The Missing Theorem: Can sampling error be estimated?

Michael C. Nelson

RTI International

{michaelnelson} @ rti.org

Abstract

If statistics is about sorting information from error, then roughly half of the theory of statistical estimation is missing. Estimation theory, at its beginning, paid sampling error (also known as random error or discrepancy, as opposed to bias) little consideration, other than to state that it exists and must be overcome. Nowhere, it seems, did our statistical forebears consider so basic a question as: “Is it possible to estimate the size and sign of sampling error in a parameter estimate, using only the sample from which it was computed?” Once the issue did appear in the literature, it seems to have been immediately and summarily answered in the negative, offering neither a proof nor a citation to a proof—a status that has remained unchanged and, apparently, unquestioned. In this paper, I argue that the omission of a theorem on sampling error estimation diminishes estimation theory, and that our collective blind spot in this respect represents a persistent limitation on our reason and imagination. I propose that a good start would be to restate the question as a scientific query, in unambiguous and falsifiable terms, and I provide my own effort to that end.

1 Introduction

It is widely understood that sampling error, defined as the deviation of a statistic from its parameter that arises from observing only a subset of the population, cannot

be estimated from a single, isolated sample. It is universally unknown why this should be so. Look, as I have, deep into our collective literature and you will find no hint that the question even occurred to those who first labored to piece together the field of statistics from disparate mathematical theorems and scientific applications. Which is fine: they had a on their plates.

Yet, in the ensuing century, we have not just perpetuated their oversight, we have elevated it to outright fallacy: in the modern literature, it is not at all difficult to find confident, if unsourced, assertions that sampling error cannot be estimated from a single sample. No proofs are presented, no theorems expounded, only the claim and, perhaps, a hint of circular reasoning. In this paper, I argue that this claim, that we cannot estimate sampling error outside of meta-analyses or multi-sample studies, is neither self-evident nor a logical consequence of estimation theory.¹ It is, at best, an axiom, although to be an axiom it must be stated as such; otherwise, it is just an ordinary assumption. Axioms, also, may be discarded.

In the first part of this paper, I provide historical context for the claim of sampling error estimation's impossibility. I do so chiefly through the lens of Fisher's 1925 "Theory of Statistical Estimation." I do not mean to imply that this particular paper, or Fisher himself, is or ever was the absolute standard for understanding statistical inference. It is only that the paper establishes the general understanding of statistical estimation at that time, and it serves as an early, influential example of the paradigm within which we work today.²

The next section reviews and responds to modern arguments against sampling error estimation—or would, if there were any. Instead, in the absence of explicit arguments, I dissect authors' implicit arguments as well as several I anticipate. The third section counters what may be the most convincing argument

¹ It has occurred to me, belatedly, that what means for a theorem to be "missing" may not be self-evident, and that the reason for wanting to find it might be equally elusive. I took these points for granted in my early drafts, and my argument proceeds without making them. I have now added an explanatory note in the Appendix. The short version is: researchers have repeated a claim about the nature of the world without even being tempted to verify it, despite their having devoted their careers to skeptical inquiry and discovery, and our understanding of statistics is diminished for it.

² It is also relevant that the paper's author was famously antagonistic and therefore unlikely to have let slide the opportunity to criticize, or at least correct, any claims by his peers on the topic of sampling error about which he disagreed.

that a theorem proving the impossibility of sampling error estimation is unnecessary: if it were essential, or even valuable, one of the many great minds in statistics and related fields over the last century would surely have at least said so. I suggest a number of potential explanations for the ongoing oversight.

Finally, I introduce a theoretical framework for conceptualizing sampling error estimation as a claim to be proven or disproven. The main thrust of the framework is simply that sampling error estimators are easy to find: every estimator of every parameter also estimates sampling error, as it must covary to some degree with both quantities. Statisticians have, of course, designed statistics that maximize the parameter information from the sample while minimizing the impact of sampling error. I suggest that, if we could construct a variable that covaries far more with sampling error than with the parameter, even as sample size increases, then its value when computed from the same data as a conventional estimator ought to serve as an indicator of sampling error in the conventional estimator. The collection of ideas I put forward in this regard, though minimal, provides a basis for conceiving of the properties of a useful sampling error covariate not as abstractions but as concrete and testable.

2 Foundations for missing a theorem

The following, I believe, is a fair distillation of the most elementary points of estimation theory in the context of frequentist statistics, both at its inception and as it stands today. Certainly, it encapsulates what I was taught and what I have taught others.

Suppose that θ is a parameter defining a population distribution, and that sample X consists of n observations $(x_1, x_2, x_3 \dots x_n)$ drawn at random³ from that population. X contains some amount of information about θ , which we call the Fisher information $I(\theta, X)$. We may summarize some or all of the Fisher information in the sample by computing a statistical estimate of θ , $\hat{\theta}$. We should expect that the sample estimate will not be exactly equal to

³ Randomness being an unnecessary but convenient simplification. As Fisher said, even a non-random sample is equivalent to a random sample of some population.

the parameter, a discrepancy that results from observing only a subset of the entire population. This is the sampling error of $\hat{\theta}$, $\varepsilon = \hat{\theta} - \theta$.

