}_n; \tilde{\mu}_{\alpha,n}^*, \tilde{\sigma}_{post,n}^2, \tilde{\mathcal{Z}}_{post,n}) = 1 - \alpha.$$

Since the truncated normal CDF is strictly decreasing in μ , the above implies

$$\tilde{\mu}_{\alpha,n}^* \geq l'_{post}\tilde{\beta}_n(P) \iff 1 - \alpha \leq F_{TN}(l'_{post}\tilde{\beta}_n; l'_{post}\tilde{\beta}_n(P), \tilde{\sigma}_{post,n}^2, \tilde{\mathcal{Z}}_{post,n}).$$

Observe that

$$\begin{aligned}
F_{TN}(l'_{post}\tilde{\beta}_n; l'_{post}\tilde{\beta}_n(P), \tilde{\sigma}_{post,n}^2, \tilde{\mathcal{Z}}_{post,n}) &= \frac{\int_{-\infty}^{l'_{post}\tilde{\beta}_n} \phi\left(\frac{t - l'_{post}\tilde{\beta}_n(P)}{\tilde{\sigma}_{post,n}}\right) \mathbb{1}\{t \in \tilde{\mathcal{Z}}_{post,n}\} dt}{\int_{-\infty}^{\infty} \phi\left(\frac{t - l'_{post}\tilde{\beta}_n(P)}{\tilde{\sigma}_{post,n}}\right) \mathbb{1}\{t \in \tilde{\mathcal{Z}}_{post,n}\} dt} \\
&= \frac{\int_{-\infty}^{l'_{post}\tilde{\beta}_n^*} \phi\left(\frac{u}{\tilde{\sigma}_{post,n}}\right) \mathbb{1}\{u \in \tilde{\mathcal{Z}}_{post,n} - l'_{post}\tilde{\beta}_n(P)\} du}{\int_{-\infty}^{\infty} \phi\left(\frac{u}{\tilde{\sigma}_{post,n}}\right) \mathbb{1}\{u \in \tilde{\mathcal{Z}}_{post,n} - l'_{post}\tilde{\beta}_n(P)\} du} \\
&= F_{TN}(l'_{post}\tilde{\beta}_n^*; 0, \tilde{\sigma}_{post,n}^2, \tilde{\mathcal{Z}}_{post,n} - l'_{post}\tilde{\beta}_n(P)),
\end{aligned}$$

where the second equality follows from substituting $u = t - l'_{post}\tilde{\beta}_n(P)$ and $du = dt$. Thus,

$$\tilde{\mu}_{\alpha,n}^* \geq l'_{post}\tilde{\beta}_n(P) \iff 1 - \alpha \leq F_{TN}(l'_{post}\tilde{\beta}_n^*; 0, \tilde{\sigma}_{post,n}^2, \tilde{\mathcal{Z}}_{post,n} - l'_{post}\tilde{\beta}_n(P)),$$

as desired. ■

Lemma 2 provides a representation of the coverage event $\{\tilde{\mu}_{\alpha,n}^* \geq l'_{post}\tilde{\beta}_n(P)\}$ in terms of the quantity $F_{TN}(l'_{post}\tilde{\beta}_n^*; 0, \tilde{\sigma}_{post,n}^2, \tilde{\mathcal{Z}}_{post,n}^*)$, which depends on normalized variables

$$(\tilde{\beta}_n^*, \tilde{\Sigma}_n, \tilde{\mathcal{Z}}_{post,n}^*) = (\sqrt{n}(\hat{\beta}_n - \beta(P)), n\hat{\Sigma}_n, \sqrt{n}(\hat{\mathcal{Z}}_{post,n} - l'_{post}\beta(P))).$$

We now characterize the limiting/asymptotic behavior of the RHS as $n \rightarrow \infty$. This will allow us to leverage Lemma 1 to prove Proposition 7.

