



A Heisenberg-esque Uncertainty Principle for Simultaneous (Machine) Learning and Error Assessment?

Xiao-Li Meng

Department of Statistics, Harvard University

Received: 28 September 2024; Revised: 31 November 2024; Accepted: 02 December 2024

Abstract

A highly cited and inspiring article by Bates *et al.* (2024) demonstrates that the prediction errors estimated through cross-validation, Bootstrap or Mallorw's C_P can all be independent of the actual prediction errors. This essay hypothesizes that these occurrences signify a broader, Heisenberg-like uncertainty principle for learning: optimizing learning and assessing actual errors using the same data are fundamentally at odds. Only suboptimal learning preserves untapped information for actual error assessments, and vice versa, reinforcing the 'no free lunch' principle. To substantiate this intuition, a Cramér-Rao-style lower bound is established under the squared loss, which shows that the relative regret in learning is bounded below by the square of the correlation between any unbiased error assessor and the actual learning error. Readers are invited to explore generalizations, develop variations, or even uncover genuine 'free lunches.' The connection with the Heisenberg uncertainty principle is more than metaphorical, because both share an essence of the Cramér-Rao inequality: marginal variations cannot manifest individually to arbitrary degrees when their underlying co-variation is constrained, whether the co-variation is about individual states or their generating mechanisms, as in the quantum realm. A practical takeaway of such a learning principle is that it may be prudent to reserve some information specifically for error assessment rather than pursue full optimization in learning, particularly when intentional randomness is introduced to mitigate overfitting.

Key words: C. R. Rao; Cramér-Rao bound; Cross validation; Epistemology; Heisenberg uncertainty principle; Machine learning; Quantum mechanics; Uniformly minimum variance unbiased estimator.

AMS Subject Classifications: 62K05, 05B05

1. A Rao-esque apology and a quantum-leap excuse

Many of the advances in statistics and machine learning are about using data as efficiently and reliably as possible to achieve a host of learning objectives, such as inference, prediction, classification, *etc.* Being statistically efficient typically means to optimize over some criterion that amounts to minimizing learning errors based on the data at hand, whether

in a brute-force fashion, such as minimizing a χ^2 distance or adopting the L^2 -loss directly on the target of learning, or through deeper principles, *e.g.*, by maximizing a likelihood function or a posterior density. Since the actual learning errors themselves cannot be known without an external benchmark, we seek clever and reliable ways to assess them, whether for training machine learning algorithms, constructing confidence intervals, or checking Bayesian models.

Naturally, we wish to be able to optimally use our data for both purposes: to most efficiently learn whatever we can learn, and to most reliably assess the errors in whatever we cannot learn. However, since any information on the actual learning error can be used to improve the learning itself, we should be mindful that optimizing one endeavor comes at the expense of the other. To emphasize this no-free lunch principle, this essay first revisits seemingly quaint examples and classical results to remind ourselves that this principle has been in action for as long as statistical inference exists. However, such an issue has not received much emphasis apparently because principled statistical methods, such as likelihood or Bayesian methods, automatically prioritize optimal learning over error assessment.

Yet time has changed. Machine learning and other pattern-seeking methods require much intuition and judgment to tune well, when their theoretical guiding principles are not well developed or digested. Substituting—not merely supplementing—virtual trials and errors for sapient contemplation and introspection is becoming increasingly habitual, making us more vulnerable to wishful thinking, misinformed intuitions, and misguided common sense. To better prepare students and newcomers to our progressively empiricism-slanted culture of learning, this essay then recasts a classical result regarding UMVUE to the broader class of problems of unbiased learning, and establishes a mathematical inequality that captures the aforementioned Heisenberg-esque uncertainty principle for simultaneous learning and error assessment under the squared loss.

This inequality is a low-hanging fruit in establishing a general theory for understanding the competing nature between optimal learning and actual error assessing. Nevertheless, it can help us anticipate and better appreciate further results such as those obtained in Bates *et al.* (2024), which show that the error estimates from cross validation and other popular methods can be independent of actual learning error. The uncertainty principle tells us that this should not come as a surprise. Rather, the independence is an indication that the corresponding learning is optimal in some sense.

Since this essay was prepared for this special issue in memory of Professor C. R. Rao, it seems fitting to quote Rao (1962), a discussion article presented¹ to the Royal Statistical Society in England (RSS):

“While thanking the Royal Statistical Society for giving me an opportunity to read a paper at one of its meetings, I must apologize for choosing a subject which may appear somewhat classical. But I hope this small attempt intended to state in precise terms what can be claimed about m.l. estimates, in large samples, will at least throw some light on current controversies.”

Rao (1962) was a paper on “Efficient estimates and optimum inference procedures in

¹As a reminder of C. R. Rao’s remarkable personal and professional longevity, this presentation took place before my parents had decided to conceive me.

large samples” (and his “m.l.” referred to maximum likelihood, not machine learning), one of a series of fundamental articles that he authored during what is now considered an era of classical mathematical statistics. Therefore, initially I was somewhat surprised by Rao’s apologetic sentiment—one that I ought to adopt myself for bringing up UMVUE in an era where few statistics students would recognize the acronym without Googling it. However, upon reflection, and considering his training under R. A. Fisher and the characteristically wry culture of RSS discussion at that time, I suspect Rao’s apology was more of a gentle reminder to not ignore established literature or wisdom when facing new problems. I am therefore grateful to the editors of this special issue, especially Bhramar Mukherjee, for the opportunity to honor Professor C. R. Rao with one more example of the value of such a reminder: how classical statistical results can offer insights and contextualization for modern work in data science like Bates *et al.* (2024).

I am also deeply grateful to Bhramar for her extraordinary patience in allowing me two extra months to complete this essay, without which I would have embarrassed myself significantly more by writing about the Heisenberg Uncertainty Principle (HUP) while knowing almost surely nothing even about classical mechanics². The connection between Cramér-Rao inequality and HUP has long been suspected, but I was unaware of any statistical literature on the connection between the two (however, during this work, I was made aware of such results in information theory—see Section 7).

Unfortunately, I had found neither the time nor the courage to explore quantum physics. Bhramar’s invitation gave me a great excuse to delve into it, though clearly it has been a quantum leap (or dive). I am therefore deeply grateful to the physicists, philosophers, and statisticians (see acknowledgment) who generously took the time to educate and inspire me, introducing me to numerous articles that, no doubt, will require another quantum-leap excuse to digest fully. These include physics literature on quantum Cramér-Rao bounds and quantum Fisher information (*e.g.*, Tóth and Petz, 2013; Tóth and Fröwis, 2022), as well as statistical writings on the relevance of quantum uncertainty to statistics (*e.g.*, Gelman and Betancourt, 2013), to name just a few.

Nevertheless, to set readers’ expectations realistically, this essay offers nothing about HUP that isn’t already in Wikipedia. I wrote much of it as reading notes to educate myself, so, paraphrasing a most memorable chiasmus from an RSS discussion: “The parts of the paper that are true are not new, and parts that are new are not true” (McCullagh, 1999). My hope, however, is that these notes may still be of use to those who share my curiosity (and innocence). I also hope that my attempt to extend the notion of covariance to quantum operators might encourage us to step out of our comfort zones without stepping out of our minds.

Intellectually, quantum indeterminacy is a captivating and challenging topic, especially for those of us who have been probability-law abiding citizens. To my knowledge, currently only a few statisticians—most notably Richard Gill³—have studied it systematically. Therefore, even if everything “new” in this essay ends up merely demonstrating that humans can out-hallucinate ChatGPT, I’d still be content dedicating it to the legendary C.

²Majoring in pure math in 1980s China means that I had taken no courses outside of mathematics, with the exception of mandatory ones for regulating students’ bodies or minds.

³See <https://www.math.leidenuniv.nl/~gillrd/>

R. Rao. Throughout his extraordinary career, Professor Rao applied his statistical insight and mathematical skills to establish and solidify the foundations of statistics. As quantum computing looms on the horizon, some statisticians should be leading the way in building the foundations of quantum data science, as articulated in the discussion article “When Quantum Computation Meets Data Science: Making Data Science Quantum” by Wang (2022), a prominent statistician exploring quantum computing’s role in data science. Thus, even if this essay inspires only one future C. R. Rao of quantum data science, it won’t take a quantum leap to believe that Professor Rao would embrace my dedication.

More broadly, I would find great professional satisfaction (and justification for my insomnia) if this essay serves as a reminder that time-honored statistical theory and wisdom have much to offer as we statisticians are increasingly called to step outside our comfort zones—from embracing machine learning to anticipating quantum computing. By learning from and contributing to other fields, especially time-tested ones such as philosophy and physics, we can enhance the intellectual impact of our discipline.

2. A paradox of error assessment?

Let us start with an excursion to the classical statistical sanctuary most frequently adopted in statistical research and pedagogy: we have an independently and identically distributed (i.i.d) normal sample, $X_1, \dots, X_n \stackrel{\text{iid}}{\sim} N(\mu, \sigma^2)$, and we are interested in making inferences about μ . It is well-known that the maximum likelihood estimator (MLE) for μ is the sample mean \bar{X}_n . The actual error of the MLE then is $\delta = \bar{X}_n - \mu$. It is textbook knowledge that the sample mean \bar{X}_n and the sample variance S_n^2 are independent under the normal model $N(\mu, \sigma^2)$. This fact is critical for establishing perhaps the most celebrated pivotal quantity in statistics, $t = \sqrt{n}(\bar{X} - \mu)/S_n$, *i.e.*, the t statistic, because of the existence of the parameter-free distribution of t for any $n \geq 2$, thanks to the aforementioned independence.

But this independence also implies a seemingly paradoxical fact that has received no mention in any textbook (that I am aware of): that $\hat{\delta}^2 \equiv S_n^2/n$ apparently is the worst estimate of the square of the actual error $\delta^2 = (\bar{X}_n - \mu)^2$, because $\hat{\delta}^2$ and δ^2 are independent of each other for any choice of $\theta = \{\mu, \sigma^2\}$. In what other context would a statistician (knowingly) suggest estimating an unknown with an independent quantity?

The article by Bates *et al.* (2024) reminds us that this seemingly paradoxical phenomenon is far more prevalent than we may have realized. To recast their findings in a broader setting but with a scalar estimand for notational simplicity, consider the possibly heteroscedastic linear regression setting,

$$Y_i = \theta X_i + \epsilon_i, \quad \text{where} \quad E[\epsilon_i | \mathbf{X}] = 0, V(\epsilon_i | \mathbf{X}) = \sigma_i^2, \quad i = 1, \dots, n. \quad (1)$$

and conditioning on $\mathbf{X} = \{X_1, \dots, X_n\}$, $\{\epsilon_1, \dots, \epsilon_n\}$ are mutually independent. As Bates *et al.* (2024) reminds us, when $\{\epsilon_1, \dots, \epsilon_n\}$ are i.i.d $N(0, \sigma^2)$, the least-squares estimator for θ , $\hat{\theta}_{\text{LS}} = \sum_{i=1}^n Y_i X_i / \sum_{i=1}^n X_i^2$ is independent of the residual $R = \{\hat{r}_i = Y_i - \hat{\theta} X_i, i = 1, \dots, n\}$, for any given $\{\theta, \sigma^2\}$. Consequently, since the true predictive error depends on the data only through $\hat{\theta}_{\text{LS}}$, and cross-validation error estimators are functions only of the residuals, the true and estimated errors are independent of each other. The results obviously apply to any

error estimates that depend on data only through R , which is the case virtually for all the common estimators in practice, as demonstrated in Bates *et al.* (2024).

