

Bayesian Data Analysis in Empirical Software Engineering Research

Carlo A. Furia, Robert Feldt, and Richard Torkar

Abstract—Statistics comes in two main flavors: frequentist and Bayesian. For historical and technical reasons, frequentist statistics have traditionally dominated empirical data analysis, and certainly remain prevalent in empirical software engineering. This situation is unfortunate because frequentist statistics suffer from a number of shortcomings—such as lack of flexibility and results that are unintuitive and hard to interpret—that curtail their effectiveness when dealing with the heterogeneous data that is increasingly available for empirical analysis of software engineering practice.

In this paper, we pinpoint these shortcomings, and present Bayesian data analysis techniques that provide tangible benefits—as they can provide clearer results that are simultaneously robust and nuanced. After a short, high-level introduction to the basic tools of Bayesian statistics, we present the reanalysis of two empirical studies on the effectiveness of automatically generated tests and the performance of programming languages. By contrasting the original frequentist analyses with our new Bayesian analyses, we demonstrate the concrete advantages of the latter. To conclude we advocate a more prominent role for Bayesian statistical techniques in empirical software engineering research and practice.

Index Terms—Bayesian data analysis, statistical analysis, statistical hypothesis testing, empirical software engineering.

1 INTRODUCTION

Compared to other fields such as the social sciences, *empirical software engineering* has been using standard statistical practices for a relatively short period of time [42]. In light of the replication crisis mounting in several of the more established research fields [77], empirical software engineering’s relative statistical immaturity can actually be an opportunity to embrace more powerful and flexible statistical methods—which can elevate the impact of the whole field and the solidity of its research findings.

In this paper, we focus on the difference between so-called *frequentist* and *Bayesian* statistics. For historical and technical reasons [26], [30], [80], frequentist statistics are the most widely used, and the customary choice in empirical software engineering (see Sect. 1.1) as in other experimental research areas. In recent years, many statisticians have repeatedly pointed out [22], [49] that frequentist statistics

suffer from a number of shortcomings, which limit their scope of applicability and usefulness in practice, and may even lead to drawing flat-out unsound conclusions in certain contexts. In particular, the widely popular techniques for null hypothesis statistical testing—based on computing the infamous p -values—have been *de facto* deprecated [7], [41], but are still routinely used by researchers who simply lack practical alternatives: techniques that are rigorous yet do not require a wide statistical know-how, and are fully supported by easy-to-use flexible analysis tools.

Bayesian statistics has the potential to replace frequentist statistics, addressing most of the latter’s intrinsic shortcomings, and supporting detailed and rigorous statistical analysis of a wide variety of empirical data. To illustrate the effectiveness of Bayesian statistics in practice, we present two *reanalyses* of previous work in empirical software engineering. In each reanalysis, we first describe the original data; then, we replicate the findings using the same frequentist techniques used in the original paper; after pointing out the blind spots of the frequentist analysis, we finally perform an internal replication [17] that analyzes the same data using Bayesian techniques—leading to outcomes that are clearer and more robust. The bottom line of our work is not that frequentist statistics are intrinsically unreliable—in fact, if properly applied in some conditions, they might lead to results that are close to those brought by Bayesian statistics [94]—but that they are generally less intuitive, depend on subtle assumptions, and are thus easier to misuse (despite looking simple superficially).

Our first reanalysis, presented in Sect. 3, targets data about debugging using manually-written vs. automatically-generated tests. The frequentist analysis by Ceccato et al. [18] is based on (frequentist) linear regression, and is generally sound and insightful. However, since it relies on frequentist statistics to assess statistical significance, it is hard to consistently map some of its findings to practical significance and, more generally, to interpret the statistics in the quantitative terms of the application domain. By simply switching to Bayesian linear regression, we show many of these issues can be satisfactorily addressed: Bayesian statistics estimate actual probability distributions of the measured data, based on which one can readily assess practical significance and even build a *prediction model* to be used for estimating the outcome of similar experiments, which can be fine-tuned with any new experimental data.