Limit Problem. For subsequences $\{n_s\} \subseteq \{n\}$ and $\{P_{n_s}\} \in \times_{s=1}^{\infty} \mathcal{P}_{n_s}$ satisfying the conditions of Lemma 1, we show that the corresponding sequence of finite sample problems, where we observe samples of size n_s from distributions P_{n_s} and employ conditional inference procedures, is well approximated by a limit problem where we observe a single normal vector with known covariance matrix and employ the analogous conditional inference procedures. In considering this limit problem, we will be able to understand the asymptotic behavior of

$$(\tilde{\beta}_{n_s}^*, \tilde{\Sigma}_{n_s}, \tilde{\mathcal{Z}}_{post, n_s}^*) = (\sqrt{n_s}(\hat{\beta}_{n_s} - \beta(P_{n_s})), n_s \hat{\Sigma}_{n_s}, \sqrt{n_s}(\hat{\mathcal{Z}}_{post, n_s} - l'_{post} \beta(P_{n_s}))),$$

and hence the asymptotic behavior of $F_{TN}(l'_{post} \tilde{\beta}_{n_s}^*; 0, \tilde{\sigma}_{post, n_s}^2, \tilde{\mathcal{Z}}_{post, n_s}^*)$.

To derive the limit problem, note that under Assumptions 1-3, and for subsequences that satisfy the conditions of Lemma 1, we obtain the following convergence results, which we will use in the proof of Proposition 7:

$$\begin{aligned} \tilde{\beta}_{n_s}^* &\xrightarrow{d} N(0, \Sigma^*), & A_{post, k} \tilde{\beta}_{n_s}^*(P_{n_s}) &\rightarrow a_{post, k}^*, \\ \tilde{\Sigma}_{n_s} &\xrightarrow{p} \Sigma^*, & \tilde{c}_{post, k, n_s} &\xrightarrow{p} c_{post, k}^*, \\ \Sigma(P_{n_s}) &\rightarrow \Sigma^*, & c_{post, k}(P_{n_s}) &\rightarrow c_{post, k}^*, \end{aligned}$$

where $\Sigma^* \in \mathcal{M}$, $a_{post, k}^* \in [-\infty, +\infty]^{\dim(c_k)}$, $c_{post, k}^* \in [-\bar{\lambda}_c, \bar{\lambda}_c]^{\dim(c_k)}$, and $k = 1, \dots, K$. Thus, stacking the vectors and matrices over k , we obtain

$$\begin{pmatrix} \tilde{\beta}_{n_s}^* \\ \tilde{\Sigma}_{n_s} \\ A_{post} \tilde{\beta}_{n_s}^*(P_{n_s}) \\ \tilde{c}_{post, n_s} \end{pmatrix} \xrightarrow{d} \begin{pmatrix} \xi^* \\ \Sigma^* \\ a_{post}^* \\ c_{post}^* \end{pmatrix}, \quad \xi^* \sim N(0, \Sigma^*). \quad (30)$$

Letting $(A_{post, k})_j$ denote rows of $A_{post, k}$, the continuous mapping theorem (CMT) implies

$$\begin{aligned} \max_k \mathbb{1} \{A_{post, k} \tilde{\beta}_{n_s} \leq \tilde{c}_{post, k, n_s}\} &= \max_k \mathbb{1} \{A_{post, k} \tilde{\beta}_{n_s}^* \leq \tilde{c}_{post, k, n_s} - A_{post, k} \tilde{\beta}_{n_s}(P_{n_s})\} \\ &= \max_k \min_j \mathbb{1} \{(A_{post, k})_j \tilde{\beta}_{n_s}^* \leq (\tilde{c}_{post, k, n_s})_j - (A_{post, k})_j \tilde{\beta}_{n_s}(P_{n_s})\} \\ &\xrightarrow{d} \max_k \min_j \mathbb{1} \{(A_{post, k})_j \xi^* \leq (c_{post, k}^*)_j - (a_{post, k}^*)_j\} \\ &= \max_k \mathbb{1} \{A_{post, k} \xi^* \leq c_{post, k}^* - a_{post, k}^*\}, \end{aligned} \quad (31)$$

where we use the continuity of $(A_{post,k})_j \xi^* \sim N(0, (A_{post,k})_j \Sigma^* (A_{post,k})'_j)$, which we obtain from $(A_{post,k})_j \neq 0$ for all (k, j) . Consider the union of polyhedra

$$\mathcal{B}_{post}^* = \bigcup_{k=1}^K \{\xi^* : A_{post,k} \xi^* \leq c_{post,k}^* - a_{post,k}^*\}.$$