It is well-known (*e.g.*, Casella and Berger, 2024) that under the i.i.d normal setting, $\hat{\theta}_{LS}$ is the MLE and indeed UMVUE (uniformly minimum variance unbiased estimator) because its variance reaches the Cramér-Rao bound. Even without the normality, we know that $\hat{\theta}_{LS}$ is BLUE (best linear unbiased estimator) and it is linearly uncorrelated with the residual R under the squared loss, because it is the orthogonal projection of Y onto the space expanded by \mathbf{X} when σ_i is invariant of i .

Although rarely mentioned in textbooks, this optimality-orthogonality duality appears in essentially all inferential paradigms. Geometrically speaking, the equivalence is due to the fact that the linear correlation between two variables is the cosine of the angle between them in the L^2 space, and optimal projection is the orthogonal projection. Probabilistically, the ubiquity of this duality is manifested by the so-called “Eve’s law” (Blitzstein and Hwang, 2014), an instance of the Pythagorean theorem in the L^2 space.

That is, under any joint distribution, $p(H, G)$, as long as it generates finite second moments, $\text{Cov}[H - \mathbb{E}(H|G), \mathbb{E}(H|G)] = 0$, because $\mathbb{E}(H|G)$ is the orthogonal projection of H to the space of L^2 functions that are measurable with respect to the σ -field generated by G . Consequently, the Pythagorean theorem is in force:

$$\begin{aligned} V(H) &= \mathbb{E}[H - \mathbb{E}(H)]^2 = \mathbb{E}[H - \mathbb{E}(H|G)]^2 + \mathbb{E}[\mathbb{E}(H|G) - \mathbb{E}(H)]^2 \\ &= \mathbb{E}[V(H|G)] + V[\mathbb{E}(H|G)], \end{aligned} \quad (2)$$

which is Eve’s law. The ubiquity of the duality is due to the fact that the expectation operator in (2) can be taken with any kind of distribution: posterior (predictive) distributions for Bayesian inferences, super-population distributions as typical for likelihood inference (as in the $N(\mu, \sigma^2)$ example), or randomization distributions as in finite-population calculations (as adopted in Meng, 2018).

Nevertheless, this duality is a qualitative statement, as it does not quantify what happens for non-optimal estimation or learning. As demonstrated below, this duality can be extended quantitatively by tethering the deficiency in learning with the relevancy in assessing the actual learning errors. This quantification crystallizes the reason for the apparent paradox, and it can help reduce wasted efforts in pursuit of the impossible. It also makes it clearer that there is no real paradox, much like how Simpson’s paradox is not a paradox once its workings are revealed and understood (*e.g.* Liu and Meng, 2014; Gong and Meng, 2021).

The title of the next section says it all: there is no free lunch. If there is any data information left—after learning—for assessing the actual error, then we can reduce the actual error by removing the part that can be predicted by the untapped data information. This implies our learning is not optimal, and vice versa. Section 3 illustrates this fact in the context of heteroscedastic regression, followed by a broad reflection in Section 4 on its implications in the context of error assessment without external benchmarks, a statistical magic. Sections 5 and 6 then establish respectively the exact and asymptotic inequalities that capture the learning uncertainty principle under the squared loss.

To facilitate a formal comparison with HUP using the notion of co-variation, Sec-

tion 7 discusses the generalization of the measure of co-variance from real-valued variables to complex-valued variables and functions. Section 8 then applies the generalization to the case of HUP by defining co-variances between mechanisms (*e.g.*, the position and momentum *operators*) rather than between the states they generate (*e.g.*, the actual position and momentum states). With these preparations, Section 9 compares the learning-error inequality, Cramér-Rao inequality, and HUP inequality, highlighting their shared essence from a statistical perspective.

Section 10 reflects on various philosophical issues surrounding uncertainty principles in general, and HUP in particular, with insights from the encyclopedic essay by Hilgevoord and Uffink (2024). Section 11 briefly touches on the trade-off between quantitative and qualitative studies, prompted by a discussion in Hilgevoord and Uffink (2024), and how intellectual inquires can benefit from their happy marriage. This leads to a piece of advice from Professor Rao on living a happy life, which serves as a fitting conclusion to this essay in his memory. However, to encourage students to engage with this essay to the fullest extent of their attention spans, Section 12 provides a prologue, especially for those who may not enjoy technical appendices but wish the essay were even longer.

3. Once again, there is no free lunch

Consider the heteroscedastic setting (1), where we know that BLUE is given by the weighted LS, in the form of

$$\hat{\theta}_w = \frac{\sum_{i=1}^n w_i Y_i X_i}{\sum_{i=1}^n w_i X_i^2}, \quad (3)$$

when the weights $w_i \propto \sigma_i^{-2}, i = 1, \dots, n$. Now consider an arbitrarily weighted $\hat{\theta}_w$, and its correlation—denoted by ρ —with the corresponding residual $R_w = \{\hat{r}_{w,i} = Y_i - \hat{\theta}_w X_i; i = 1, \dots, n\}$. For conveying the main idea, the case of $n = 2$ is sufficient. As a special case of the general expression given in Appendix A, we have, conditioning on \mathbf{X} (but we suppress this conditioning notation-wise unless necessary),

$$\rho^2(\hat{\theta}_w, \hat{r}_{w,i}) = \frac{X_1^2 X_2^2 (w_1 \sigma_1 \sigma_2^{-1} - w_2 \sigma_2 \sigma_1^{-1})^2}{(w_1^2 X_1^2 \sigma_1^2 + w_2^2 X_2^2 \sigma_2^2)(X_1^2 \sigma_1^{-2} + X_2^2 \sigma_2^{-2})}, \quad i = 1, 2, \quad (4)$$

which is zero if and only if $w_i \propto \sigma_i^{-2}, i = 1, 2$ (as long as $X_i \neq 0, i = 1, 2$). That is, $\hat{\theta}_w$ is BLUE (or the MLE if we assume normality) if and only if $\hat{\theta}_w$ is uncorrelated with $\hat{r}_{w,i}$. More importantly, expression (4) tells us exactly how the statistical efficiency of $\hat{\theta}_w$ is directly linked to this correlation.

Specifically, let $\hat{\theta}_{\text{BLUE}}$ be the optimally weighted LS estimator with weight $w_i \propto \sigma_i^{-2}, i = 1, 2$, and RR_w be the relative regret of an arbitrarily weighted $\hat{\theta}_w$ under the squared loss, that is,

$$RR_w = \frac{V(\hat{\theta}_w) - V(\hat{\theta}_{\text{BLUE}})}{V(\hat{\theta}_w)} = 1 - \frac{(w_1 X_1^2 + w_2 X_2^2)^2}{(w_1^2 X_1^2 \sigma_1^2 + w_2^2 X_2^2 \sigma_2^2)(X_1^2 \sigma_1^{-2} + X_2^2 \sigma_2^{-2})}. \quad (5)$$

Whereas it may not be immediate from (4) and (5), one can verify directly that

$$\rho^2(\hat{\theta}_w, \hat{r}_{w,i}) = RR_w, \quad i = 1, 2, \quad (6)$$

for any choice of weights w or values of $\{\sigma_i^2, i = 1, 2\}$. This means that if we want to increase the magnitude of the correlation between $\hat{\theta}_w$ and $\hat{r}_{w,i}$, we must sacrifice the efficiency of $\hat{\theta}_w$, and vice versa.

However, why might we want to increase $|\rho(\hat{\theta}_w, \hat{r}_{w,i})|$? Consider the case where our learning target is $c\theta$, with c being a constant. For example, we take $c = 1$ when the regression coefficient θ is the target, or $c = X^*$ when the learning target is the mean of Y when $X = X^*$. In such cases, the actual error is given by $\delta_w = c(\hat{\theta}_w - \theta)$. We can assess δ_w via $\hat{\delta}_w = \tilde{c}\hat{r}_{w,1}$ for some choice of \tilde{c} (recall $\hat{r}_{w,1} + \hat{r}_{w,2} = 0$ and hence a single residual suffices). Because

$$\rho^2(\delta_w, \hat{\delta}_w) = \rho^2(c\hat{\theta}_w, \tilde{c}\hat{r}_{w,1}) = \rho^2(\hat{\theta}_w, \hat{r}_{w,1}), \tag{7}$$

we see that by moving $\rho^2(\hat{\theta}_w, \hat{r}_{w,1})$ away from zero, we will have an assessment $\hat{\delta}_w$ of the actual error δ_w that has a degree of conditional relevancy, that is, $\hat{\delta}_w$ is at least correlated with δ_w conditioning on the setting (1). But this gain of relevancy is achieved necessarily by increasing the relative regret (recall the relative regret for $c\hat{\theta}_w$ is invariant to the value of c), that is, by sacrificing the efficiency of $\hat{\theta}_w$, because

$$\rho^2(\delta_w, \hat{\delta}_w) = RR_w, \tag{8}$$

thanks to (6)-(7).

If our learning target is to predict (a new) Y^* when $X = X^*$, then the actual prediction error is $\delta_w^* = Y^* - \hat{\theta}_w X^*$. In such cases, the prediction risk under the squared loss is

$$E(Y^* - \hat{\theta}_w X^*)^2 = V(Y^*) + (X^*)^2 V(\hat{\theta}_w).$$

Because $V(Y^*)$ and $(X^*)^2$ are invariant to the weights, we obtain the relative regret for prediction $RR_w^* = \gamma RR_w$, where RR_w is from (5) and the adjustment factor γ is given by

$$\gamma = \frac{(X^*)^2 V(\hat{\theta}_w)}{V(Y^*) + (X^*)^2 V(\hat{\theta}_w)}. \tag{9}$$

Furthermore, because $\hat{\delta}_w = \tilde{c}\hat{r}_{w,1}$ is independent of Y^* , $\text{Cov}(\delta_w^*, \hat{\delta}_w) = -X^* \text{Cov}(\hat{\theta}_w, \hat{\delta}_w)$. Hence,

$$\rho^2(\delta_w^*, \hat{\delta}_w) = \frac{(X^*)^2 \text{Cov}^2(\hat{\theta}_w, \hat{\delta}_w)}{[V(Y^*) + (X^*)^2 V(\hat{\theta}_w)] V(\hat{\delta}_w)} = \gamma \rho^2(\hat{\theta}_w, \hat{\delta}_w). \tag{10}$$

Consequently, the identity (8) holds for both estimation and prediction, implying the same trade-off between optimal learning and relevant error assessment.

Section 5 below will provide a general inequality that captures this trade-off under squared loss, for which identity (8) is a special case. But before presenting that result, we must ask: if we cannot relevantly assess the actual error δ , then what kind of errors have we been assessing? And that is exactly one of the two questions raised in the title of Bates *et al.* (2024): Cross-validation: what does it estimate and how well does it do it? The following section supplements Bates *et al.* (2024) to answer this question more broadly and more pedagogically.

4. Jay Leno's irony and a statistical magic

During one of the years the United States census took place (likely 2000-2001), comedian Jay Leno brought up the issue of under-counting on his *Tonight Show*. He began by informing the audience that the U.S. Census Bureau had just reported that approximately p percentage of the population had not been counted. With an arch smile, he then quipped, "But I don't understand—if they knew they missed p percentage of people, why didn't they just add it back?" (The actual value p he used now lies deep in my memory.)

The audience was amused, as was I, though perhaps for different reasons—what amused me was the very appearance of such a nerdy joke on a mainstream comedy show. Humor is often rooted in life's ironies, and whoever crafted this joke clearly understood the irony in announcing both an estimate and its error. In the case of the U.S. Census, the irony—or more accurately, the *magic*—is not as profound as it may seem. The estimation of undercount relies on external data, such as demographic analysis, post-enumeration surveys, administrative records, and other sources. The term *magic* is used here because statistical inference can appear magical to uninitiated yet inquisitive minds. How can one estimate an unknown quantity, and then estimate the error of that estimation, without any external knowledge of the true value?