X is only one of a potentially infinite number of samples ($X_1, X_2, X_3 \dots$) of size n that we might draw from the population, which collectively constitute a sampling distribution for the statistic with mean equal to the expected value of $\hat{\theta}$, $E(\hat{\theta}) = \theta$. A non-zero value of the standard deviation of the sampling distribution, $\sigma_{\hat{\theta}}$, means that $\hat{\theta}$ will vary from one sample to another around θ even without measurement error. We therefore assess the precision of $\hat{\theta}$ by estimating from the same sample a second quantity, $\hat{\sigma}_{\hat{\theta}}$, called the standard error of $\hat{\theta}$.

Here, sampling error is mentioned only once, to give it a definition and to identify it as the source of our unavoidable uncertainty about the parameter's true value. The discussion then immediately moves on to standard error as a way of estimating the degree of uncertainty. There is some (unintentional) sleight of hand here: our focus has been shifted from the parameter estimate's accuracy—its probable distance from the parameter in a particular sample—to the parameter estimator's precision—the consistency with which it approaches the parameter across infinite samples. That it is pointless to directly estimate sampling error from the sample is kept implicit.

Fisher's (1925) *Theory of Statistical Estimation* exemplifies this critical omission, as it omits any rigorous attempt to provide a conceptual, mathematical or practical basis for understanding sampling error and its implications. In particular, Fisher fails to note, or apparently even to realize, that a reasonable person could ask the simple question, "Is it possible to estimate sampling error?" and expect an answer just as thoroughly considered and rigorously proven as for any other fundamental question in statistics. Fisher's only substantial consideration of sampling error in the 1925 paper is in distinguishing it from estimation error, the inaccuracy of an estimate due to the inefficiency of the estimator.

It may be that Fisher believed he had answered the question of sampling error estimation—in the negative—through his earlier work (Fisher, 1922) establishing the concepts of parametric information (later called Fisher information in his honor) in general and statistical sufficiency in particular. His definition and discussion of sufficiency in the 1922 paper could be taken as defining away sampling error as a source of information. In the formal list of definitions opening

the paper, Fisher says that “A statistic satisfies the criterion of sufficiency when no other statistic which can be calculated from the same sample **provides any additional information as to the value, of the parameter** to be estimated” (Fisher, 1922, p. 310, emphasis added). He later gives a definition “in mathematical language” that is consistent:

...if θ be the parameter to be estimated, θ_1 a statistic which contains the whole of the information as to the value of θ , which the sample supplies, and θ_2 any other statistic, then the surface of distribution of pairs of values of θ_1 and θ_2 for a given value of θ is such that for a given value of θ_1 , the distribution of θ_2 **does not involve θ** .

Fisher, 1922, pp. 316-317 (emphasis added).

It follows from these definitions that only Fisher information describes a parameter or, equivalently, a distribution that is free of Fisher information is independent of the parameter (see bolded text).

Fisher proves each of these points with rigor and precision. Yet, elsewhere in the paper, Fisher adopts less precise language: he refers to Fisher information as “**any information** given by the sample, which is **of use in estimating the values of...parameters**” (Fisher, 1922, p. 311, emphasis added). And he summarizes his mathematical definition of sufficiency thus: “In other words, when θ_1 is known, knowledge of the value of θ_2 **throws no further light upon the value of θ** ” (Fisher, 1922, p. 317, emphasis added). The distinction is subtle: in the principal definitions, Fisher is describing the *object* of the information, while in the informal summaries of those definitions, he is referring to the *usefulness* of the information.

A reasonable reader might interpret Fisher’s wording, absent context, as indicating that information *about* a parameter is the only quantity in the sample that is useful for estimating the parameter. Fisher, of course, knew this to be false, or at least would determine it to be false by the time of writing the 1925 paper, wherein he introduces the principle of statistical ancillarity.⁴ He clearly understood that

⁴ Ancillary statistics are those which contain no Fisher information. Ancillary compliments are ancillary statistics that, when conditioned upon by an estimator that is not sufficient, allows that estimator to recover some or all of the parameter information from the sample that otherwise would be lost. Thus, ancillary information can be useful for parameter estimation, but only to the extent that it liberates additional Fisher information from the sample.

information other than Fisher information could be of use in estimation, but only in so far as it allows us to recover more information about the parameter.

Now, it would be foolish to take Fisher's shorthand descriptions of what constitutes useful information as seriously or as literally as his technical definitions. Nonetheless, the impression left is that Fisher may have believed he had proved more than he in fact had. He had proved that a parameter estimate improves as it recovers more of a sample's parametric information, and that there is a limited amount of such information in any sample. Fisher may also have believed he had proved a much stronger claim, that a parameter estimate can *only* be improved by increasing the amount of parametric information the estimator recovers from the sample. Once the limit of parametric information has been reached, then in the absence of measurement error, the estimate can no longer be improved. Had he managed to prove the stronger claim, that proof would certainly have excluded the possibility of sampling error estimation. Otherwise, we could improve even the estimate given by a sufficient statistic by removing some or all of the random sampling error in the estimate.

In a world where the stronger claim holds, sampling error would have to be unpredictable and unknowable and unusable. It would represent not information but the absence of information. This conclusion is consistent with the inattention to the concept of sampling error by Fisher, his contemporaries and statisticians for generations to follow.