By equation (31) and the conditions of Lemma 1, we obtain

$$\begin{aligned} \mathbb{P}_{P_{n_s}} \left\{ \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\} &= \mathbb{E}_{P_{n_s}} \left[\max_k \mathbb{1} \left\{ A_{post,k} \tilde{\beta}_{n_s} \leq \tilde{c}_{post,k, n_s} \right\} \right] \\ &\rightarrow \mathbb{E} \left[\max_k \mathbb{1} \left\{ A_{post,k} \xi^* \leq c_{post,k}^* - a_{post,k}^* \right\} \right] \\ &= \mathbb{P} \left\{ \xi^* \in \mathcal{B}_{post}^* \right\} \\ &= p^* > 0. \end{aligned} \tag{32}$$

For the limit problem residual step, define

$$r^* = \xi^* - \gamma^* l'_{post} \xi^*, \quad \gamma^* = \frac{\Sigma^* l_{post}}{\sigma_{post}^{*2}}, \quad \sigma_{post}^{*2} = l'_{post} \Sigma^* l_{post}.$$

For conditioning, define truncation quantities

$$\begin{aligned} Z_{post,k}^{*-} &= \max_{j \in J_k^-} \frac{(c_{post,k}^* - a_{post,k}^*)_j - (A_{post,k} r^*)_j}{(A_{post,k} \gamma^*)_j}, & J_k^- &= \{j : (A_{post,k} \gamma^*)_j < 0\}, \\ Z_{post,k}^{*+} &= \min_{j \in J_k^+} \frac{(c_{post,k}^* - a_{post,k}^*)_j - (A_{post,k} r^*)_j}{(A_{post,k} \gamma^*)_j}, & J_k^+ &= \{j : (A_{post,k} \gamma^*)_j > 0\}, \\ Z_{post,k}^{*0} &= \min_{j \in J_k^0} (c_{post,k}^* - a_{post,k}^*)_j - (A_{post,k} r^*)_j, & J_k^0 &= \{j : (A_{post,k} \gamma^*)_j = 0\}, \end{aligned}$$

such that

$$\mathcal{Z}_{post}^* = \bigcup_{k=1}^K \mathcal{Z}_{post,k}^*, \quad \mathcal{Z}_{post,k}^* = \begin{cases} [Z_{post,k}^{*-}, Z_{post,k}^{*+}], & Z_{post,k}^{*0} \geq 0, \\ \emptyset, & Z_{post,k}^{*0} < 0. \end{cases}$$

In summary, the limit problem consists of a normal draw $\xi^* \sim N(0, \Sigma^*)$ with (unknown) mean zero and (known) covariance matrix $\Sigma^* \in \mathcal{M}$, and considers inference for post estimates $l'_{post} \xi^*$ conditional on ξ^* falling into a union of polyhedra:

$$\xi^* \in \mathcal{B}_{post}^*, \quad \mathcal{B}_{post}^* = \bigcup_{k=1}^K \{\xi^* : A_{post,k} \xi^* \leq c_{post,k}^* - a_{post,k}^*\},$$

where the conditioning event $\{\xi^* \in \mathcal{B}_{post}^*\}$ is equivalent to event $\{l'_{post}\xi^* \in \mathcal{Z}_{post}^*\}$. This conditioning event has positive probability, as noted in (32).