The magic begins with a sleight of hand—in this case, the word *error* does not refer to the actual error, as a layperson might assume. Instead, we aim to understand the statistical properties of the actual error by imagining its variations across hypothetical replications. The construction of these replications depends on the philosophical framework one subscribes to, with the two main schools being frequentist and Bayesian (but see Lin (2024b) for a spectrum between them). Perhaps surprisingly, the key to resolving the apparent paradox in Section 2 lies in adopting insights from both perspectives.

To see this, consider again the normal example where the true error is $\delta = \bar{X}_n - \mu$. In the frequentist framework, the hypothetical replications consist of all possible copies of $D = \mathbf{X} = \{X_1, \dots, X_n\}$ generated from $N(\mu, \sigma^2)$ with the *same* but unknown parameter values $\theta = \{\mu, \sigma^2\}$. In this replication setting, the expected value of δ^2 , which is the sampling variance of \bar{X}_n , equals σ^2/n . It is well-known that under the same replication framework, the expectation of $\hat{\delta}^2 = S_n^2/n$ is also σ^2/n .

Thus, while δ^2 and $\hat{\delta}^2$ are independent of each other for any given $\theta = \{\mu, \sigma^2\}$, they share the same expectation within the frequentist framework. By invoking the same leap of faith that underpins the frequentist approach—trusting and transferring average behaviors to assess individual cases—we justify $\hat{\delta}^2$ as an estimate of δ^2 . Such a leap of faith exists regardless of the goal of our data exercise, be it prediction, estimation, or attribution (significance testing), albeit with increased levels of intolerance to the inaccuracy in error assessing, as revealed by the insightful article of Efron (2020).

For Bayesians, such a leap of faith is unconvincing or even "irrelevant" in the sense of Dempster (1963), as the actual error can differ significantly from its expectation. The independence between $\hat{\delta}^2$ and δ^2 suggests that accepting this leap would require a religious level of faith. In the Bayesian framework, the relevant hypothetical replications include all possible values of $\theta = \{\mu, \sigma^2\}$ (and their associated probabilities) that could have generated the same data set D , and therefore the same $\{\bar{X}_n, S_n^2\}$.

However, for such a replication setting to be realized—for instance, via a simulation—a prior distribution for $\theta = \{\mu, \sigma^2\}$ must be assumed. This postulation represents the Bayesian leap of faith in actual implementations, since it is virtually certain that a part of the assumption is faith-based instead of knowledge-driven; for a broader discussion on the necessity of such leaps across all major schools of statistical inference—Bayesian, Fiducial, and Frequentist (BFF)—see Craiu *et al.* (2023) and more comprehensively the *Handbook on BFF Inference* edited by Berger *et al.* (2024).

Although we shall not take a Bayesian excursion here, we can borrow the Bayesian concept of allowing $\theta = \{\mu, \sigma^2\}$ to have a distribution in order to establish a joint replication setting, where both D and $\theta = \{\mu, \sigma^2\}$ vary. This framework is relevant (for frequentists) when recommending the same statistical procedure across multiple studies with normal data, where both D and $\theta = \{\mu, \sigma^2\}$ may differ from study to study. In the machine learning world—or any domain reliant on training data—such a joint replication setting can be visualized as potential training datasets drawn from related populations, which makes transfer learning a meaningful endeavor (*e.g.*, Abba *et al.*, 2024).

For our normal example, given any proper prior on θ , it can be shown (see Appendix B) that over any proper joint replication of $\{D, \theta\}$,

$$\rho(\hat{\delta}^2, \delta^2) = \frac{\gamma_{\sigma^2}^2}{\sqrt{\frac{n+1}{n-1}\gamma_{\sigma^2}^2 + \frac{2}{n-1}\sqrt{3\gamma_{\sigma^2}^2 + 2}}}, \quad (11)$$

where γ_{σ^2} is the coefficient of variation of σ^2 with respect to the (proper) prior distribution of σ^2 . This correlation is non-negative, providing a plausible measure of how relevant $\hat{\delta}^2$ is for assessing δ^2 . It is zero if and only if $V(\sigma^2) = 0$, meaning that we revert to the situation of conditioning on a fixed σ^2 : since S_n^2 is invariant to μ , $\hat{\delta}^2$ and δ^2 remain independent when conditioned on σ^2 alone. The fact that (11) is a monotonic increasing function of γ_{σ^2} implies that the relevance of $\hat{\delta}^2$ for assessing δ^2 increases as the heterogeneity among the studies—in terms of the within-study variation indexed by σ^2 —grows. This monotonicity is intuitive, given that S_n^2 is an unbiased and asymptotically efficient estimator of σ^2 , and $\hat{\delta}^2$ is useful for comparing the magnitudes of δ^2 across studies with different σ^2 values. However, the fact that this correlation can never exceed $1/\sqrt{3} \approx 0.577$ is unexpected. For those of us who believe that mathematical results are never coincidental, contemplating the intricacies of this bound might induce insomnia (while serving as a cure for many others).

This joint replication framework clarifies the role of $\hat{\delta}^2$ as an *adaptive benchmark* for assessing the statistical properties of δ^2 over the hypothetical replications. *That* is statistical magic—the ability to establish cross-study comparisons based on a single study. More broadly, the magic lies in creating hypothetical “control” replications $\{\tilde{D}, \tilde{\theta}\}$ from the actual “treatment” $\{D, \theta\}$ at hand, as elaborated in Liu and Meng (2016), borrowing the metaphor of individualized treatment.

Generally speaking, the magic relies on two tricks: (I) creating replications within D , and (II) linking those replications to the imagined variations of D through the within- D replications from (I). The first trick is applicable when the mechanism generating the data D inherently includes (higher resolution) replications, either by design (*e.g.*, simple random sampling) or by declaration (*e.g.*, imposing an i.i.d. structure as a working assumption).

The second trick is enabled by theoretical understanding (*e.g.*, the relationship between the distribution of the sample mean and the distribution of the individual samples) or by simulations and approximations that are enabled by (I), such as the Bootstrap (see Craiu *et al.*, 2023, for a discussion).

The magic metaphor also serves as a reminder that magic relies on illusions, and interpreting average errors as actual ones is one such illusion. With that understanding, we might wonder if it's possible to assess the actual error with greater relevance. For example, in the normal case, one might ask whether a different error estimate $\check{\delta}$ could be more relevant for $\delta = \bar{X}_n - \mu$, in the sense that $\rho(\check{\delta}, \delta) > 0$ given any value of $\theta = \{\mu, \sigma^2\}$. The classical statistical literature offers a fairly clear answer to this question, as discussed below.

5. From UMVUE to an uncertainty principle for unbiased learning

The celebrated Cramér–Rao bound, more broadly known as the information inequality (see Lehmann and Casella, 2006, Ch. 2), tells us that if $\hat{\theta}$ is an unbiased estimator for θ under a parametric model $f(D|\theta)$, then under mild conditions, $V(\hat{\theta}) \geq I^{-1}(\theta)$, where $I(\theta)$ is the expected Fisher information. For the normal example, when we take $\theta = \mu$ (temporarily assuming σ^2 is known), we have $V(\bar{X}_n) = \sigma^2/n = I^{-1}(\mu)$, where $I(\mu)$ is the expected Fisher information from $f(X_1, \dots, X_n|\mu)$. Thus, we know \bar{X}_n is UMVUE for μ .

It is well-known that an estimator $\hat{\theta}$ is UMVUE if and only if it is uncorrelated with any unbiased estimator U for zero for any θ (see Lehmann and Casella, 2006, Ch. 2), that is, $E_\theta[(\hat{\theta} - \theta)U] = 0$, whenever $E_\theta(U) = 0$. Since $\hat{\theta} - \theta$ is simply the actual error δ , this result implies that conditioning on θ , it is impossible to have an error assessment $\hat{\delta}$ for δ that is both unbiased and relevant at the same time, *i.e.*, $E_\theta(\hat{\delta}) = 0$ and $\rho_\theta^2(\hat{\delta}, \delta) > 0$ cannot hold simultaneously for any θ , where we inject the subscript θ in ρ_θ to explicate that the correlation is with respect to $f(D|\theta)$ for fixed θ .

Intuitively, if any unbiased error assessment $\hat{\delta}$ is correlated with δ , then some part of the actual error δ is predictable by $\hat{\delta}$. This means that we could improve $\hat{\theta}$ without losing its unbiasedness, which contradicts the fact that $\hat{\theta}$ is already an UMVUE. An astute reader may quickly recognize that this insight has much broader implications than merely for UMVUEs. The following result is a proof of this realization, using the same proof strategy as for UMVUE, but establishes a broader quantitative result than the aforementioned qualitative “if and only if” result for UMVUE. The result is presented in the scalar case for simplicity, but its multivariate counterpart can be derived easily using corresponding matrix notation.

Specifically, let $Q \in \mathbb{R}$ be our target of learning, which could represent a future outcome, a model parameter, a latent trait, *etc.* Suppose the state space of our data D is Ω and $\hat{Q} : \Omega \rightarrow \mathbb{R}$ is our learning algorithm, or a *learner* for Q . For any learner \hat{Q} , let $\hat{\delta}_{\hat{Q}} : \Omega \rightarrow \mathbb{R}$ be an assessment (*e.g.*, an estimator) of the exact (additive) error of \hat{Q} , namely, $\delta_{\hat{Q}} = \hat{Q} - Q$. Let $L(\hat{Q}, Q)$ be the loss function, and $\mathcal{P} = \{P_s(D; Q), s \in S\}$ be the family of distributions under which we calculate the learning risk: $R_s(\hat{Q}) = E_s[L(\hat{Q}, Q)]$. Note that Q may be a function of s (*e.g.*, when estimating the model parameter s) or it may be a random variable itself (*e.g.*, a future realization), in which case the notation $P_s(D; Q)$ represents the joint distribution over D and Q .

Theorem 1: Let $L(\hat{Q}, Q) = (\hat{Q} - Q)^2$ be the squared loss, and let $L_{\mathcal{P}}^2$ denote the collection of all square-integrable functions with respect to \mathcal{P} . Define

$$\mathcal{Q} = \{\hat{Q}(D) \in L_{\mathcal{P}}^2 : E_s(\hat{Q} - Q) = 0, \forall s \in S\} \quad (12)$$

as the collection of unbiased learners of Q with respect to \mathcal{P} . For any $\hat{Q} \in \mathcal{Q}$, define

$$\mathcal{E}(\hat{Q}) = \{\hat{\delta}_{\hat{Q}}(D) \in L_{\mathcal{P}}^2 : E_s(\hat{\delta}_{\hat{Q}}) = 0, \forall s \in S\} \quad (13)$$

as the collection of corresponding unbiased error assessors for $\delta_{\hat{Q}}$. Suppose there exists an optimal learner $\hat{Q}^{\text{opt}} \in \mathcal{Q}$, with risk $R_s^{\text{opt}} < \infty$ under $f_s, s \in S$. Then:

(I) For any $\hat{Q} \in \mathcal{Q}$ and any corresponding $\hat{\delta}_{\hat{Q}} \in \mathcal{E}(\hat{Q})$, we have

$$\rho_s^2(\delta_{\hat{Q}}, \hat{\delta}_{\hat{Q}}) \leq \frac{R_s(\hat{Q}) - R_s^{\text{opt}}}{R_s(\hat{Q})} \equiv RR_s(\hat{Q}), \quad \forall s \in S, \quad (14)$$

where $RR_s(\hat{Q})$ is the relative regret of \hat{Q} under distribution P_s , and it is set to zero if $R_s(\hat{Q}) = 0$.