3 The missing theorem no one missed

Of course, I cannot know Fisher's thought process on the question⁵—he may simply have taken for granted that sampling error cannot be estimated, axiomatically, without giving the matter any particular thought. But it is plausible that the minor role Fisher gave sampling error in his theory, together with his occasionally imprecise language, set a precedent for later authors, teachers and practitioners to accept, and to perpetuate, Fisher's oversight. I say it has been

⁵ Reading the entirety of the extremely prolific Fisher is beyond me, so it very well could be that I have missed some key comment, albeit one that has gone largely uncited and was never catalogued in a searchable way in popular databases.

perpetuated because, after a thorough search of the literature, I find no proof, no citation or allusion to a proof, not even a mention of the need for a proof of this apparently universal conviction.⁶ Instead, in the rare cases when the possibility of sampling error estimation is raised⁷, it is invariably in the form of a cursory dismissal.

To illustrate, I have culled the following quotations from the work of deservedly-lauded experts in statistics and its application in psychology, evaluation and related fields (bolded emphasis added):

"Since the actual **sampling error in an estimate cannot be known**, we obtain a measure of the average magnitude over all possible samples of the samples of the sampling error in the estimate." P.V. Sukhatme, 1954, p 28.

"**It is impossible to correct for the effect of sampling error in an isolated study**. The confidence interval gives an idea of the potential size of the sampling error, but **the magnitude of the sampling error in any one study is unknown and hence cannot be corrected**." Cooper and Hedges, 1993, p. 324. (See also Cooper, Hedges and Valentine, 2019, p. 317.)

"The amount of [sampling] **error associated with a particular sample is not known in practice**, since to know the error we need to know the value of the parameter that is to be estimated." Biemer and Lyberg, 2003, p. 320.

⁶ Here I will quote Guttman (1977): "Empirical proof of non-existence is in principle difficult, but proof of existence is comparatively simple--requiring exhibiting but a single example. Thus, when an assertion below is of the form: 'No textbook proves that...', and should one wish to document such an assertion, one would have to refer to all extant textbooks. On the other hand, should a reader believe he can prove the assertion to be false, all he need do is provide a single correct reference. I would appreciate receiving any such rectifying reference from any interested reader..." (p. 84). The difficulty of documenting "what does not exist" (p. 84) is both a limitation on my argument and, perhaps, a reason that so few have pursued the question of sampling error estimation at all.

⁷ Curiously, the earliest reference of this sort I have found is from 1956, three decades after Fisher's Theory of Statistical Estimation (1925). Even searches for related terms like "discrepancy" and "random error" and "entropy" were not fruitful prior to that date, though the last of these did uncover several early attempts to reconcile Fisher information with Shannon information. Perhaps earlier authors were just more cautious; see footnote #5.

“While the error in any given study is random, we *can* estimate the sampling distribution of the errors.” (italicized emphasis original) Borenstein, Hedges, Higgins and Rothstein, 2011, p. 65.

“At the level of the single study, sampling error is a random event. If the observed correlation is .30, then the unknown population correlation could be higher than .30 or lower than .30, and **there is no way that we can know the sampling error or correct for it.**” Schmidt and Hunter, 2014, p. 90.

“Although a confidence interval for the effect size provides an unbiased estimate of potential sampling error for the single study, **there can be no correction for sampling error at the level of the single study.**” Schmidt and Hunter, 2014, p. 275-76.

My inclusion of these quotes is not intended as a criticism of their authors but to highlight the ubiquity of a claim for which no rational or empirical basis is provided⁸ or cited. As a community of scholars and users of statistics, we define ourselves by our skepticism of the unquantified claim, yet we appear to have unskeptically embraced this one.

The omission is not trivial. It represents a massive blind spot that may be indicative of deeper flaws in our statistical paradigm. It means our understanding of statistics is incomplete, and our progress toward greater understanding is hobbled. This is so even if the conclusion, that sampling error is inestimable, should ultimately hold true. Per Guttman (1977): “...progress is certainly not made easier, and may even be prevented, by having researchers subservient to or continually misdirected by thinking and practices that are not what they are presumed to be.”

4 Obstacles to conceiving of the missing theorem

Every claim to knowledge should be appraised with skepticism—the more consequential the claim, the greater the skepticism. This is true of the claim that sampling error cannot be estimated from an isolated sample. It is equally true of

⁸ A couple of the quotations do offer statements resembling a rationale, e.g., sampling error is random, the parameter is unknown. But as I demonstrate in the next section, these are tautologies at best and fallacies at worst.

my own assertion, that all statisticians, mathematicians and experts of every kind have overlooked a fundamental principle of estimation theory for over a century. I believe I have made the factual case that this is true. Nonetheless, I also feel compelled to make the case that it is *plausible*, and to that end I now present a number of possible explanations for how this state of affairs could have persisted for so long.

I frame these explanations as five obstacles that have, perhaps, prevented us from perceiving the sizeable gap in our statistical foundations. To be clear, the thesis of this paper in no way hangs on plausibility or verification of these obstacles. Their value lies only in their potential to persuade the skeptic, and to serve the persuaded as a source of insights as to where we, as scholars, may have gone wrong, and how we might right ourselves.