Representing Inclusion in Truncation Set. As a last step before proving the main result, it is useful to represent membership $\mathbb{1}\{t \in \mathcal{Z}_{post}^*\}$ in the truncation set for a given $t \in \mathbb{R}$ in terms of the limiting variables $(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*)$. To do so, note that if $t \in \mathcal{Z}_{post,k}^*$ for some k , then this is equivalent to (i) $Z_{post,k}^{*0} \geq 0$ and (ii)

$$\max_{j \in J_k^-} \frac{(c_{post,k}^* - a_{post,k}^*)_j - (A_{post,k} r^*)_j}{(A_{post,k} \gamma^*)_j} \leq t \leq \min_{j \in J_k^+} \frac{(c_{post,k}^* - a_{post,k}^*)_j - (A_{post,k} r^*)_j}{(A_{post,k} \gamma^*)_j}.$$

Condition (i) is equivalent to

$$(A_{post,k} \gamma^*)_j t + (A_{post,k} r^*)_j \leq (c_{post,k}^* - a_{post,k}^*)_j, \quad \forall j \in J_k^0.$$

Condition (ii) is equivalent to the combination of these two statements:

$$\begin{aligned} \frac{(c_{post,k}^* - a_{post,k}^*)_j - (A_{post,k} r^*)_j}{(A_{post,k} \gamma^*)_j} &\leq t, \quad \forall j \in J_k^-, \\ \frac{(c_{post,k}^* - a_{post,k}^*)_j - (A_{post,k} r^*)_j}{(A_{post,k} \gamma^*)_j} &\geq t, \quad \forall j \in J_k^+. \end{aligned}$$

Together, the equivalent expressions for conditions (i) and (ii) show that, if $t \in \mathcal{Z}_{post,k}^*$ for some k , then we equivalently have $(A_{post,k} \gamma^*)_j t + (A_{post,k} r^*)_j \leq (c_{post,k}^* - a_{post,k}^*)_j$ for all j . In other words,

$$\mathbb{1}\{t \in \mathcal{Z}_{post,k}^*\} = \mathbb{1}\{A_{post,k}(\gamma^* t + r^*) \leq c_{post,k}^* - a_{post,k}^*\}.$$

Since $t \in \mathcal{Z}_{post}^*$ is equivalent to the existence of k such that $t \in \mathcal{Z}_{post,k}^*$, we obtain

$$\mathbb{1}\{t \in \mathcal{Z}_{post}^*\} = \max_k \mathbb{1}\{A_{post,k}(\gamma^* t + r^*) \leq c_{post,k}^* - a_{post,k}^*\}. \quad (33)$$

Proving the Main Result. We are now prepared to prove Proposition 7.

Proof of Proposition 7.

Under Assumptions 1-3, to prove that

$$\lim_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_n} \left| \mathbb{P}_P \left\{ \hat{\mu}_{\alpha,n}^* \geq l'_{post} \beta(P) \mid \hat{\beta}_n \in \hat{\mathcal{B}}_{post,n} \right\} - \alpha \right| \mathbb{P}_P \left\{ \hat{\beta}_n \in \hat{\mathcal{B}}_{post,n} \right\} = 0,$$

it suffices to show that for all subsequences $\{n_s\} \subseteq \{n\}$ and $\{P_{n_s}\} \in \times_{s=1}^{\infty} \mathcal{P}_{n_s}$ that satisfy the

conditions of Lemma 1, we have

$$\lim_{s \rightarrow \infty} \mathbb{P}_{P_{n_s}} \left\{ \tilde{\mu}_{\alpha, n_s}^* \geq l'_{post} \tilde{\beta}_{n_s} (P_{n_s}) \mid \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\} = \alpha.$$

By Lemma 2, the above is equivalent to

$$\lim_{s \rightarrow \infty} \mathbb{P}_{P_{n_s}} \left\{ F_{TN}(l'_{post} \tilde{\beta}_{n_s}^*; 0, \tilde{\sigma}_{post, n_s}^2, \tilde{\mathcal{Z}}_{post, n_s}^*) \geq 1 - \alpha \mid \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\} = \alpha.$$