(II) Equality $\rho_s^2(\delta_{\hat{Q}}, \hat{\delta}_{\hat{Q}}) = RR_s(\hat{Q})$ holds for any particular s if and only if R_s^{opt} is attainable in the sub-class $\mathcal{Q}(\hat{Q}, \hat{\delta}_{\hat{Q}}) = \{\hat{Q} - \lambda\hat{\delta}_{\hat{Q}} : \forall \lambda \in \mathbb{R}\} \subset \mathcal{Q}$.

Proof: For any given $\hat{Q} \in \mathcal{Q}$ (which is non-empty since $\hat{Q}^{\text{opt}} \in \mathcal{Q}$) and any $\hat{\delta}_{\hat{Q}} \in \mathcal{E}(\hat{Q})$ (which is non-empty since $\hat{\delta}_{\hat{Q}} \equiv 0$ is always included), we define $\hat{Q}_{\lambda} = \hat{Q} - \lambda\hat{\delta}_{\hat{Q}}$ for any constant $\lambda \in \mathbb{R}$. Under our assumptions, $E_s(\hat{Q}_{\lambda} - Q) = 0$, and $\hat{Q}_{\lambda} \in L_{\mathcal{P}}^2$, implying $\hat{Q}_{\lambda} \in \mathcal{Q}$. Since $\hat{Q}_{\lambda} - Q = \delta_{\hat{Q}} - \lambda\hat{\delta}_{\hat{Q}}$ and it has mean zero under $f_s(D; Q)$, we have

$$R_s^{\text{opt}} \leq R_s(\hat{Q}_{\lambda}) = V_s(\delta_{\hat{Q}} - \lambda\hat{\delta}_{\hat{Q}}) = V_s(\delta_{\hat{Q}}) + \lambda^2 V_s(\hat{\delta}_{\hat{Q}}) - 2\lambda \text{Cov}_s(\delta_{\hat{Q}}, \hat{\delta}_{\hat{Q}}), \quad \forall s \in S. \quad (15)$$

Since the left-hand side of this inequality is free of λ , the inequality holds when we minimize the right-hand side over $\lambda \in \mathbb{R}$, which is achieved at $\lambda = \lambda^* = \text{Cov}_s(\delta_{\hat{Q}}, \hat{\delta}_{\hat{Q}}) / V_s(\hat{\delta}_{\hat{Q}})$, assuming $V_s(\hat{\delta}_{\hat{Q}}) > 0$. (When $V_s(\hat{\delta}_{\hat{Q}}) = 0$, $\rho_s(\delta_{\hat{Q}}, \hat{\delta}_{\hat{Q}}) = 0$; hence (14) holds trivially, and we can set $\lambda^* = 0$.) Thus, we obtain

$$R_s^{\text{opt}} \leq V_s(\delta_{\hat{Q}}) \left[1 - \rho_s^2(\delta_{\hat{Q}}, \hat{\delta}_{\hat{Q}}) \right], \quad \forall s \in S,$$

which yields (14) since $R_s(\hat{Q}) = V_s(\delta_{\hat{Q}})$ when $E_s(\delta_{\hat{Q}}) = 0$. This proves part (I).

Part (II) follows from (15) as well, because the equality holds there if and only if R_s^{opt} is attainable by $\hat{Q}_{\lambda^*} \in \mathcal{Q}(\hat{Q}, \hat{\delta}_{\hat{Q}})$. This includes the case with $V_s(\hat{\delta}_{\hat{Q}}) = 0$, where the result holds trivially, because then $\rho_s(\delta_{\hat{Q}}, \hat{\delta}_{\hat{Q}}) = 0$ and $R_s(\hat{Q}) = R_s^{\text{opt}}$, *i.e.*, \hat{Q} itself is optimal. \square

The immediate implication of inequality (14) is that there is no free lunch. If we want to increase the relevance of our assessment $\hat{\delta}_{\hat{Q}}$ for the actual error $\delta_{\hat{Q}}$ by increasing their

correlation, we must also increase the relative regret for \hat{Q} , effectively sacrificing degrees of freedom of learning for the error assessment. Conversely, the less regret in \hat{Q} , the less relevant its error assessment will be to the actual error. In the extreme case, when $\hat{Q} = \hat{Q}^{\text{opt}}$, we arrive at the following result, where by a *relevant error assessor* we mean it is linearly correlated with the actual error of the learner.

Corollary 1: Under the same setup as in Theorem 1, the following two assertions cannot hold simultaneously:

- (A) $\hat{Q} \in \mathcal{Q}$ is an *optimal and unbiased learner* for Q under P_s ; and
- (B) \hat{Q} has an *unbiased and relevant error assessor* $\hat{\delta}_{\hat{Q}} \in \mathcal{E}(\hat{Q})$.

6. Beyond unbiased learning and error assessing

A key limitation of Theorem 1 is the requirement that both the learner and error assessor must be unbiased. An immediate generalization is to consider cases where both are asymptotically unbiased, under an asymptotic regime with respect to some information index ι , such as the size of data. Mathematically, given a sequence of error order e_ι such that $\limsup_{\iota \rightarrow \infty} |e_\iota| = 0$, we can modify the classes of the learners and error assessors in (12) and (13) respectively by

$$\begin{aligned} \mathcal{Q}_\iota &= \{\hat{Q}(D) \in L_{\mathcal{P}}^2 : E_s[\hat{Q}(D) - Q] = O(e_\iota), \forall s \in S\}, & (16) \\ \mathcal{E}_\iota(\hat{Q}) &= \{\hat{\delta}_{\hat{Q}}(D) \in L_{\mathcal{P}}^2 : E_s(\hat{\delta}_{\hat{Q}}) = O(e_\iota), \forall s \in S\}, & (17) \end{aligned}$$

where $O(e_\iota)$ is the standard notation for being of the same order as e_ι . That the error assessor $\hat{\delta}_{\hat{Q}}$ must share the same order of expectation as the actual error $\delta_{\hat{Q}}$ is a necessary requirement to render the term ‘error assessor’ meaningful, as otherwise anything could be regarded as $\hat{\delta}_{\hat{Q}}$. With these modifications, we have the following asymptotic counterpart of Theorem 1.

Theorem 2: Assume the same setup as Theorem 1, but with \mathcal{Q} and $\mathcal{E}(\hat{Q})$ extended respectively to \mathcal{Q}_ι and $\mathcal{E}_\iota(\hat{Q})$. We then have

$$\rho_s^2(\delta_{\hat{Q}}, \hat{\delta}_{\hat{Q}}) \leq RR_s(\hat{Q}) + O(e_\iota^2), \quad \forall s \in S, \quad (18)$$

where e_ι is a sequence of vanishing error rates that determines the asymptotic regime.

Proof: For $\hat{Q} \in \mathcal{Q}_\iota$ and $\hat{\delta}_{\hat{Q}} \in \mathcal{E}_\iota(\hat{Q}_\iota)$, we can write $E_s(\delta_{\hat{Q}}) = a_\iota$ and $E_s(\hat{\delta}_{\hat{Q}}) = b_\iota$ where $a_\iota = O(e_\iota)$ and $b_\iota = O(e_\iota)$ by our assumption. Hence for $\hat{Q}_\lambda = \hat{Q} - \lambda \hat{\delta}_{\hat{Q}}$, $E_s(\hat{Q}_\lambda - Q) = a_\iota - \lambda b_\iota = O(e_\iota)$ for any λ , implying that $\hat{Q}_\lambda \in \mathcal{Q}_\iota$. Let λ^* be the minimizer of $V_s[\delta_{\hat{Q}} - \lambda \hat{\delta}_{\hat{Q}}]$, as defined in the proof of Theorem 1. The optimality of R_s^{opt} then implies that

$$\begin{aligned} R_s^{\text{opt}} &\leq R_s(\hat{Q}_{\lambda^*}) = V_s[\delta_{\hat{Q}} - \lambda^* \hat{\delta}_{\hat{Q}}] + [E_s(\delta_{\hat{Q}} - \lambda^* \hat{\delta}_{\hat{Q}})]^2 \\ &= V_s(\delta_{\hat{Q}}) [1 - \rho_s^2(\delta_{\hat{Q}}, \hat{\delta}_{\hat{Q}})] + (a_\iota - \lambda^* b_\iota)^2. \\ &\leq R_s(\delta_{\hat{Q}}) [1 - \rho_s^2(\delta_{\hat{Q}}, \hat{\delta}_{\hat{Q}})] + (a_\iota - \lambda^* b_\iota)^2. \end{aligned}$$

But this proves the inequality (18) because $(a_\iota - \lambda^* b_\iota)^2 = O^2(e_\iota) = O(e_\iota^2)$. \square

A major application of Theorem 2 is for the maximum likelihood estimator \hat{Q}_{MLE} , which under regularity conditions is efficient and asymptotically normal (e.g., Lehmann and Casella, 2006) and hence it is asymptotically optimal under the squared loss. Theorem 2 says that asymptotically, there cannot be any relevant error assessor $\hat{\delta}_{\text{MLE}} \in \mathcal{E}_\iota(\hat{Q}_{\text{MLE}})$ that is asymptotically correlated with the actual error $\delta_{\text{MLE}} = \hat{Q}_{\text{MLE}} - Q$. When $\{\hat{\delta}_{\text{MLE}}, \delta_{\text{MLE}}\}$ are jointly asymptotically normal, then Theorem 2 would imply that any such $\hat{\delta}_{\text{MLE}}$ will be asymptotically independent of δ_{MLE} . It is worthy noting that the same would hold for any estimator that is asymptotically normal and optimal (under quadratic loss), such as those studied in the classic work by Wald (1943) and Le Cam (1956).

Because the asymptotic variance of the MLE can be well approximated by the inverse of Fisher information, especially the observed Fisher information (Efron and Hinkley, 1978), the preceding result might lead some readers to wonder if the MLE and the observed Fisher information are asymptotically independent, or at least the MLE and the inverse of the observed Fisher information $I_{\text{obs}}^{-1}(\hat{Q})$ are asymptotically uncorrelated. The normal example given in Section 2 may be especially suggestive, since the MLE for μ , \bar{X}_n , is independent of $I_{\text{obs}}^{-1}(\hat{\mu}) = n/\hat{\sigma}_{\text{MLE}}^2 = n^2/[(n-1)S_n^2]$. However it will be a mistake to generalize from this example.

Consider the same normal model $N(\mu, \sigma^2)$, but our goal now is to estimate the variance σ^2 . The MLE for σ^2 is $\hat{\sigma}_{\text{MLE}}^2 = (n-1)S_n^2/n$, and the corresponding observed Fisher information (pretending μ is known) is $I^{-1}(\hat{\sigma}_{\text{MLE}}^2) = 2\hat{\sigma}_{\text{MLE}}^4/n$; hence they have a deterministic relationship. However, this is not a contradiction to Theorem 2 because $I^{-1}(\hat{\sigma}_{\text{MLE}}^2)$ is not an unbiased assessment of the actual error, but rather its variance. Since the variance is effectively an index of the *problem difficulty* for estimation (as termed in Meng, 2018), it is entirely natural to expect that the variance can vary closely with the value of the estimand. The normal mean problem is a special case because it is a location family, for which shifting the mean only changes the value of the estimand, but does not alter the difficulty of its estimation. This point is reinforced if we reparameterize σ^2 via $\eta = \log \sigma^2$, which yields $\hat{\eta}_{\text{MLE}} = \log \hat{\sigma}_{\text{MLE}}^2$ and $I^{-1}(\hat{\eta}_{\text{MLE}}) = 2/n$, and they are now trivially independent of each other, because $\hat{\eta}_{\text{MLE}} - \eta \sim \log \chi_{n-1}^2 - \log(n-1)$ is a location family.