The obstacle of language. Probably the most familiar criticism of statistical language is its use of lay terms for technical concepts, with definitions that diverge from, or even contradict, their common definitions (e.g., “errors” are not mistakes, “significant” findings are not necessarily meaningful, “independent” variables are generally interrelated, etc.; Guttman, 1977). The linguistic obstacle I refer to here involves a different sort of confusion, one arising from the use of nearly identical terms for distinct concepts in estimation. Consider the following “clarifications”:

- Sampling error is one part of a sample’s error.
- We assess sampling error using the sample standard error.
- The population standard error is the standard deviation of the sampling distribution.
- We compute the sample standard error from the standard deviation of the sample’s distribution.

Each of these statements is accurate in its use of statistical terminology, and each means something distinctly different. Any of these several sentences might plausibly appear by itself in a paragraph from an introductory statistics text, and all of them might appear throughout a chapter. The experienced instructor will sympathize when I say that I have had many intelligent, attentive students flummoxed by more straightforward passages in standard textbooks. Not to mention my own lecture slides.

Complaints about phrasing may seem trivial or pedantic, but Fisher himself recognized that progress in statistics could be held back by “purely verbal confusion” (Fisher, 1922, p. 311). He even coined the term “statistic” as a replacement for “statistical parameter” specifically so that researchers would stop confusing statistics with parameters.⁹ Evidence that undifferentiated terminology has led to fuzzy thinking (or at least poor communication) about sampling error can be found throughout the literature. A disturbingly large number of works (journal articles, textbooks, classroom materials) appear to conflate sampling error—a quantity of definite magnitude and sign—with standard error—a variance. For example, the vast majority of results returned in a search of scholarly databases for the phrase “sampling error” pertain to authors reporting the square root of a sample variance divided by sample size. Authors in many other cases use the terms without meaningful distinction (e.g., “Sampling error is the accuracy of an estimate that is expressed as the standard error...”).

It is likely that many of these authors actually do understand the difference between sampling error and standard error. But it is clear that they see little harm in conflating the concepts, probably because they assume actual sampling error can have no practical role in the calculation of statistical estimates. And those who absorb this ambiguity are likely to assume the same. Thus, misapprehension continues to be passed from author to reader, and from teacher to student.

The obstacle of assumptions. It is difficult to make a positive case for the necessity of the missing theorem, given that I can find no real examples of arguments to the contrary—if the truth of a thing is universally assumed, who would bother to argue for it? In the absence of prior debate, I have attempted to anticipate the potential counterarguments, which includes several I raised against myself in the course of this work.

All of these counterarguments must, under my thesis, rest upon one or more false assumptions. It is these assumptions I must therefore negate. I acknowledge the risk of my unwittingly constructing straw men only to claim

⁹ I very nearly followed Fisher’s example by using some other term throughout this paper in place of “sampling error.” I chose not to do so, however, out of concern that this could compound confusion, perhaps even leading some to assume that I am introducing an entirely new concept. Nonetheless, I have begun referring to epsilon or epsilon-hat in my other work.

victory when they slump at my feet. I therefore entreat the reader to subject my reasoning to the deepest scrutiny.

Assumption #1: A truly random outcome by its nature is unpredictable.

It is true that we cannot predict the future outcomes of random events, like the result of a particular coinflip or die roll; however, nothing precludes us from estimating the probability that a dataset (sequences of H's and T's, or of the integers from one to six) was generated from a random process. This is the difference between inference as prediction and inference as estimation. Computer scientists have developed a wide variety of tests of randomness in data, so the idea of estimating the degree or probability of randomness in sample data ought not seem implausible, especially if we can make parametric assumptions about the non-random components.

Assumption #2: The random variance of a statistic cannot be estimated because it is perfectly confounded with the true variance.

This is just restating the premise as argument: the question is whether they are perfectly confounded, so saying they that are assumes the conclusion.

Assumption #3: The amount of sampling error in an estimate from a single sample cannot itself be estimated because that would require knowing the true value of the parameter being estimated.

Estimating quantities with unknowable true values is literally the point of statistical inference. Knowing the true value of the parameter is no more necessary for sampling error estimation than knowing the true value of the sampling error is necessary for parameter estimation.

Assumption #4: The very idea of sampling error estimation is absurd. It cannot be addressed in estimation theory because the claim is not a coherent statistical or mathematical statement.

Again, this is just restating the impossibility, not an argument in itself. Even then, it ought to be easy to falsify an absurd statement, mathematical or not.

Assumption #5: The proof already exists: sampling error estimation is precluded by the principle of statistical sufficiency.

Superficially, this is a strong argument: We know that a sufficient statistic $\hat{\theta}$ uses all of the Fisher information in a sample, $I(\theta, X)$; yet, if we estimate the sampling error of $\hat{\theta}$ and define a new estimator of θ as the difference $\hat{\theta} - \hat{\varepsilon} = \hat{\theta}^*$, then $\hat{\theta}^*$ will be a better estimator of θ than $\hat{\theta}$. This implies that $\hat{\theta}^*$ contains more Fisher information than does a sufficient statistic computed for the same sample, a contradiction.