To verify this convergence, note that (i) $F_{TN}(l'_{post} \xi^*; 0, \sigma_{post}^{*2}, \mathcal{Z}_{post}^*) \mid \xi^* \in \mathcal{B}_{post}^* \sim U(0, 1)$, as noted in the main text and (ii) $\mathbb{P}\{\xi^* \in \mathcal{B}_{post}^*\} > 0$, as per equation (32). Thus, if we can show that

$$\begin{aligned} & \frac{\mathbb{E}_{P_{n_s}} \left[\mathbb{1} \left\{ F_{TN}(l'_{post} \tilde{\beta}_{n_s}^*; 0, \tilde{\sigma}_{post, n_s}^2, \tilde{\mathcal{Z}}_{post, n_s}^*) \geq 1 - \alpha \right\} \mathbb{1} \left\{ \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\} \right]}{\mathbb{P}_{P_{n_s}} \left\{ \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\}} \\ & \rightarrow \frac{\mathbb{E} \left[\mathbb{1} \left\{ F_{TN}(l'_{post} \xi^*; 0, \sigma_{post}^{*2}, \mathcal{Z}_{post}^*) \geq 1 - \alpha \right\} \mathbb{1} \left\{ \xi^* \in \mathcal{B}_{post}^* \right\} \right]}{\mathbb{P} \left\{ \xi^* \in \mathcal{B}_{post}^* \right\}} \\ & = \mathbb{P} \left\{ F_{TN}(l'_{post} \xi^*; 0, \sigma_{post}^{*2}, \mathcal{Z}_{post}^*) \geq 1 - \alpha \mid \xi^* \in \mathcal{B}_{post}^* \right\} = \alpha, \end{aligned}$$

then we will have verified the desired convergence.

Note that $\mathbb{P}_{P_{n_s}} \left\{ \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\} \rightarrow \mathbb{P} \left\{ \xi^* \in \mathcal{B}_{post}^* \right\} > 0$ by equation (32). Thus, it suffices to show that the LHS numerator converges to the RHS numerator.

To establish convergence of the numerator, it suffices to show

$$\begin{aligned} & \mathbb{1} \left\{ F_{TN}(l'_{post} \tilde{\beta}_{n_s}^*; 0, \tilde{\sigma}_{post, n_s}^2, \tilde{\mathcal{Z}}_{post, n_s}^*) \geq 1 - \alpha \right\} \mathbb{1} \left\{ \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\} \\ & \xrightarrow{d} \mathbb{1} \left\{ F_{TN}(l'_{post} \xi^*; 0, \sigma_{post}^{*2}, \mathcal{Z}_{post}^*) \geq 1 - \alpha \right\} \mathbb{1} \left\{ \xi^* \in \mathcal{B}_{post}^* \right\}. \end{aligned} \tag{34}$$

To this end, note that $\mathbb{P} \left\{ F_{TN}(l'_{post} \xi^*; 0, \sigma_{post}^{*2}, \mathcal{Z}_{post}^*) \geq 1 - \alpha \mid \xi^* \in \mathcal{B}_{post}^* \right\} = 1$, since

$$F_{TN}(l'_{post} \xi^*; 0, \sigma_{post}^{*2}, \mathcal{Z}_{post}^*) \mid \xi^* \in \mathcal{B}_{post}^* \sim U(0, 1).$$

Thus, if we can show that

$$\begin{aligned} & (F_{TN}(l'_{post} \tilde{\beta}_{n_s}^*; 0, \tilde{\sigma}_{post, n_s}^2, \tilde{\mathcal{Z}}_{post, n_s}^*) \mathbb{1} \left\{ \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\}, \mathbb{1} \left\{ \tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post, n_s} \right\}) \\ & \xrightarrow{d} (F_{TN}(l'_{post} \xi^*; 0, \sigma_{post}^{*2}, \mathcal{Z}_{post}^*) \mathbb{1} \left\{ \xi^* \in \mathcal{B}_{post}^* \right\}, \mathbb{1} \left\{ \xi^* \in \mathcal{B}_{post}^* \right\}), \end{aligned} \tag{35}$$

then CMT will yield the convergence in (34).