The consideration of the relationship between the MLE and the Fisher information provides a natural segue to the following discussion involving the relationship between inequity (14) and the Cramér-Rao low bound. As is well documented⁴, the seminal work by Rao (1945) was prompted by a question raised during a lecture Rao gave in 1943 on whether there could be a small-sample counterpart of the asymptotic efficiency for MLE as captured by the Fisher information. However, the significance of this work goes beyond accenting the role of Fisher information, because the Cramér-Rao inequality can be viewed as a statistical counterpart of the fundamental Heisenberg Uncertainty Principle (Griffiths and Schroeter, 2018) via the notion of *co-variation*, as explored in the next three sections.

⁴See the video on C.R. Rao: A Life in Statistics II at <https://www.youtube.com/watch?v=eaxjUxoCx5w&t=324s>

7. Measuring co-variation without probabilistic joint-state specifications

In statistical and (ordinary) probabilistic literature, the most commonly adopted measure of the co-variation of two real-valued random variables G and H is their covariance $\text{Cov}(G, H)$ (which includes correlation once G and H are standardized) defined via their joint probabilistic distribution $F_{G,H}(g, h)$:

$$\text{Cov}(G, H) = \int \int (g - \mu_G)(h - \mu_H) F_{G,H}(dg, dh) = \langle (g - \mu_G), (h - \mu_H) \rangle_F, \quad (19)$$

where μ_G and μ_H are respectively the means of G and H , which, without loss of generality, we will assume to be zero for the subsequent discussions for notational simplicity. The subscript F in the inner product notation highlights the critical dependence of $\text{Cov}(G, H)$ on their joint distribution $F(g, h)$. The elegant Hoeffding identity (Hoeffding, 1940)

$$\text{Cov}(G, H) = \int \int [F_{G,H}(g, h) - F_G(g)F_H(h)] dgdh, \quad (20)$$

where F_G and F_H are the marginal (cumulative) distributions, further highlights how the covariance measures the co-variation in G and H as captured by their joint distribution, with respect to their benchmark distribution under the assumption of independence.

For HUP, it seems natural to take $G = x$, the position of a particle, and $H = p$, its momentum, to follow the standard notation in quantum mechanics. It is textbook knowledge (*e.g.* Landau and Lifshitz, 2013; Griffiths and Schroeter, 2018) that densities of the position x and momentum p are given by $|\psi(x)|^2$ and $|\varphi(p)|^2$ respectively, where $\psi(x)$ is a complex-valued position wave function, and the momentum wave function $\varphi(p)$ is a scaled Fourier transform of $\psi(x)$ in the form of

$$\varphi(p) = \frac{1}{\sqrt{2\pi\hbar}} \int_{-\infty}^{\infty} \psi(x) e^{-ipx/\hbar} dx, \quad (21)$$

where the scale factor $\hbar = h/(2\pi)$, with $h = 6.6260701 \times 10^{-34}$, the Planck's constant. Clearly, $\psi(x)$ is the inverse Fourier transform of $\varphi(p)$, and together x and p form a pair of the so-called conjugate variables (Stam, 1959).

As a statistician, once I understood how the marginal distributions for x and p were constructed, I naturally asked for their joint distribution. This is where things become intriguing or puzzling to those of us who are trained to model non-deterministic relationships via probability, because (quantum) physicists' answer would be that there is no joint probability distribution for x and p —not that they are unknown, but that there cannot be one. Unlike the mystery of deep learning to statisticians—and its winning of the Nobel Prize in physics only makes it more intriguing or puzzling—I found good clues to the inadequacy of ordinary probability for dealing the quantum world by the very fact that its mathematical modeling involves non-commutative relationships, such as between operators or matrices.

Perhaps the easiest way to see potential complications with non-commutative relationship is to consider the problem of generalizing the notion of variance to co-variance with complex-valued variables. With real-valued random variables G and H having a joint distribution F , we know variance is the co-variance of a variable with itself, that

is, $V(G) = \text{Cov}(G, G)$. In other words, when we link variance with an inner product, *i.e.*, $V(G) = \langle G, G \rangle_F$, there is a natural extension for covariance by defining $\text{Cov}(G, H) = \langle G, H \rangle_F$. However, with the ordinary definition of the co-variance, this extension works only if the inner product is symmetric, that is, $\langle G, H \rangle_F = \langle H, G \rangle_F$, since $\text{Cov}(G, H) = \text{Cov}(H, G)$ in the *real* world.

This is where the complex world is, literally, more complex than the real world. For two complex-valued L^2 functions $u(y)$ and $v(y)$ on $y \in \Omega$, the inner product is not symmetric, because it is defined by

$$\langle u|v \rangle_\mu \equiv \int_\Omega \bar{u}(y)v(y)\mu(dy) \neq \langle v|u \rangle_\mu \equiv \int_\Omega \bar{v}(y)u(y)\mu(dy), \quad (22)$$

where \bar{u} is the complex conjugate of u , and μ is a baseline measure, which does not need to be a probabilistic measure. This non-commutative property is at the heart of quantum mechanics, as reviewed in the next Section. It can also be seen with matrix mechanics, since for any two matrices A and B or more broadly operators, in general $AB \neq BA$. The very fact that a regular joint probability specificity must render $\text{Cov}(u, v) = \text{Cov}(v, u)$ should remind us that whatever ‘joint specification’ of u and v we come up with, it will be more nuanced than a direct probabilistic distribution for $\{u, v\}$ whenever (22) rears its head. This phenomenon is not unique to the quantum world, since a similar situation happens with the notion of quasi-score functions, which can violate a symmetry requirement for genuine score functions, as reviewed in Appendix C.

However, this complication does not imply that probabilistic thinking is out the window. Because $\overline{\langle v|u \rangle_\mu} = \langle u|v \rangle_\mu$, we see that if we define $\text{Cov}(u, v) = \langle u|v \rangle_\mu$, then its magnitude, $|\text{Cov}(u, v)| = |\text{Cov}(v, u)|$ is symmetric. Therefore, as long as $|\text{Cov}(u, v)|$ is used as a measure of the magnitude of the co-variation between u and v , we can treat it as if it were the magnitude of a standard probabilistic co-variance. In other words, the concept, or at least the essence of *co-variance*, can be extended to non-probabilistic settings, and this extension perhaps can help our appreciation of HUP from a statistical perspective, as detailed in the next Section.

8. A lower resolution co-variation: co-variance of generating mechanisms

In the quantum world, we have seen that a particle’s position and momentum have their respectively well-defined probability distribution, and we can express $V(x) = \langle f|f \rangle_\mu$ and $V(p) = \langle g|g \rangle_\mu$, where $f(x) = x\psi(x)$ and $g(p) = p\phi(p)$. It is then mathematically tempting to define $\text{Cov}(x, p) = \langle f|g \rangle_\mu$ and $\text{Cov}(p, x) = \langle g|f \rangle_\mu$, using the notation of the previous section. This construction is problematic starting from the very notation $\text{Cov}(x, p)$, since it may suggest that we are measuring the co-variance between the position and momentum as *states*, which creates an epistemic disconnect with the understanding that a *joint statehood* of x and p does not exist or cannot be constructed in the quantum world.

However, x and p clearly have physical relationships. Indeed the so-called Stam’s uncertainty principle (Stam, 1959) establishes that

$$C^2V(x) - J(p) \geq 0 \quad \text{and} \quad C^2V(p) - J(x) \geq 0, \quad (23)$$

where $C = 4\pi$ for standard Fourier transform, and $C = 2/\hbar$ when we use the \hbar -scaled Fourier

transform (21). Here $J(p)$ is the Fisher information for the density of p , $f(p)$, that is,

$$J(p) = \int_{-\infty}^{\infty} \left[\frac{d \log f(p)}{dp} \right]^2 f(p) dp, \tag{24}$$

and similarly for $J(x)$. For readers who are unfamiliar with defining Fisher information for a density itself instead of its parameter, $J(p)$ is the same as the Fisher information for the location family $f(p - \theta)$, where θ shares the same state space as p (in the current case, the real line). In the same vein, the Cramér-Rao inequality can be applied to the density itself, which leads to $V(x) \geq J^{-1}(x)$ and $V(p) \geq J^{-1}(p)$. Consequently, as shown in Dembo (1990) and Dembo *et al.* (1991),

$$V(x)V(p) \geq C^{-2} = \frac{\hbar^2}{4}, \tag{25}$$

which is the same as the usual expression of HUP proved in Kennard (1927):

$$\Delta x \Delta p \geq \frac{\hbar}{2}, \tag{26}$$

where Δx and Δp denote respectively the standard deviation of x and p . Dembo (1990) and Dembo *et al.* (1991) also used (23) to prove that HUP implies the Cramér-Rao inequality.

The Stam’s uncertainty principle is elegant, and it reveals a kind of relationship between two marginal distributions that is not commonly studied in statistical literature, because it bypasses the specification of a joint distribution between x and p . However, this does not rule out—and indeed it suggests—that we can consider quantifying the relationships between the *mechanisms* that generate x and p . A mechanism can generate a single state, many states, or no states at all—which is equivalent to presenting itself as a whole—at any given circumstance, such as temporal instance. Hence quantifying relationships among mechanisms is a broader construct than that for the states they generate.

For statistical readers, a reasonable analogy is to think about the notion of *likelihood*. When we employ a likelihood, we can consider a single likelihood value (*e.g.*, at the MLE), several likelihood values (*e.g.*, likelihood ratio tests), or not any particular value but the likelihood function as a whole (*e.g.*, for Bayesian inference). By considering co-variations at the (resolution) level of mechanisms instead of states, we may find it less foreign to contemplate indeterminacy of relationship, such as between two sets—including empty ones—of the states generated by related mechanisms.

Of course, one may wonder if any relationship between two mechanisms itself can be indeterminable. The logical answer is yes, but fortunately for quantum mechanics we do not need go that far. As any useful quantum mechanics textbook (Landau and Lifshitz, 2013; Griffiths and Schroeter, 2018) teaches us, the position mechanism and momentum mechanism can be represented mathematically via the so-called position operator \hat{x} and momentum operator \hat{p} , to follow the notation in quantum mechanics, and they are tethered together when being applied to the same wave function $\psi(x)$ (in the position space⁵), that is

$$\hat{x} \circ \psi(x) = x\psi(x), \quad \text{and} \quad \hat{p} \circ \psi(x) = -i\hbar\psi'(x). \tag{27}$$

⁵One can define the operators equivalently in the conjugate momentum space via $\hat{p} \circ \varphi(p) = p\varphi(p)$ and $\hat{x} \circ \varphi(p) = i\hbar\varphi'(p)$, where the momentum wave function $\varphi(p)$ is the Fourier transform of $\psi(x)$ (Griffiths and Schroeter, 2018).

That is, the position operator acts on ψ by multiplying ψ with its argument, and the momentum operator acts on ψ by differentiating it, and multiplying it by $-i\hbar$, where $i = \sqrt{-1}$.

With these representations of the mechanisms, we can measure their co-variations induced by changing the state x in real line (as a univariate case) via the inner products, with respect to a common measure μ , typically Lebesgue measure. That is, we can define

$$\text{Cov}(\hat{x}, \hat{p}) = \langle \hat{x} \circ \psi | \hat{p} \circ \psi \rangle_{\mu} = -i\hbar \int_{-\infty}^{\infty} x \bar{\psi}(x) \psi'(x) dx; \tag{28}$$

$$\text{Cov}(\hat{p}, \hat{x}) = \overline{\text{Cov}(\hat{x}, \hat{p})} = i\hbar \int_{-\infty}^{\infty} x \psi(x) \bar{\psi}'(x) dx = -i\hbar \left(1 + \int_{-\infty}^{\infty} x \bar{\psi}(x) \psi'(x) dx \right). \tag{29}$$

Here the last equality is obtained by integration by parts and by using the fact that $|\psi(x)|^2$ is a probability density and that $x|\psi(x)|^2$ vanishes at $x = \pm\infty$ (because physicists assume the mean position is finite). Together, expressions (28)-(29) imply that

$$\text{Cov}(\hat{x}, \hat{p}) - \text{Cov}(\hat{p}, \hat{x}) = i\hbar, \tag{30}$$

which is also the consequence of the so-called *canonical commutation relation* (Griffiths and Schroeter, 2018),

$$\hat{x} \circ \hat{p} - \hat{p} \circ \hat{x} = i\hbar, \tag{31}$$

which holds because $\hat{x} \circ (\hat{p} \circ f(x)) - \hat{p} \circ (\hat{x} \circ f(x)) = i\hbar f(x)$ for any differentiable function f .