The second reanalysis, presented in Sect. 4, targets

- C.A. Furia is with the Software Institute, Faculty of Informatics, Università della Svizzera italiana (USI), Lugano, Switzerland. E-mail: see bugcounting.net
- R. Feldt and R. Torkar are with the Software Engineering Division, Department of Computer Science and Engineering, Chalmers University of Technology and the University of Gothenburg, Gothenburg, Sweden.

data about the performance of different programming languages implementing algorithms in the Rosetta Code repository [83]. The frequentist analysis by Nanz and Furia [71] is based on a combination of null hypothesis testing and effect sizes, which establishes several interesting results—but leads to inconclusive or counterintuitive outcomes in some of the language comparisons: applying frequentist statistics sometimes seems to reflect idiosyncrasies of the experimental data rather than intrinsic differences between languages. By performing full-fledged Bayesian modeling, we infer a more coherent picture of the performance differences between programming languages: besides the advantage of having actual estimates of average speedup of one language or another, we can quantitatively analyze the *robustness* of the analysis results, and point out what can be improved by collecting additional data in new experiments.

Contributions. This paper makes the following contributions:

- a discussion of the shortcomings of the most common frequentist statistic techniques often used in empirical software engineering;
- a high-level introduction to Bayesian analysis;
- detailed analyses of two case studies [18], [71] from empirical software engineering, discussing how the limitations of frequentist statistics weaken clarity and generality of the results, while Bayesian statistics make for robust and rich analyses.

Prerequisites. This paper only requires a basic familiarity with the fundamental notions of probability theory. Then, Sect. 2 introduces the basic principles of Bayesian statistics (including a concise presentation of Bayes’ theorem), and illustrates how frequentist techniques such as statistical hypothesis testing work and the shortcoming they possess. Our aim is that our presentation be accessible to a majority of empirical software engineering researchers.

Scope. This paper’s main goal is demonstrating the shortcomings of commonly used frequentist techniques, and to show how Bayesian statistics could be a better choice for data analysis. This paper is *not* meant to be:

- 1) A tutorial on applying Bayesian statistic tools; we refer readers to textbooks [37], [57], [62] for such step-by-step introductions.
- 2) A philosophical comparison of frequentist vs. Bayesian interpretation of statistics.
- 3) A criticism of the two papers (Ceccato et al. [18] and Nanz and Furia [71]) whose frequentist analyses we peruse in Sect. 3 and Sect. 4. We chose those papers because they carefully apply generally accepted best practices, in order to demonstrate that, even when applied properly, frequentist statistics have limitations and bring results that can be practically hard to interpret.

Availability. All machine-readable data and analysis scripts used in this paper’s analyses are available online at

<https://bitbucket.org/caf/bda-in-ese>

1.1 Related Work

Empirical research in software engineering. Statistical analysis of empirical data has become commonplace in

software engineering research [4], [42], [99], and it is even making its way into software development practices [54].

As we discuss below, the overwhelming majority of statistical techniques that are being used in software engineering empirical research are, however, of the frequentist kind, with Bayesian statistics hardly even mentioned.

Of course, Bayesian statistics is a fundamental component of many machine learning techniques [8], [46]; as such, it is used in software engineering research indirectly whenever machine learning is used. In this paper, however, we are concerned with the direct usage of statistics to analyze empirical data from the scientific perspective—a pursuit that seems mainly confined to frequentist techniques in software engineering [42]. As we argue in the rest of the paper, this is a lost opportunity because Bayesian techniques do not suffer from several limitations of frequentist ones, and can support rich, robust analyses in several situations.

Bayesian analysis in software engineering? To validate the impression that Bayesian statistics are not normally used in empirical software engineering, we carried out a small literature review of ICSE papers.¹ We selected all papers from the main research track of the latest six editions of the International Conference on Software Engineering (ICSE 2013 to ICSE 2018) that mention “empirical” in their title or in their section’s name in the proceedings. This gave 25 papers, from which we discarded one [90] that turned out not to be an empirical study. The experimental data in the remaining 24 papers come from various sources: the output of analyzers and other tools [19], [20], [21], [69], [76], the mining of repositories of software and other artifacts [15], [59], [61], [84], [102], [103], the outcome of controlled experiments involving human subjects [75], [93], [100], interviews and surveys [6], [9], [26], [53], [55], [60], [79], [86], and a literature review [88].