To this end, recall the convergence in (30) yields

$$\begin{pmatrix} \tilde{\beta}_{n_s}^* \\ \tilde{\Sigma}_{n_s} \\ A_{post} \tilde{\beta}_{n_s}(P_{n_s}) \\ \tilde{c}_{post, n_s} \end{pmatrix} \xrightarrow{d} \begin{pmatrix} \xi^* \\ \Sigma^* \\ a_{post}^* \\ c_{post}^* \end{pmatrix}. \quad (36)$$

Let $\mathcal{V} = \mathbb{R}^{\dim(\beta)} \times \mathcal{M} \times [-\infty, +\infty]^{\dim(c)} \times [-\bar{\lambda}_c, \bar{\lambda}_c]^{\dim(c)}$. Given some function $g: \mathcal{V} \rightarrow \mathbb{R}^{\dim(g)}$, let $\mathcal{C}(g) \subseteq \mathcal{V}$ denote the set of points at which g is continuous. If we can represent the RHS of (35) as a function $g(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*)$ for which the set of continuity points $\mathcal{C}(g)$ has probability one under $\xi^* \sim N(0, \Sigma^*)$, $\Sigma^* \in \mathcal{M}$, $a_{post}^* \in [-\infty, +\infty]^{\dim(c)}$, and $c_{post}^* \in [-\bar{\lambda}_c, \bar{\lambda}_c]^{\dim(c)}$ in that

$$\mathbb{P}\{(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*) \in \mathcal{C}(g)\} = 1,$$

then the convergence in (36) combined with CMT yields the desired convergence in (35).

To work towards a continuous representation, we first expand the truncated CDF term. By equation (33), we have

$$\begin{aligned} F_{TN}(l'_{post} \xi^*; 0, \sigma_{post}^{*2}, \mathcal{Z}_{post}^*) \mathbb{1}\{\xi^* \in \mathcal{B}_{post}^*\} &= \frac{\int_{-\infty}^{l'_{post} \xi^*} \phi\left(\frac{t}{\sigma_{post}^*}\right) \mathbb{1}\{t \in \mathcal{Z}_{post}^*\} dt}{\int_{-\infty}^{\infty} \phi\left(\frac{t}{\sigma_{post}^*}\right) \mathbb{1}\{t \in \mathcal{Z}_{post}^*\} dt} \mathbb{1}\{\xi^* \in \mathcal{B}_{post}^*\} \\ &= \frac{\int_{-\infty}^{\infty} \mathbb{1}\{t \leq l'_{post} \xi^*\} \phi\left(\frac{t}{\sigma_{post}^*}\right) \max_k \mathbb{1}\{A_{post, k}(\gamma^* t + r^*) \leq c_{post, k}^* - a_{post, k}^*\} dt}{\int_{-\infty}^{\infty} \phi\left(\frac{t}{\sigma_{post}^*}\right) \max_k \mathbb{1}\{A_{post, k}(\gamma^* t + r^*) \leq c_{post, k}^* - a_{post, k}^*\} dt} \mathbb{1}\{\xi^* \in \mathcal{B}_{post}^*\}. \end{aligned}$$

We now derive continuous representations for the numerator and denominator. To this end, we consider sequences $(\xi^m, \Sigma^m, a_{post}^m, c_{post}^m) \in \mathcal{V}$ for which

$$(\xi^m, \Sigma^m, a_{post}^m, c_{post}^m) \rightarrow (\xi^0, \Sigma^0, a_{post}^0, c_{post}^0),$$

where $(\xi^0, \Sigma^0, a_{post}^0, c_{post}^0)$ is such that the corresponding set $\mathbb{R} \setminus \mathcal{T}_{g_1}^0$ given below is finite.

For the numerator, consider the function

$$\begin{aligned} g_1(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*) \\ = \int_{-\infty}^{\infty} \mathbb{1}\{t \leq l'_{post} \xi^*\} \phi\left(\frac{t}{\sigma_{post}^*}\right) \max_k \min_j \mathbb{1}\{(A_{post, k})_j(\gamma^* t + r^*) \leq (c_{post, k}^*)_j - (a_{post, k}^*)_j\} dt. \end{aligned}$$