An immediate consequence of (30) is that the magnitude of the covariances between \hat{x} and \hat{p} is bounded below regardless of the form of the wave function $\psi(x)$. This is because for any complex number z , $|z|^2 \geq |\text{Im}(z)|^2 = |(z - \bar{z})/2i|^2$. Hence the identity (30) implies that

$$|\text{Cov}(\hat{x}, \hat{p})|^2 \geq \left[\frac{\text{Cov}(\hat{x}, \hat{p}) - \text{Cov}(\hat{p}, \hat{x})}{2i} \right]^2 = \frac{\hbar^2}{4}. \tag{32}$$

As reviewed in the next section, inequality (32) implies HUP in the form of (26), just as Stam's uncertainty principle does. For that purpose, it is worth pointing out that marginally,

$$\text{V}(\hat{x}) = \langle \hat{x} \circ \psi | \hat{x} \circ \psi \rangle_{\mu} = \int_{-\infty}^{\infty} x^2 \bar{\psi}(x) \psi(x) dx = \int_{-\infty}^{\infty} x^2 |\psi(x)|^2 dx; \tag{33}$$

$$\text{V}(\hat{p}) = \langle \hat{p} \circ \psi | \hat{p} \circ \psi \rangle_{\mu} = \hbar^2 \int_{-\infty}^{\infty} \bar{\psi}'(x) \psi'(x) dx = \int_{-\infty}^{\infty} p^2 |\varphi(p)|^2 dp, \tag{34}$$

where the last equation in (34) is due to the fact that $\varphi(p)$ is the (\hbar -scaled) Fourier transformation of $\psi(x)$, as given in (21). These two equalities tell us that when we consider either the position or the momentum by itself, its mechanism-level variance, $\text{V}(\hat{x})$ or $\text{V}(\hat{p})$, and the state-level variance, $\text{V}(x)$ or $\text{V}(p)$, are the same. This renders the unity between the mechanism-level representation (as a distribution or operator) and the state-level representation (as an observable or latent variable), a distinction seldom made conceptually under the ordinary probability framework. However, this distinction can be crucial once we go outside the regular probability framework, as in the current context of measuring co-variations between the position and momentum.

9. Bounding co-variations: A commonality of uncertainty principles

With co-variances constructed broadly, we can study the similarities and differences between inequality (14) and the Cramér-Rao inequality, as well as their intrinsic connections with HUP. Specifically, both inequalities are based on bounding joint variations of two random objects, say, G and H , by their marginal variations. For (14), under the unbiasedness assumptions and using the notation given in Section 5, if we write $G = \delta_{\hat{Q}}$ and $H = \hat{\delta}_{\hat{Q}}$, then inequality (14) is the consequence of (omitting subscript s):

$$\text{Cov}^2(G, H) \leq V(H) [V(G) - R^{\text{opt}}]. \quad (35)$$

For the Cramér-Rao inequality, we can take the same $G = \delta_{\hat{Q}} = \hat{Q} - Q$, where \hat{Q} is an unbiased estimator for Q . We then let $H = S(\theta|D)$, the score function from a sampling model of our data D , $f(D|\theta)$, with $Q = Q(\theta)$. It is known that the Cramér-Rao inequality is the same as (*e.g.*, Lehmann and Casella, 2006)

$$[Q'(\theta)]^2 = \text{Cov}^2(G, H) \leq V(H)V(G), \quad (36)$$

where $Q'(\theta)$ is the derivative for $Q(\theta)$. (When $Q(\theta)$ is not differentiable, we can apply the bound given by Chapman and Robbins (1951)) in terms of likelihood ratio or elasticity.)

Evidently, inequality (36) is an application of the Cauchy-Schwartz inequality. In contrast, inequality (35) delivers a more precise bound because of the subtraction of the term R^{opt} . Indeed, inequality (35) is often an equality because the condition in (II) of Theorem 1 frequently holds in practice. Given the two inequalities share the same type of G , the difference must be attributable to something distinctive between the two H 's. Whereas both H 's have zero expectation, the first $H = \hat{\delta}_{\hat{Q}}$ is a *statistic*, required to be a function of data D only. In contrast, the second $H = S(\theta|D)$ is a *random function*, depending on both data D and the unknown θ . Since the actual error $\delta_{\hat{Q}} = \hat{Q} - Q(\theta)$ is also a random function, the second H can co-vary with G to a greater extent than the first H can. Consequently, $\text{Cov}^2(G, H)$ can reach a looser upper bound in (36) than in (35). As an illustrative example, for estimating the normal mean under $N(\mu, \sigma^2)$, $Q = \bar{X}_n - \mu$ and $H = S(\mu|X) = n(\bar{X}_n - \mu)/\sigma^2 = nG/\sigma^2$, and hence (36) becomes equality, whereas such an H is clearly not permissible for (35).

Nevertheless, both inequalities reveal the tension between individual variations—features of their respective marginal distributions—and their co-variation, which reflects their relationships, probabilistic or not. For (36), in order to keep $\text{Cov}^2(G, H)$ at the value of $[Q'(\theta)]^2 > 0$, the two variances $V(H)$ and $V(G)$ cannot be simultaneously small to an arbitrary degree, just as a rectangle cannot have arbitrarily small sides simultaneously when its area is bounded away from zero. This restriction leads to the Cramér-Rao lower bound. In (36), we purposefully write the Fisher information as the variance of the score function instead of the expectation of its negative derivative. The variance expression makes it clearer the co-variation essence of Cramér-Rao inequality, and draws a direct parallel with the inequality underlying HUP.

Specifically, using the notation and the inequality (32) of Section 8 and taking $G = \hat{x}$ and $H = \hat{p}$, we have

$$\frac{\hbar^2}{4} \leq |\text{Cov}(G, H)|^2 \leq V(H)V(G), \quad (37)$$

Comparing (37) with (36), we see that the Cramér-Rao bound and the Heisenberg uncertainty principle are consequences of essentially the same statistical phenomenon, that is, two marginal variances necessarily compete with each for being arbitrarily small, when the corresponding covariance is constrained in magnitude from below.

In contrast, for (35), the trade-off is between the covariance and one of the marginal variances. To see this clearly, we can assume $V(H) = 1$, which does not offend the assumption that $E(H) = 0$. Inequality (35) then becomes

$$\text{Cov}^2(G, H) \leq V(G) - R^{\text{opt}} = R_G, \quad (38)$$

where R_G is the regret of G . On the surface, the changes of covariance and $V(G)$ appear to be coordinated instead of in competition, because the larger $\text{Cov}^2(G, H)$, the larger $V(G)$. The reverse holds when the inequality is equality (which often is the case), and more broadly larger $V(G)$ —and hence larger regret—at least allows more room for $\text{Cov}^2(G, H)$ to grow. But this is exactly where the tension lies when we want to improve both the learning and error assessment; improving learning means to reduce R_G and hence have a *smaller* $V(G)$, but improving error assessment requires a *larger* $\text{Cov}^2(G, H)$.

10. Elementary mathematics, advanced statistics, and inspiring philosophy

Mathematically, the proof of either (36) or (37) is elementary, yet the implications of either inequality, as we know, are profound. Similarly, the inequality (35) is built upon equally elementary mathematics, and the work of Bates *et al.* (2024) has already suggested its potential impact. However, many more studies remain, particularly regarding alternative loss functions, where the relevance of error assessment may not align with covariance. From a probabilistic standpoint, a thorough theoretical exploration of the relevance of an error assessor, $\hat{\delta}$, for the true error δ should involve investigating the joint distribution of $\hat{\delta}$ and δ . In this context, irrelevance can be characterized by the independence between $\hat{\delta}$ and δ .

On a broader level, formulating a general trade-off between learning and error assessment remains a complex task. This challenge stems from the need to define and measure the actual information utilized during learning and to identify relevant replications when assessing errors. Both ‘information’ and ‘learning’ are elusive notions, having taken on numerous interpretations throughout history, many of which require a refined understanding. For instance, even in the case of classical likelihood inference within parametric models, the role of conditioning in error assessment continues to provoke theoretical and practical debates.

I was reminded of this reality by an astrostatistics project involving correcting conceptual and methodological errors in astrophysics for conducting model fitting and goodness-of-fit assessment via the popular C-statistics, which is the likelihood ratio statistic under a Poisson regression model (Cash, 1979). When the project started, I naively believed that it would be merely an exercise of applying classical likelihood theory and methods, perhaps with some clever computational tricks or approximations to render them practically efficient and hence appealing to astrophysicists.

As reported in Chen *et al.* (2024), however, the issue about whether one should condition on the MLE itself or not in the context of goodness-of-fit testing, is a rather nuanced one. The issue is closely related to the issue of conditioning on ancillary statistics,

since for testing distributional shape, the parametric parameters are *nuisance objects* (as termed in Meng, 2024) and their MLE can be intuitively perceived as locally ancillary (Cox, 1980; Severini, 1993) because the distribution shape of the MLE will be normal to the first order (under regularity conditions) despite the shape of the distribution being tested. However, it is not exactly ancillary, and to decide when conditioning is beneficial (*e.g.*, leading to a more powerful test) in any sample settling is not a straightforward matter. Higher order asymptotics can help provide insight, but communicating them intuitively is a tall order even for statisticians, let alone for astrophysicists or any scientists (including data scientists). However, regardless of whether low-level mathematics or high/tall order of statistics are involved, the ultimate challenge of contemplating and formulating uncertainty principles is epistemological, or even metaphysical. For readers interested in philosophical contemplation—and I'd expect that statisticians should be in that group because statistics is essentially *applied epistemology*⁶, I highly recommend the over 50 pages entry titled “The Uncertainty Principle” by Hilgevoord and Uffink (2024) in *The Stanford Encyclopedia of Philosophy*.⁷ It is an erudite and thought-provoking essay about the intellectual journey of Heisenberg's uncertainty principle. Even or perhaps especially the name “uncertainty principle” has an interesting story behind it, because initially the name did not contain either ‘uncertainty’ or ‘principle’.

As Hilgevoord and Uffink (2024) discussed, the term *uncertainty* has multiple meanings, and it is not obvious in which sense the phenomena revealed by Heisenberg (1927) qualifies as ‘uncertainty’; indeed, historically terms such as “inaccuracy, spread, imprecision, indefiniteness, indeterminateness, indeterminacy, latitude” were used by various writers for what is now known as HUP. More intriguingly, Heisenberg did not postulate the finding as any kind of *principle*, but rather as *relations*, such as “inaccuracy relations” or “indeterminacy relations”. The discussions in Section 8 certainly reflect the relational nature of HUP, because it is fundamentally about the co-variation of position and momentum at the mechanism level.