The error in this reasoning was discussed earlier in connection with Fisher's introduction of sufficiency: he clearly proved that increasing the information about a parameter makes its estimator more accurate, but that does not imply that increasing parameter information is the *only* way to improve estimator accuracy. Certainly, we cannot increase the Fisher information in a statistic above the total Fisher information contained in the sample. But accounting for the sampling error in an estimate could never add Fisher information to $\hat{\theta}$ anyway. We know that any information added to $\hat{\theta}$ by subtracting $\hat{\varepsilon}$ cannot be Fisher information because the two quantities, one relating to a fixed parameter value and the other to a random process, are statistically independent.

Recall that Fisher information is a long-run property of an estimator, inversely related to standard error, such that a higher-information estimator will deviate less from the true value of its parameter across infinitely many samples of equal size. Sampling error, in contrast, does not even exist in the long run and so cannot vary across samples. Instead, sampling error is generated anew each time a new sample is drawn, even from the same population. My point is that these are two distinct, even incomparable, constructs. If subtracting $\hat{\varepsilon}$ improves $\hat{\theta}$, as any sampling error estimator must, it does so by a mechanism that in no way involves Fisher information.

Now, one might claim this proves the impossibility of estimating sampling error, arguing that any increase in accuracy must result from an increase in Fisher (parametric) information, which we have just established $\hat{\varepsilon}$ cannot achieve. Once again, this is a restating of the thing to be proven, not a proof of it. Saying that only parametric information improves a parameter estimate is equivalent to saying that sampling error cannot be estimated, as is saying that sampling error is confounded with parametric information, as is saying that the concept of sampling

error estimation is inherently unprovable. Present a proof of any of these statements and we will have our missing theorem.

The obstacle of definitions. This obstacle is akin to the Kuhnian idea (Kuhn, 1962) that the development of knowledge in a field may be limited by a narrow and static paradigm. Within the frequentist statistical paradigm, there are only so many kinds of fundamental objects we can consider. There are *observations*, which are grouped into *samples*, which are summarized by *statistics*, which estimate *parameters*, which define the *population distributions* from which observations are drawn.¹⁰ These objects are defined, individually and collectively, in such a way as to seem a complete whole, admitting no other constructs of equal significance. Almost everything else considered in estimation theory is defined as a particular example of one of these kinds of objects or a function thereof.

The obstacle of definitions arises because sampling error is...weird. It fits the definition of neither a parameter—it exists only with respect to a single sample—nor a statistic—it is inherently unobservable—nor any other primary statistical object. It also has no meaningful interpretation in the long run. Because the focus of classical statistics in general and Fisher’s estimation theory in particular is almost entirely long run or asymptotic outcomes, it only seems natural to minimize discussion of sampling error. Consideration of sampling error estimation then seems downright unnatural, if one even considers it at all. (Clearly, few have.) Sampling error, when viewed through the frequentist lens, is just the space between a parameter and its estimate. Why should a statistician conceive of that space as the object of estimation? Should a linguist think of the spaces between words in a sentence also as words?

The obstacle of practicality. We are all applied statisticians. As were Fisher and Gossett, Neyman and Pearson. These men developed important statistical theories, but they were not, I argue, theoretical statisticians.¹¹ A true theorist is as interested in the abstract and the absurd as in the concrete and the practical. The favored tool of the theorist is deductive logic, and she will use that tool to push any and every idea to its most extreme conclusion. Hence, the

¹⁰ What is or is not fundamental is highly debatable—is information a fundamental object or just an attribute of one?

¹¹ It has been said that “theoretical statistics is the theory of applying statistics.” (Gelman, 2017)

greatest theoretical physicists have concerned themselves with questions like, What would happen if a tiny demon were to sort particles by their relative entropy (Maxwell)? and, What would you see if you could ride on a beam of light (Einstein)? We have seldom seen statisticians pursue such fanciful lines of inquiry--can you imagine Fisher, for example, digressing into an exploration of a hypothetical universe wherein the result of the central limit theorem is a bimodal distribution?

No, Fisher actually warned against the contamination of statistical thought by the “modern abstract tendencies” characteristic of pure mathematicians (Savage, 1976, p. 441). Fisher denigrated the application of “exact, precise deductive reasoning” (Fisher, 1958, p. 26) to statistics as “fruitless” (Fisher, 1958, p. 38) and “muddle-headed” (p. 37), and those who felt otherwise as suffering from “a grave confusion of thought” (p. 38). Only the “inductive logical processes” (p. 261) of the empiricist could provide the basis for making inferences from observations. Most ominously, Fisher foretold a world in which a generation of statisticians “with a dense fog where their brains ought to be” (p. 38) would be so misled by the impractical, so distracted by the abstract, as to threaten everything from medical care to the national defense.

Fisher's attitude toward the theoretical may have been singularly grim among early statisticians, but an early focus on the practical and the empirical across the field has unquestionably shaped our modern understanding of what statistics is and how it should be used. Nearly all of our institutions—education, funding and publication; professional advancement; reputational cache—push statisticians toward programs of research that develop and extend empirical applications. None of which is necessarily problematic. Consider the dilemma of a statistician working on a problem with two possible solutions: one far more likely to work, the other far less. *Of course* she should try the former solution first. If that one works, *of course* she would not circle back and try the implausible approach. The problem comes when every other statistician working on a similar problem also chooses the plausible approach.