For $t \in \mathcal{T}_{g_1}^0 = \{t \in \mathbb{R} : t \neq l'_{post} \xi^0, (A_{post,k})_j \gamma^0 t \neq (c_{post,k}^0)_j - (a_{post,k}^0)_j - (A_{post,k})_j r^0, \forall k, j\}$, we have

$$\begin{aligned} & \mathbb{1}\{t \leq l'_{post} \xi^m\} \phi\left(\frac{t}{\sigma_{post}^m}\right) \max_k \min_j \mathbb{1}\{(A_{post,k})_j (\gamma^m t + r^m) \leq (c_{post,k}^m)_j - (a_{post,k}^m)_j\} \\ & \rightarrow \mathbb{1}\{t \leq l'_{post} \xi^0\} \phi\left(\frac{t}{\sigma_{post}^0}\right) \max_k \min_j \mathbb{1}\{(A_{post,k})_j (\gamma^0 t + r^0) \leq (c_{post,k}^0)_j - (a_{post,k}^0)_j\}, \end{aligned}$$

since the latter is continuous at $(\xi^0, \Sigma^0, a_{post}^0, c_{post}^0)$. Since $\mathbb{R} \setminus \mathcal{T}_{g_1}^0$ is finite, its Lebesgue measure is zero; moreover, under the eigenvalue bounds for \mathcal{M} , the terms $\phi(t/\sigma_{post}^m)$ are bounded above by a finite constant, uniformly over (t, m) . Therefore, dominated convergence theorem implies

$$\begin{aligned} & g_1(\xi^m, \Sigma^m, a_{post}^m, c_{post}^m) \\ &= \int_{-\infty}^{\infty} \mathbb{1}\{t \leq l'_{post} \xi^m\} \phi\left(\frac{t}{\sigma_{post}^m}\right) \max_k \min_j \mathbb{1}\{(A_{post,k})_j (\gamma^m t + r^m) \leq (c_{post,k}^m)_j - (a_{post,k}^m)_j\} dt \\ & \rightarrow \int_{-\infty}^{\infty} \mathbb{1}\{t \leq l'_{post} \xi^0\} \phi\left(\frac{t}{\sigma_{post}^0}\right) \max_k \min_j \mathbb{1}\{(A_{post,k})_j (\gamma^0 t + r^0) \leq (c_{post,k}^0)_j - (a_{post,k}^0)_j\} dt \\ &= \int_{-\infty}^{\infty} \mathbb{1}\{t \leq l'_{post} \xi^0\} \phi\left(\frac{t}{\sigma_{post}^0}\right) \max_k \mathbb{1}\{A_{post,k} (\gamma^0 t + r^0) \leq c_{post,k}^0 - a_{post,k}^0\} dt \\ &= g_1(\xi^0, \Sigma^0, a_{post}^0, c_{post}^0). \end{aligned}$$

The above argument holds for any $(\xi^0, \Sigma^0, a_{post}^0, c_{post}^0)$ where $\mathbb{R} \setminus \mathcal{T}_{g_1}^0$ is finite. This finiteness holds whenever

$$(c_{post,k}^0 - a_{post,k}^0)_j - (A_{post,k})_j r^0 \neq 0, \quad \forall j : (A_{post,k} \gamma^0)_j = 0, \quad \forall k.$$

Thus the set of continuity points for $g_1(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*)$ contains

$$\begin{aligned} & \{(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*) : (A_{post,k})_j r^* \neq (c_{post,k}^*)_j - (a_{post,k}^*)_j, \forall j \in J_k^0, \forall k\} \\ &= \{(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*) : (A_{post,k})_j \xi^* \neq (c_{post,k}^*)_j - (a_{post,k}^*)_j, \forall j \in J_k^0, \forall k\} \\ &\subseteq \mathcal{C}(g_1). \end{aligned}$$

For the denominator, analogous arguments show that the set of continuity points for

$$\begin{aligned} & g_2(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*) \\ &= \int_{-\infty}^{\infty} \phi\left(\frac{t}{\sigma_{post}^*}\right) \max_k \min_j \mathbb{1}\{(A_{post,k})_j (\gamma^* t + r^*) \leq (c_{post,k}^*)_j - (a_{post,k}^*)_j\} dt \end{aligned}$$

also contains

$$\{(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*) : (A_{post,k})_j \xi^* \neq (c_{post,k}^*)_j - (a_{post,k}^*)_j, \forall j \in J_k^0, \forall k\} \subseteq \mathcal{C}(g_2).$$