As one would expect from a top-tier venue like ICSE, the 24 papers follow recommended practices in reporting and analyzing data, using significance testing (6 papers), effect sizes (5 papers), correlation coefficients (5 papers), frequentist regression (2 papers), and visualization in charts or tables (23 papers). None of the papers, however, uses Bayesian statistics. In fact, no paper but two [26], [102] even mentions the terms “Bayes” or “Bayesian”. One of the exceptions [102] only cites Bayesian machine-learning techniques used in related work to which it compares. The other exception [26] includes a presentation of the two views of frequentist and Bayesian statistics—with a critique of p -values similar to the one we make in Sect. 2.2—but does not show how the latter can be used in practice. The aim of [26] is investigating the relationship between empirical findings in software engineering and the actual beliefs of programmers about the same topics. To this end, it is based on a survey of programmers whose responses are analyzed using frequentist statistics; Bayesian statistics is mentioned to frame the discussion about the relationship between evidence and beliefs, but does not feature past the introductory second section. Our paper has a more direct aim: to concretely show how Bayesian analysis can be applied in practice in empirical software engineering

1. [86] has a much more extensive literature survey of empirical publications in software engineering.

research, as an alternative to frequentist statistics; thus, its scope is complementary to [26]’s.

As additional validation based on a more specialized venue for empirical software engineering, we also inspected all 105 papers published in Springer’s Empirical Software Engineering (EMSE) journal during the year 2018. Only 22 papers mention the word “Bayes”: 17 of them refer to Bayesian machine learning classifiers (such as naive Bayes or Bayesian networks); 2 of them discuss Bayesian optimization algorithms for machine learning (such as latent Dirichlet allocation word models); 3 of them mention “Bayes” only in the title of some bibliography entries. None of them use Bayesian statistics as a replacement of classic frequentist data analysis techniques.

More generally, we are not aware of any direct application of Bayesian data analysis to empirical software engineering data with the exception of [31], [32] and [29]. The technical report [31] and its short summary [32] are our preliminary investigations along the lines of the present paper. Ernst [29] presents a conceptual replication of an existing study to argue the analytical effectiveness of multi-level Bayesian models.

Criticism of the p -value. Statistical hypothesis testing—and its summary outcome, the p -value—has been customary in experimental science for many decades, both for the influence [80] of its proponents Fisher, Neyman, and Pearson, and because it offers straightforward, ready-made procedures that are computationally simple². More recently, criticism of frequentist hypothesis testing has been voiced in many experimental sciences, such as psychology [22], [87] and medicine [44], that used to rely on it heavily, as well as in statistics research itself [36], [97], [98]. The criticism, which we articulate in Sect. 2.2, concludes that p -value-based hypothesis testing should be abandoned [2], [63].³ There has been no similar explicit criticism of p -values in software engineering research, and in fact statistical hypothesis testing is still regularly used [42].

Guidelines for using statistics. Best practices of using statistics in empirical software engineering are described in books [65], [99] and articles [4], [51]. Given their focus on frequentist statistics,⁴ they all are complementary to the present paper, whose main goal is showing how Bayesian techniques can add to, or replace, frequentist ones, and how they can be applied in practice.

2 AN OVERVIEW OF BAYESIAN STATISTICS

Statistics provides models of *events*, such as the output of a randomized algorithm; the probability function \mathbb{P} assigns probabilities—values in the real unit interval $[0, 1]$, or equivalently percentages in $[0, 100]$ —to events. Often, events are the values taken by *random variables* that follow a certain probability *distribution*. For example, if X is a random

2. With modern computational techniques it is much less important whether statistical methods are computationally simple; we have CPU power to spare—especially if we can trade it for stronger scientific results.

3. Even statisticians who still accept null-hypothesis testing recognize the need to change the way it is normally used [11].