In contrast, it is a theoretician's job to not just demonstrate that one solution works but whether both work, and if so, how they compare. And if not, *what if it did?* The fact that a question can be asked and has no answer is rationale enough for the theoretician to answer it. This is not to say that statisticians are

incurious, only that we lack the incentives to pursue knowledge that has no obvious practical application. Or to write theorems proving things that are self-evidently true.

The obstacle of intentionally hiding science from scientists. The effort to conceal scientific knowledge from researchers across all fields, commonly referred to as “paywalls,” has inflicted severe limitations on both the quality and depth of research on this topic. I know this because paywalls have slowed the progress of this paper repeatedly. My own difficulties aside, however, the most painful consequence posed by the mass disappearing of scientific contributions is not that it hinders individual researchers’ access to essential articles and other publications. It is that we cannot know which publications are essential without accessing them.

The good news on this obstacle is that authors, who once felt they had no choice but to cooperate, are increasingly subverting the paywall system by posting preprints, publishing in open-access journals, and adding their previously-published work to open-access databases. The bad news is that a century’s worth of work remains locked away, awaiting their incremental release to the public domain.

5 A new perspective on sampling error in estimation

Assumption #4 (under the obstacle of assumptions) contains a nugget of truth: that the question of sampling error estimation cannot be formulated in a practical, testable way is mere assumption, but it is true that the question has never before been posed in such a way. It is a fact that suggests yet another obstacle: how can we possibly conceive of a theorem, when no one has managed to state its premise formally and with precision? Much easier to equate inconceivability with impossibility and move on. In this final section, my objective is to present a rough framework—not fully formed, and certainly not definitive—for imagining how a sampling error estimator might work, and where it might fit into our paradigm with minimal bending.

In dismantling another assumption, that sufficiency rules out sampling error estimation, we have already found two principles for this nascent theory. First, the

existence of a sampling error estimator implies the existence of some kind of information in a sample, other than Fisher information, that we can use to improve a parameter estimate. Second, sampling error is not easily modeled as a parameter. Incorporating it into estimation theory would seem to require either a broadening of existing definitions or the creation of new ones.

What type of statistical object is sampling error? As previously mentioned, there are only so many kinds of fundamental statistical objects¹² within the frequentist statistical paradigm.¹³ Sampling error is unlike any other of these objects. Unlike a parameter, it exists only with respect to a single sample; unlike a statistic, it is inherently unobservable. Sampling error neither belongs to a population nor has a distribution across all samples. It has no long-run interpretation at all since it belongs to a single sample, comes into existence only when the sample is drawn, and disappears when the sample is discarded.

Some might say sampling error is not an object at all, just the distinction between two objects—the statistic and its parameter. This perspective may or may not be wrong, but it certainly something worse than wrong—it is not useful. For the moment, it useful to suppose that sampling error is some kind of statistical object, as it would seem to be necessary if we are to conceive of its estimation.

What type of information is about sampling error? In the introduction, I showed that Fisher and statistical theorists in general conceived of a statistical paradigm with no practical role for sampling error in estimation. The only sort of statistical object permitted to play such a role is one that contains information, a reasonable position. Fisher (1925) asserted that the only useful information in a sample is either information about the parameter of interest or ancillary information that increases parametric information in the estimator. The only other kind of sample information Fisher mentions is that pertaining to irrelevant (nuisance) parameters.

¹² “Statistical object” is an awkward term, but I cannot find an adequate substitute in common usage. “Quantity,” “factor,” “term,” “characteristic,” “value,” and “component” are sometimes used, but all seem to skirt describing statistical objects as concrete things, as opposed to mere mathematical conveniences.

¹³ What is or is not fundamental is highly debatable—is information a fundamental object or just an attribute of one?—but certainly the objects given here underlie the vast majority of principles conveyed in the typical Intro Stats curriculum.

This information typology excludes the possibility of sampling error information. Sampling error is independent of any parameter and, as I say above, sampling error cannot itself be a parameter because it does not exist independent of a particular sample—it is an unobservable sample attribute. It therefore follows that if we are to estimate sampling error, we must add a new information type, which I will call non-Fisher information. Non-Fisher information does not describe any parameter but is nonetheless useful in parameter estimation.

What type of estimator conveys sampling error information usefully?

Suppose $\hat{\theta}$ and $\hat{\theta}^*$ are two estimators of the parameter θ . $\hat{\theta}$ is a terrific estimator, in terms of sufficiency, efficiency, and consistency: it contains all the information about θ from the sample, it has a much smaller standard error compared to $\hat{\theta}^*$ for any sample size, and its value rapidly converges to θ as sample size increases. $\hat{\theta}^*$ on the other hand is a terrible estimator: the proportion of information about θ from the sample captured by $\hat{\theta}^*$ is incredibly small and easily overwhelmed by sampling error at any sample size, and the tendency of $\hat{\theta}^*$ to converge to θ reaches an asymptotic limit rapidly as sample size increases. I will call $\hat{\theta}^*$ a *deficient statistic* with respect to θ .