To use these results, notice $\mathbb{1}\{\xi^* \in \mathcal{B}_{post}^*\} = \max_k \min_j \mathbb{1}\{(A_{post,k})_j \xi^* \leq (c_{post,k}^*)_j - (a_{post,k}^*)_j\}$, and consider that

$$g_3(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*) = \max_k \min_j \mathbb{1}\{(A_{post,k})_j \xi^* \leq (c_{post,k}^*)_j - (a_{post,k}^*)_j\}$$

is continuous over the set

$$\{(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*) : (A_{post,k})_j \xi^* \neq (c_{post,k}^*)_j - (a_{post,k}^*)_j, \forall k, j\} \subseteq \mathcal{C}(g_3).$$

Under $\xi^* \in \mathcal{B}_{post}^*$, we have $0 < g_1 \leq g_2$ so that

$$\begin{aligned} & (F_{TN}(l'_{post} \xi^*, 0, \sigma_{post}^{*2}, \mathcal{Z}_{post}^*) \mathbb{1}\{\xi^* \in \mathcal{B}_{post}^*\}, \mathbb{1}\{\xi^* \in \mathcal{B}_{post}^*\}) \\ &= \left(\frac{g_1(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*)}{g_2(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*)} \mathbb{1}\{\xi^* \in \mathcal{B}_{post}^*\}, \mathbb{1}\{\xi^* \in \mathcal{B}_{post}^*\} \right) \\ &= \left(\frac{g_1(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*)}{g_2(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*)} g_3(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*), g_3(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*) \right) \\ &= g(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*). \end{aligned}$$

By the above analysis, the set of continuity points for g contains

$$\mathcal{C}(g_1) \cap \mathcal{C}(g_2) \cap \mathcal{C}(g_3) = \{(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*) : (A_{post,k})_j \xi^* \neq (c_{post,k}^*)_j - (a_{post,k}^*)_j, \forall k, j\} \subseteq \mathcal{C}(g).$$

Thus,

$$\begin{aligned} & \mathbb{P}\{(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*) \in \mathcal{C}(g)\} \\ & \geq \mathbb{P}\{(A_{post,k})_j \xi^* \neq (c_{post,k}^*)_j - (a_{post,k}^*)_j, \forall k, j\} \\ & = 1 - \mathbb{P}\{\exists(k, j) : (A_{post,k})_j \xi^* = (c_{post,k}^*)_j - (a_{post,k}^*)_j\} \\ & \geq 1 - \sum_{(k,j)} \mathbb{P}\{(A_{post,k})_j \xi^* = (c_{post,k}^*)_j - (a_{post,k}^*)_j\} \\ & = 1, \end{aligned}$$

where the last equality follows from continuity of $(A_{post,k})_j \xi^* \sim N(0, (A_{post,k})_j \Sigma^* (A_{post,k})'_j)$, since

$(A_{post,k})_j \neq 0$ for all (k,j) . To conclude, CMT combined with (36) yields the desired convergence in (35):

$$\begin{aligned}
& (F_{TN}(l'_{post}\tilde{\beta}_{n_s}^*; 0, \tilde{\sigma}_{post,n_s}^2, \tilde{\mathcal{Z}}_{post,n_s}^*) \mathbb{1}\{\tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post,n_s}\}, \mathbb{1}\{\tilde{\beta}_{n_s} \in \tilde{\mathcal{B}}_{post,n_s}\}) \\
&= g(\tilde{\beta}_{n_s}^*, \tilde{\Sigma}_{n_s}, A_{post}\tilde{\beta}_{n_s}(P_{n_s}), \tilde{c}_{post,n_s}) \\
&\xrightarrow{d} g(\xi^*, \Sigma^*, a_{post}^*, c_{post}^*) \\
&= (F_{TN}(l'_{post}\xi^*; 0, \sigma_{post}^{*2}, \mathcal{Z}_{post}^*) \mathbb{1}\{\xi^* \in \mathcal{B}_{post}^*\}, \mathbb{1}\{\xi^* \in \mathcal{B}_{post}^*\}).
\end{aligned}$$

■