4. [4], [51], [99] do not mention Bayesian techniques; [5] mentions their existence only to declare they are out of scope; one chapter [10] of [65] outlines Bayesian networks as a machine learning technique.

variable modeling the throwing of a six-face dice, it means that $\mathbb{P}[x] = 1/6$ for $x \in [1..6]$, and $\mathbb{P}[x] = 0$ for $x \notin [1..6]$ —where $\mathbb{P}[x]$ is a shorthand for $\mathbb{P}[X = x]$, and $[m..n]$ is the set of integers between m and n .

The probability of variables over discrete domains is described by *probability mass functions* (p.m.f. for short); their counterparts over continuous domains are probability density functions (p.d.f.), whose integrals give probabilities. The following presentation mostly refers to continuous domains and p.d.f., although notions apply to discrete-domain variables as well with a few technical differences. For convenience, we may denote a distribution and its p.d.f. with the same symbol; for example, random variable X has a p.d.f. also denoted X , such that $X[x] = \mathbb{P}[x] = \mathbb{P}[X = x]$.

Conditional probability. The *conditional* probability $\mathbb{P}[h \mid d]$ is the probability of h given that d has occurred. For example, d may represent the empirical *data* that has been *observed*, and h is a *hypothesis* that is being tested.

Consider a static analyzer that outputs \top (resp. \perp) to indicate that the input program never overflows (resp. may overflow); $\mathbb{P}[\text{OK} \mid \top]$ is the probability that, when the algorithm outputs \top , the input program is indeed free from overflows—the data is the output “ \top ” and the hypothesis is “the input does not overflow”.

2.1 Bayes’ Theorem

Bayes’ theorem connects the conditional probabilities $\mathbb{P}[h \mid d]$ (the probability that the hypothesis is correct given the experimental data) and $\mathbb{P}[d \mid h]$ (the probability that we would see this experimental data given that the hypothesis actually was true). The famous theorem states that

$$\mathbb{P}[h \mid d] = \frac{\mathbb{P}[d \mid h] \cdot \mathbb{P}[h]}{\mathbb{P}[d]}. \quad (1)$$

Here is an example of applying Bayes’ theorem.

Example 2.1. Suppose that the static analyzer gives true positives and true negatives with high probability ($\mathbb{P}[\top \mid \text{OK}] = \mathbb{P}[\perp \mid \text{ERR}] = 0.99$), and that many programs are affected by some overflow errors ($\mathbb{P}[\text{OK}] = 0.01$). Whenever the analyzer outputs \top , what is the chance that the input is indeed free from overflows? Using Bayes’ theorem, $\mathbb{P}[\text{OK} \mid \top] = (\mathbb{P}[\top \mid \text{OK}] \cdot \mathbb{P}[\text{OK}]) / \mathbb{P}[\top] = (\mathbb{P}[\top \mid \text{OK}] \cdot \mathbb{P}[\text{OK}]) / (\mathbb{P}[\top \mid \text{OK}] \cdot \mathbb{P}[\text{OK}] + \mathbb{P}[\top \mid \text{ERR}] \mathbb{P}[\text{ERR}]) = (0.99 \cdot 0.01) / (0.99 \cdot 0.01 + 0.01 \cdot 0.99) = 0.5$, we conclude that we can have a mere 50% confidence in the analyzer’s output.

Priors, likelihoods, and posteriors. In Bayesian analysis [28], each factor of (1) has a special name:

- 1) $\mathbb{P}[h]$ is the *prior*—the probability of the hypothesis h before having considered the data—written $\pi[h]$;
- 2) $\mathbb{P}[d \mid h]$ is the *likelihood* of the data d under hypothesis h —written $\mathcal{L}[d; h]$;
- 3) $\mathbb{P}[d]$ is the *normalizing constant*;
- 4) and $\mathbb{P}[h \mid d]$ is the *posterior*—the probability of the hypothesis h after taking data d into account—written $\mathcal{P}_d[h]$.

With this terminology, we can state Bayes’ theorem (1) as “the posterior is proportional to the likelihood times the prior”, that is,

$$\mathbb{P}[h \mid d] \propto \mathbb{P}[d \mid h] \times \mathbb{P}[h], \quad (2)$$

and hence we *update* the prior to get the posterior.