We can attribute most of the variation in the value of $\hat{\theta}$ to variation in θ , with some perturbation from sampling error. The same is not true of $\hat{\theta}^*$: most of the cross-sample variation observed in $\hat{\theta}^*$ will be due to sampling error, thus it can provide only very limited information about θ . A statistician with $\hat{\theta}$ in her toolbox would seem to have no incentive to ever employ an estimator with the properties of $\hat{\theta}^*$; if given no alternative estimator of θ , she would be extremely motivated to find one. Fisher rightly saw great value in an estimation theory that provides tools for the discovery and discernment of estimators like $\hat{\theta}$ that maximize Fisher information and minimize sampling error contamination.

What Fisher missed is the value of a theory that also accomplishes the equal and opposite objective, providing tools for the discovery and discernment of estimators like $\hat{\theta}^*$ with maximum non-Fisher information about sampling error and minimal contamination from the parameter. Such a terrible parameter estimator will covary far more with sampling error than with the parameter, and it would, in theory, quantify the sampling error. In fact, $\hat{\theta}^*$ need not be maximally deficient to be useful—it may contain some small amount of information about the parameter just as $\hat{\theta}$ must invariably contain some amount of sampling error. I propose that

such a statistic may provide an estimate of sampling error so long as each of the five criteria below are met for some meaningful (read: useful) range of parameter values:

1. The correlation between θ^* and the parameter θ is small¹⁴;
2. The correlation between θ^* and the sampling error ε of $\hat{\theta}$, an estimator of θ , is not small;
3. The correlation between $\hat{\theta}$ and θ is not small;
4. The general functional form of the joint distribution among unknowns (θ , ε) and estimators ($\hat{\theta}$, θ^*) is known (or reasonably assumed);
5. The relationships among these variables and their estimators are relatively consistent even as θ varies within some range.

The first three criteria refer to the relative strengths of the relationships among the quantities θ , $\hat{\theta}$, θ^* and ε , implying a model roughly like the one in Figure 2:

Figure 2. The relative relationships among the parameter of interest θ , its estimator $\hat{\theta}$, the sampling error ε in $\hat{\theta}$, and the hypothetical θ^* , a deficient estimator of θ that functions as a

¹⁴ “Small” is obviously a relative term. One might prefer to say that the correlation between the sampling error estimator and the parameter is nearer zero than are the correlations described in 2 and 3 as “not small.”

sampling error estimator for θ . Note that the given relationship between θ and ε is not postulated but is, nonetheless, unavoidably true.

In the model, the parameter estimator and the sampling error estimator share the same sources of variance, but the latter has much less in common with the parameter. Assuming the absence of measurement error, we know that all of the variance in θ is explained by θ and ε , whereas θ barely informs θ^* . It should not be inferred that the paths labeled “ > 0 ” are all equal; indeed, if θ is reasonably consistent, the ratio of Fisher information to sampling error will increase with sample size.

The last two criteria, relating to the functional form of the quantities’ joint distribution, are not unique to estimators of sampling error: any time we want to make inferences from observables to unobservables, we must know how they covary (fourth criterion), and that relationship has to hold up across a range of plausible values (fifth criterion). Given that we currently have no way to directly obtain information about the size of sampling error in a particular sample, a sampling error estimator need only meet these criteria to the least degree to be of use to us. At a minimum, we want to be able to make general statements like “as ε increases, so does the expected value of θ^* ” or “the strength of the relationship between ε and θ^* is proportional to the size of θ ,” without knowing the true values of either ε or θ , and to have such statements apply with a degree of certainty that we can describe probabilistically.

The five criteria do not include a requirement that a sampling error estimator be completely deficient or that the parameter estimator be a sufficient statistic, as they were in the original scenario. However, for the purposes of identifying a sampling error estimator, searching for a variable that is as deficient as possible is a good place to start. Likewise, if we want to prove that θ^* is estimating sampling error and not merely some stray Fisher information in the sample that θ fails to account for, it is helpful if θ is a sufficient statistic for θ .

6 Conclusion

Imagine you are giving an introductory lecture on statistical inference in the social sciences. The intended audience is college students who have never taken a statistics course, but your colleagues from various fields—psychology, learning

sciences, quantitative methods—are in attendance as well. You begin by explaining that the basic procedure for estimating a parameter from a sample of observations is this: First, compute a quantity from the data that contains as much information about the parameter as possible. This is the parameter estimate. Second, compute another quantity from the same data that contains as much information about the degree of (irrelevant) randomness as possible. This is the estimate of the accuracy of the parameter estimate. The difference of the first and second quantities is the revised parameter estimate. Finally, compute a third quantity, still from the same data, that contains as much information about variability in the sample as possible. This is the estimate of the precision of the revised parameter estimate.

The students have no reason to doubt you—the whole thing sounds perfectly reasonable to them. *Of course* one would want to find out how far off the initial estimate is, and then correct for that discrepancy. Your colleagues, however, immediately realize what you have said is not just wrong but demonstrably absurd. There is much head shaking in the back of the room, where more than one hand stridently rises.