The entry by Hilgevoord and Uffink (2024) invites readers to consider a fundamental question that underpins these onomasiological reflections: Is the HUP a mere epistemic constraint, or a metaphysical limitation in nature? Unsurprisingly, this question is a source of ongoing dispute among philosophers of physics and even among physicists themselves. The most well-known historical debates are Heisenberg and Bohr's Copenhagen interpretation emphasizing the metaphysical indeterminacy, and the contrasting deterministic interpretation developed by de Broglie and Bohm, known as Bohmian mechanics (Hilgevoord and Uffink, 2024).

Given I have already greatly exceeded the deadline to submit this essay, I will refrain from revealing any further thrills provided in Hilgevoord and Uffink (2024), such as more recent debates about HUP, leaving readers to enjoy their own treasure hunt. But I will

⁶This was a characterization given by philosopher Hanti Lin during the JSM 2024, where Hanti and I co-organized a session where each philosopher presented for 20 minutes followed by a 15-min discussion by a statistician, and there were three pairs in total. (I made a mistake that embodied the statisticians' modesty: the estimated room size I provided to the JSM meeting department had an unacceptably negative bias.)

⁷SEP is simply a fountain of afflatus and a *Who's Who* in philosophy. Indeed SEP was where I came across Hanti Lin's 115-page entry on “Bayesian Epistemology” (Lin, 2024a), and led to my invitation to Hanti to serve as a co-editor to establish the “Meta Data Science” column (<https://hdsr.mitpress.mit.edu/meta-data-science>) for *Harvard Data Science Review*.

mention that this question has prompted me to wonder whether inequality (14) also suggests that any effort to assess the actual error is antithetical to probabilistic learning.

This is because the crux of probabilistic learning—unlike deterministic approaches, such as solving algebraic equations—lies in using distributions as our fundamental mathematical vehicles for carrying our states of knowledge (or lack thereof) and for transporting data into information that furthers learning. From this distributional perspective, assessing the actual error means to assess the *distribution* of the actual error, which is all we need to, for example, provide the usual confidence regions. It does suffer from the leap of faith problem as discussed in Section 4, but then that is a universal predicament to any form of empirical learning, as far as I can imagine.

11. From uncertainty principles to happy marriages...

A further inspiration from Hilgevoord and Uffink (2024) is its discussion on the relationship between the original semi-quantitative argument made by Heisenberg (1927) and the mathematical formalism established by Kennard (1927). Kennard's inequality (26) is precise, but can be perceived as narrow, for instance, in its reliance on standard deviation to describe “uncertainty.” A similar limitation applies to inequality (14), which assesses relevance through linear correlation, a measure that surely is not universally appropriate for capturing the notion of relevance.

More broadly, much remains to be examined regarding the trade-offs between the flexibility of qualitative frameworks, which embrace the nuances and ambiguities of natural language, and the rigor of quantitative formulations, which offer the precision of mathematical language but often at the risk of being overly restrictive or idealized. Reflecting on these trade-offs is essential to learning. Statisticians and data scientists, in particular, can draw from centuries of philosophical inquiry into epistemology, as exemplified by the discussions surrounding the HUP and the like. In truth, when thoughtfully practiced, data science embodies—or ought to embody—a harmonious blend of quantitative and qualitative thinking and reasoning. This was the central theme of my *Harvard Data Science Review* editorial, “Data Science: A Happy Marriage of Quantitative and Qualitative Thinking?” (Meng, 2021), inspired by Tanweer *et al.* (2021)'s compelling article, “Why the Data Revolution Needs Qualitative Thinking.” Maintaining this harmony, akin to sustaining a functioning marriage, requires commitment from all parties and a willingness to compromise. Ultimately, it calls for the wisdom to recognize that individual fulfillment and happiness—whether in marriage, mentorship, or mind melding or mating—depends profoundly on collective well-being. Professor Rao certainly embodied this wisdom.

I vividly recall my first visit to Pennsylvania State University as a seminar speaker, shortly after Professor Rao's 72nd birthday on September 10, 1992. During the seminar lunch, Professor Rao graciously joined us. We—students and early-career researchers (myself included, back when my hair was dense almost surely everywhere)—felt honored by his presence. All questions naturally revolved around statistics, except for one that made us all chuckle: “Professor Rao, how does one live a long and happy life?”

Without missing a beat, and with his characteristic paced, confident cadence, Rao replied, “Keep your wife happy.”

12. A prologue or an invitation

For those who would like this article to conclude with a statistical Q&A: During the elevator ride following my seminar, which carried the seemingly oxymoronic title “A Bayesian p -value” (a deliberate contrast to the title of Meng (1994)), Professor Rao turned to me and asked, “Do people still use p -values?” To which I responded. . .

Well, I’ll leave that as a missing data point, inviting you to impute your own favorite answer. Alternatively, if you prefer, find a deliberately embedded mathematical (but petty) error in this article and exchange it for the answer by emailing `meng@stat.harvard.edu` (as long as God permits me to respond).

Acknowledgments

I am deeply grateful to physicists Aurore Courtoy, Louis Lyons, Thomas Junk, and Pavel Nadolsky, as well as statistician Yazhen Wang, for their careful and patient explanations regarding the non-existence of a joint probabilistic distribution of a particle’s position and momentum. I am equally indebted to Hanti Lin for elucidating the philosophical debates surrounding the Heisenberg Uncertainty Principle.

My thanks also extend to editor Bhramar Mukherjee, to whom I owe a profound debt, and to Peter Bickel, Joe Blitzstein, Radu Craiu, Walter Dempsey, Benedikt Höltgen, Peter McCullagh, Pavlos Msaouel, Steve Stigler, Robert Tibshirani, Théo Voldoire, and Bob Williamson for collectively providing insightful comments and sharing relevant literature—some of which may inspire a sequel to this essay.

I also thank Julie Vu and Sicheng Zhou for their meticulous proofreading efforts; naturally, any remaining errors are entirely my own (though I wish they weren’t!). Finally, I acknowledge partial financial support from the NSF during the period when this essay was conceived and completed.

References

- Abba, M. A., Williams, J. P., and Reich, B. J. (2024). A Bayesian shrinkage estimator for transfer learning. *arXiv:2403.17321*.
- Bates, S., Hastie, T., and Tibshirani, R. (2024). Cross-validation: what does it estimate and how well does it do it? *Journal of the American Statistical Association*, **119**, 1434–1445.
- Berger, J., Meng, X. L., Reid, N., and Xie, M. (2024). *Handbook of Bayesian, Fiducial, and Frequentist Inference*. CRC Press.
- Blitzstein, J. K. and Hwang, J. (2014). *Introduction to Probability*. CRC Press, Boca Raton, FL, 1st edition.
- Casella, G. and Berger, R. (2024). *Statistical Inference*. CRC Press.
- Cash, W. (1979). Parameter estimation in astronomy through application of the likelihood ratio. *The Astrophysical Journal*, **228**, 939.
- Chapman, D. G. and Robbins, H. (1951). Minimum variance estimation without regularity assumptions. *The Annals of Mathematical Statistics*, **22**, 581–586.

- Chen, Y., Li, X., Meng, X. L., van Dyk, D. A., Bonamente, M., and Kashyap, V. (2024). Boosting C-statistics in astronomy via conditioning: More power, less computation. Technical report, Department of Statistics, University of Michigan.
- Cox, D. R. (1980). Local ancillarity. *Biometrika*, **67**, 279–286.
- Craiu, R. V., Gong, R., and Meng, X. L. (2023). Six statistical senses. *Annual Review of Statistics and Its Application*, **10**, 699–725.
- Dembo, A. (1990). Information inequalities and uncertainty principles. *Department of Statistics, Stanford University, Stanford, CA, Technical Report*, **75**.
- Dembo, A., Cover, T. M., and Thomas, J. A. (1991). Information inequalities and uncertainty principles. *IEEE Transactions on Information Theory*, **37**, 1501–1518.
- Dempster, A. P. (1963). Further examples of inconsistencies in the fiducial argument. *The Annals of Mathematical Statistics*, **34**, 884–891.
- Efron, B. (2020). Prediction, estimation, and attribution. *International Statistical Review*, **88**, S28–S59.
- Efron, B. and Hinkley, D. V. (1978). Assessing the accuracy of the maximum likelihood estimator: Observed versus expected fisher information. *Biometrika*, **65**, 457–483.
- Gelman, A. and Betancourt, M. (2013). Does quantum uncertainty have a place in everyday applied statistics. *Behavioral and Brain Sciences*, **36**, 285.
- Gong, R. and Meng, X. L. (2021). Judicious judgment meets unsettling updating: Dilation, sure loss, and simpson’s paradox. *Statistical Science*, **36**, 169–214. Discussion article with rejoinder.
- Griffiths, D. J. and Schroeter, D. F. (2018). *Introduction to Quantum Mechanics*. Cambridge University Press.
- Heisenberg, W. (1927). Über den anschaulichen inhalt der quantentheoretischen kinematik und mechanik. *Zeitschrift für Physik*, **43**, 172–198.
- Hilgevoord, J. and Uffink, J. (2024). The uncertainty principle. In Zalta, E. N. and Nodelman, U., editors, *The Stanford Encyclopedia of Philosophy*. Stanford University. Spring 2024 edition.
- Hillery, M., O’Connell, R. F., Scully, M. O., and Wigner, E. P. (1984). Distribution functions in physics: Fundamentals. *Physics Reports*, **106**, 121–167.
- Hoeffding, W. (1940). Maßtabinvariante Korrelatiostheorie. *Schriften des Mathematischen Instituts und des Instituts für Angewandte Mathematik der Universität Berlin*, **5**, 179–233.
- Kennard, E. H. (1927). Zur quantenmechanik einfacher bewegungstypen. *Zeitschrift für Physik*, **44**, 326–352.
- Landau, L. D. and Lifshitz, E. M. (2013). *Quantum Mechanics: Non-relativistic Theory*, volume 3. Elsevier.
- Le Cam, L. (1956). On the asymptotic theory of estimation and testing hypotheses. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, volume 3, pages 129–157. University of California Press.
- Lehmann, E. L. and Casella, G. (2006). *Theory of Point Estimation*. Springer Science & Business Media.

- Lin, H. (2024a). Bayesian epistemology. In Zalta, E. N. and Nodelman, U., editors, *The Stanford Encyclopedia of Philosophy*. 2024 Edition, originally published 2022.
- Lin, H. (2024b). To be a Frequentist or Bayesian? Five positions in a spectrum. *Harvard Data Science Review*, **6**. <https://hdsr.mitpress.mit.edu/pub/axvcupj4>.
- Liu, K. and Meng, X. L. (2014). Comment: A fruitful resolution to simpson's paradox via multiresolution inference. *The American Statistician*, **68**, 17–29.
- Liu, K. and Meng, X. L. (2016). There is individualized treatment. Why not individualized inference? *Annual Review of Statistics and Its Application*, **3**, 79–111.
- Lorce, C. and Pasquini, B. (2011). Quark wigner distributions and orbital angular momentum. *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, **84**, 014015.
- McCullagh, P. (1999). Discussion on some statistical heresies. *Journal of the Royal Statistical Society: Series D (The Statistician)*, **48**, 34–35.
- McCullagh, P. and Nelder, J. A. (1989). *Generalized Linear Models*, volume 37 of *Monographs on Statistics and Applied Probability*. Chapman & Hall/CRC, London, 2nd edition.
- Meng, X. L. (1994). Posterior predictive p -values. *The Annals of Statistics*, **22**, 1142–1160.
- Meng, X. L. (2018). Statistical paradises and paradoxes in big data (I): law of large populations, big data paradox, and the 2016 us presidential election. *The Annals of Applied Statistics*, **12**, 685–726.
- Meng, X. L. (2021). Data science: A happy marriage of quantitative and qualitative thinking? *Harvard Data Science Review*, **3**. <https://hdsr.mitpress.mit.edu/pub/pger71uh>.
- Meng, X. L. (2024). A BFFer's exploration with nuisance constructs: Bayesian p -value, H-likelihood, and Cauchyanity. In *Handbook of Bayesian, Fiducial, and Frequentist Inference*, Eds J. Berger, XL. Meng, N. Reid and M. Xie, pages 161–187. Chapman and Hall/CRC.
- Rao, C. R. (1945). Information and the accuracy attainable in the estimation of statistical parameters. *Bulletin of the Calcutta Mathematical Society*, **37**, 81–91.
- Rao, C. R. (1962). Efficient estimates and optimum inference procedures in large samples. *Journal of the Royal Statistical Society: Series B (Methodological)*, **24**, 46–63.
- Severini, T. A. (1993). Local ancillarity in the presence of a nuisance parameter. *Biometrika*, **80**, 305–320.
- Stam, A. J. (1959). Some inequalities satisfied by the quantities of information of fisher and shannon. *Information and Control*, **2**, 101–112.
- Tanweer, A., Gade, E. K., Krafft, P., and Dreier, S. (2021). Why the data revolution needs qualitative thinking. *Harvard Data Science Review*, **3**. <https://hdsr.mitpress.mit.edu/pub/u9s6f22y>.
- Tóth, G. and Fröwis, F. (2022). Uncertainty relations with the variance and the quantum fisher information based on convex decompositions of density matrices. *Physical Review Research*, **4**, 013075.
- Tóth, G. and Petz, D. (2013). Extremal properties of the variance and the quantum fisher information. *Physical Review A—Atomic, Molecular, and Optical Physics*, **87**, 032324.
- Wald, A. (1943). Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Transactions of the American Mathematical Society*, **54**, 426–482.

Wang, Y. (2022). When quantum computation meets data science: Making data science quantum. *Harvard Data Science Review*, 4. <https://hdsr.mitpress.mit.edu/pub/kpn45eyx>.

APPENDIX

Appendix A: Derivations for the regression example in section 3

In general, the weighted estimate of θ can be written as

$$\hat{\theta}_w = \frac{\sum_{i=1}^n w_i X_i Y_i}{\sum_{i=1}^n w_i X_i^2},$$

with OLS corresponding to choosing $w_i = 1$ and BLUE given by $w_i = \sigma_i^{-2}$, for all i . Conditioning on \mathbf{X} but for notational simplicity we suppress the conditioning notation in all expectations below, we have

$$V(\hat{\theta}_w) = \frac{\sum_{i=1}^n w_i^2 X_i^2 \sigma_i^2}{[\sum_{i=1}^n w_i X_i^2]^2} = \frac{T_{w,\sigma}}{T_w^2}.$$

Let $\hat{r}_{w,j} = Y_j - \hat{\theta}_w X_j$. Because $E(\hat{r}_{w,j}) = 0$, to calculate ρ , we only need to calculate

$$\begin{aligned} E[\hat{\theta}_w(Y_j - \hat{\theta}_w X_j)] &= \frac{\sum_{i=1}^n w_i X_i E[Y_i Y_j]}{T_w} - \frac{E[\sum_{i=1}^n w_i X_i Y_i]^2 X_j}{T_w^2} \\ &= \frac{\sum_{i=1}^n w_i X_i [\text{Cov}(Y_i, Y_j) + \theta^2 X_i X_j]}{T_w} - \frac{[\sum_{i=1}^n w_i^2 X_i^2 \sigma_i^2 + \theta^2 T_w^2] X_j}{T_w^2} \\ &= \frac{(\theta^2 T_w + w_j \sigma_j^2) X_j}{T_w} - \frac{[T_{w,\sigma} + \theta^2 T_w^2] X_j}{T_w^2} = \frac{X_j}{T_w} \left[w_j \sigma_j^2 - \frac{T_{w,\sigma}}{T_w} \right]; \end{aligned}$$

and

$$\begin{aligned} V(\hat{r}_{w,j}) &= V \left[\frac{\sum_{i=1}^n w_i X_i (X_i Y_j - X_j Y_i)}{T_w} \right] = T_w^{-2} V \left[\sum_{i \neq j}^n w_i X_i (X_i Y_j - X_j Y_i) \right] \\ &= T_w^{-2} E \left\{ V \left[\sum_{i \neq j}^n w_i X_i (X_i Y_j - X_j Y_i) | Y_j \right] \right\} + V \left\{ E \left[\sum_{i \neq j}^n w_i X_i (X_i Y_j - X_j Y_i) | Y_j \right] \right\} \\ &= T_w^{-2} \left\{ \left[X_j^2 \sum_{i \neq j}^n w_i^2 X_i^2 \sigma_i^2 \right] + V \left[\sum_{i \neq j}^n w_i X_i^2 Y_j \right] \right\} \\ &= T_w^{-2} \left\{ \left[X_j^2 (T_{w,\sigma} - w_j^2 X_j^2 \sigma_j^2) \right] + [T_w - w_j X_j^2]^2 \sigma_j^2 \right\} \\ &= T_w^{-2} \left\{ X_j^2 T_{w,\sigma} + \sigma_j^2 [T_w^2 - 2T_w w_j X_j^2] \right\}. \end{aligned}$$

Putting all the pieces together, we have

$$\text{Corr}(\hat{\theta}_w, \hat{r}_{w,j}) = \frac{X_j (w_j \sigma_j^2 T_w - T_{w,\sigma})}{\sqrt{T_{w,\sigma} [X_j^2 T_{w,\sigma} + \sigma_j^2 (T_w^2 - 2T_w w_j X_j^2)]}}, \quad j = 1, 2. \quad (39)$$

For $n = 2, j = 1$, expression (39) simplifies to the desired (4) because

$$\begin{aligned} \text{Corr}(\hat{\theta}_w, r_{w,1}) &= \frac{X_1 X_2^2 w_2 (w_1 \sigma_1^2 - w_2 \sigma_2^2)}{\sqrt{[X_1^2 w_2^2 X_2^2 \sigma_2^2 + w_2^2 X_2^4 \sigma_1^2][w_1^2 X_1^2 \sigma_1^2 + w_2^2 X_2^2 \sigma_2^2]}} \\ &= \frac{X_1 |X_2| (w_1 \frac{\sigma_1}{\sigma_2} - w_2 \frac{\sigma_2}{\sigma_1})}{\sqrt{[X_1^2 \sigma_1^{-2} + X_2^2 \sigma_2^{-2}][w_1^2 X_1^2 \sigma_1^2 + w_2^2 X_2^2 \sigma_2^2]}}. \end{aligned}$$

To calculate the relative regret (RR), we have

$$V(\hat{\theta}_w) = V \left[\frac{\sum_{i=1}^n w_i X_i Y_j}{T_w} \right] = \frac{w_1^2 X_1^2 \sigma_1^2 + w_2^2 X_2^2 \sigma_2^2}{[w_1 X_1^2 + w_2 X_2^2]^2}, \tag{40}$$

which also implies, by taking $w_i \propto \sigma_i^{-2}$,

$$V(\hat{\theta}_{\text{BLUE}}) = \frac{1}{(X_1^2 \sigma_1^{-2} + X_2^2 \sigma_2^{-2})}. \tag{41}$$

Putting together (40) and (41) yields the desired (5).

Appendix B: Derivation of (11) in section 4

Because $\hat{\delta}^2$ and δ^2 are independent given $\theta = \{\mu, \sigma^2\}$ and hence $\text{Cov}(\hat{\delta}^2, \delta^2 | \mu, \sigma^2) = 0$, we see over the joint replication,

$$\text{Cov}(\hat{\delta}^2, \delta^2) = E \left[\text{Cov}(\hat{\delta}^2, \delta^2 | \mu, \sigma^2) \right] + \text{Cov} \left[E(\hat{\delta}^2 | \mu, \sigma^2), E(\delta^2 | \mu, \sigma^2) \right] = \frac{1}{n^2} V(\sigma^2),$$

as long as the prior distribution for $\theta = \{\mu, \sigma^2\}$ is proper. Furthermore, conditioning on $\theta = \{\mu, \sigma^2\}$, $\delta^2 \sim \sigma^2 \chi_1^2/n$ and $\hat{\delta}^2 \sim \sigma^2 \chi_{n-1}^2/[n(n-1)]$ (where the two chi-square variables are independent of each other), we have

$$\begin{aligned} V(\hat{\delta}^2) &= E \left[V(\hat{\delta}^2 | \mu, \sigma^2) \right] + V \left[E(\hat{\delta}^2 | \mu, \sigma^2) \right] = \frac{2}{(n-1)n^2} E(\sigma^4) + \frac{1}{n^2} V(\sigma^2); \\ V(\delta^2) &= E \left[V(\delta^2 | \mu, \sigma^2) \right] + V \left[E(\delta^2 | \mu, \sigma^2) \right] = \frac{2}{n^2} E(\sigma^4) + \frac{1}{n^2} V(\sigma^2). \end{aligned}$$

Consequently, we see over the joint replication,

$$\text{Corr}(\hat{\delta}^2, \delta^2) = \frac{V(\sigma^2)}{\sqrt{2(n-1)^{-1} E(\sigma^4) + V(\sigma^2)} \sqrt{2 E(\sigma^4) + V(\sigma^2)}},$$

which yields (11) because $E(\sigma^4) = V(\sigma^2) + [E(\sigma^2)]^2$.

Appendix C: A quasi-score analogy for understanding the lack of joint probability

For statistically oriented readers, an instructive—though far from being perfect—analogy to the issue of the non-existence of a probabilistic model due to violations of symmetry or commutativity is the generalization from likelihood inference via the score function to estimation based on quasi-score functions. The correct score function, when available, provides the most efficient inference asymptotically (under regularity conditions). However, specifying the correct data-generating model often requires more information and resources than we typically possess.

In contrast, a quasi-score function only requires the specification of the first two moments of the data-generating model. This makes it a more practical and robust alternative to exact model-based inference, particularly in the presence of model misspecification. However, this robustness comes at the cost of reduced efficiency, reflecting the trade-off inherent in this approach.

Broadly speaking there are three types of pseudo scores: (I) those that are equivalent to the actual score; (II) those that are not equivalent to the actual score, but are equivalent to the score from a misspecified data generating model, and (III) those that cannot be derived from any probabilistic model.

Type (III) exists because any (differentiable) authentic score vector $(S_1(\theta), \dots, S_d(\theta))^T$ for a d -dimension parameter $\theta = (\theta_1, \dots, \theta_d)^T$ must satisfy

$$\frac{\partial S_i(\theta)}{\partial \theta_j} = \frac{\partial S_j(\theta)}{\partial \theta_i}, \quad \forall i, j = 1, \dots, d, \quad (42)$$

because the corresponding (observed) Fisher information matrix, $-\frac{\partial^2 S(\theta)}{\partial \theta^2}$, is symmetric. However, even some most innocent looking quasi-scores, such as for certain 2×2 contingency tables, the symmetry requirement of (42) can be easily violated, as demonstrated in Chapter 9 of McCullagh and Nelder (1989), which is an excellent source for understanding quasi scores and estimation equations in general.

The fact that violating the symmetry condition (42) rules out the possibility of being an actual score may help some of us imagine how the lack of symmetry or commutativity might rule out the existence of a probability specification, at least from a mathematical perspective. Furthermore, just as one can generalize from likelihood to quasi-likelihood of many shapes and forms—again see McCullagh and Nelder (1989)—the non-existence of a probabilistic distribution does not prevent us from forming quasi-distributions for various purposes, such as the Wigner quasiprobability distribution, which permits negative values, for position and momentum (x, p) (Hillery *et al.*, 1984; Lorce and Pasquini, 2011). Whether the mechanism-level covariances as given in (28)-(29) have the same magnitude as that from the Wigner quasiprobability distribution will be left as a homework exercise.