This paper is not about whether it is the students or your colleagues who are correct. This paper is about the fact that your colleagues can provide no better justification for their belief than can the students. The students, at least, will not argue with conviction. Such is the world in which we find ourselves, where we practice or teach from atop a chair with too few legs. The ramifications of embracing this realization, I believe, are infinitely greater than those of the missing theorem itself.

References

- Biemer, P. P., Lyberg, L. E. (2003). *Introduction to Survey Quality*. United States: Wiley.
- Borenstein, M., Hedges, L. V., Higgins, J. P. T., Rothstein, H. R. (2011). *Introduction to Meta-Analysis*. Germany: Wiley.
- Cooper, H., Hedges, L. V., and Valentine, J. C. (2019). *The Handbook of Research Synthesis and Meta-Analysis*. United States: Russell Sage Foundation.

- Hedges, L. V. and Cooper, H. (1993). *The Handbook of Research Synthesis*. United States: Russell Sage Foundation.
- Fisher, R. A. (1922). On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 222(594-604), 309-368.
- Fisher, R. A. (1925, July). Theory of statistical estimation. In *Mathematical Proceedings of the Cambridge Philosophical Society* (Vol. 22, No. 5, pp. 700-725). Cambridge University Press.
- Fisher, R. A. (1959). *Smoking: the cancer controversy: some attempts to assess the evidence*. Edinburgh: Oliver & Boyd.
- Gelman, A. (2017, May 26). Theoretical Statistics is the Theory of Applied Statistics: How to Think About What We Do. *Statistical Modeling, Causal Inference, and Social Science*. <https://statmodeling.stat.columbia.edu/2017/05/26/theoretical-statistics-theory-applied-statistics-think/>
- Guttman, L. (1977). What is not what in statistics. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 26(2), 81-107.
- Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. United States: University of Chicago Press.
- Savage, L. J. (1976). On rereading R. A. Fisher. *The Annals of Statistics*, 441-500.
- Schmidt, F. L., Hunter, J. E. (2014). *Methods of Meta-Analysis: Correcting Error and Bias in Research Findings*. United States: SAGE Publications.
- Sukhatme, P. V. (1954). *Sampling theory of surveys*. Iowa State University Press. Ames, Iowa.

Appendix

What it means for the theorem to be missing, and why we should care

Each of the many statistical tools researchers have at our disposal are based on a set of core claims, most mathematical in nature, that together constitute the foundation of estimation theory. For example, Student's t -test is a tool, one that relies on the claim (among others) we call the Central Limit Theorem (CLT). These core claims are what establish which methods are valid and under what circumstances—we avoid using t -tests with $n = 5$ because the CLT cannot be relied upon in samples so small.

All statistical tools can be deduced from the core claims of statistics. It goes without saying that we take these claims to be true, though most of us have never explored too deeply *why* they are true. To those scholars whose business it is, the truth of these core claims can be evaluated using two criteria: rational consistency and empirical consistency. Rational consistency is the extent to which the logic (i.e., math) of a claim is consistent internally (it does not contradict itself) and externally (it does not contradict other core claims). Empirical consistency is the extent to which a claim corresponds to observations of the real world—the tools derived from a claim must be observed to work as predicted when applied to actual data.

The CLT, to continue that example, is mathematically provable and therefore rational, but it is also empirical because we can observe it in real life: we can add data points to a sample and see its distribution gradually approach normality. Another example is the concept of a sampling distribution for a statistic—the collection of all possible values of the statistic given all possible samples of size n , in proportion to their probability of being observed. A sampling distribution is a mathematical fiction convenient for modeling real processes but exists in no real place. We can validate it directly by assessing its rational consistency, but we can evaluate its empirical consistency only indirectly, by confirming that predictions about sampling, conditioned on the sampling distribution model for sampling, hold up in real life. Luckily, they do!

Which brings me to my point. The existence of sampling error is a core claim of statistics, one which we can confirm empirically—we might take several samples from census data and subtract the mean of the total from each sample mean, noting the differences as sampling errors. But we also make claims *about* sampling error and how it relates to other statistical concepts. In particular, we claim that the sampling error of a statistic cannot be estimated solely from the sample used to calculate the statistic. Is this claim true? Unfortunately, the claim has never before been given mathematical expression, nor has it been deduced from other, proven theorems, so its rational consistency has never been either confirmed or refuted. To evaluate the claim's empirical consistency, we would test whether the consequences of its being true or false are observed in reality. Again unfortunately, the only prediction implied by the claim is that a certain kind of event will never occur, a prediction that can be disproven (by offering a counterexample) but never proven. That is, the failure to observe sampling error being estimated is not a proof of the claim that it can never be estimated.

So the claim of sampling error's inestimability exists in a state of limbo—neither a mathematical statement nor a scientific one. Yet it is without question a statistical statement, and so part of statistical theory; and it is a statement about the real world, so it must either be true or false, with different implications either way. (Indeed, if the statement is false, the implications are huge!) In other words, a theorem proving or disproving this statement has a place beside the Central Limit Theorem and the description of sampling distributions, a place that is currently empty. Put more colorfully, the thing holding up our understanding of statistics has a big hole in it. This paper is an attempt to make clear the emptiness of that space and the necessity of filling it.