The only role of the normalizing constant is ensuring that the posterior defines a correct probability distribution when evaluated over all hypotheses. In most cases we deal with hypotheses $h \in H$ that are mutually exclusive and exhaustive; then, the normalizing constant is simply $\mathbb{P}[d] = \int_{h \in H} \mathbb{P}[d | h] \mathbb{P}[h] dh$, which can be computed from the rest of the information. Thus, it normally suffices to define likelihoods that are *proportional* to a probability, and rely on the *update rule* to normalize them and get a proper probability distribution as posterior.

2.2 Frequentist vs. Bayesian Statistics

Despite being a simple result about an elementary fact in probability, Bayes' theorem has significant implications for statistical reasoning. We do not discuss the philosophical differences between how frequentist and Bayesian statistics interpret their results [23]. Instead, we focus on describing how some features of Bayesian statistics support new ways of analyzing data. We start by criticizing statistical hypothesis testing since it is a customary technique in frequentist statistics that is widely applied in empirical software engineering research, and suggest how Bayesian techniques could provide more reliable analyses. For a systematic presentation of Bayesian statistics see Kruschke [57], McElreath [62], and Gelman et al. [37].

2.2.1 Hypothesis Testing vs. Model Comparison

A primary goal of experimental science is validating models of behavior based on empirical data. This often takes the form of choosing between alternative hypotheses, such as, in the software engineering context, deciding whether automatically generated tests support debugging better than manually written tests (Sect. 3), or whether a programming language is faster than another (Sect. 4).

Hypothesis testing is the customary framework offered by frequentist statistics to choose between hypotheses. In the classical setting, a null hypothesis H_0 corresponds to “no significant difference” between two *treatments* A and B (such as two static analysis algorithms whose effectiveness we want to compare); an alternative hypothesis H_1 is the null hypothesis's negation, which corresponds to a significant difference between applying A and applying B . A *null hypothesis significance test* [5], such as the t -test or the U -test, is a procedure that takes as input a combination D^* of two datasets D_A and D_B , respectively recording the outcome of applying A and B , and outputs a probability called the p -value.

The p -value is the probability, under the null hypothesis, of observing data at least as extreme as the ones that were actually observed. Namely, it is the probability $\mathbb{P}[D \geq D^* | H_0]$ ⁵ of drawing data D that is equal to the observed data D^* or more “extreme”, conditional on the null hypothesis H_0 holding. As shown visually in Fig. 1, data is expected to follow a certain distribution under the null hypothesis (the actual distribution depends on the statistical test that is used); the p -value measures the tail

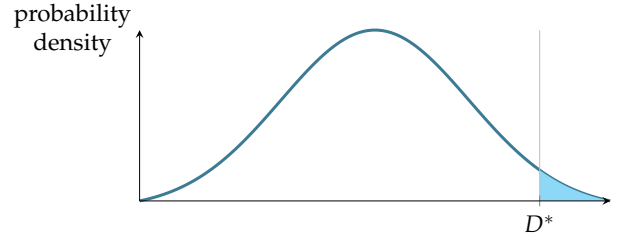


Fig. 1: The p -value is the probability, under the null hypothesis, of drawing data that is at least as extreme as the observed data D^* . Graphically, the p -value is the shaded area under the curve, which models the probability distribution under the null hypothesis.

probability of drawing data equal to the observed D^* or more unlikely than it—in other words, how far the observed data is from the most likely observations under the null hypothesis. If the p -value is sufficiently small—typically $p \leq 0.05$ or $p \leq 0.01$ —it means that the null hypothesis is an unlikely explanation of the observed data. Thus, one *rejects* the null hypothesis, which corresponds to leaning towards preferring the alternative hypothesis H_1 over H_0 : in this case, we have increased our confidence that A and B differ.

Unfortunately, this widely used approach to testing hypotheses suffers from serious shortcomings [22], which have prompted calls to seriously reconsider how it's used as a statistical practice [11], [97], or even to abandon it altogether [2], [63]. The most glaring problem is that, in order to decide whether H_0 is a plausible explanation of the data, we would need the conditional probability $\mathbb{P}[H_0 | D \geq D^*]$ of the hypothesis given the data, not the p -value $\mathbb{P}[D \geq D^* | H_0]$. The two conditional probabilities are related by Bayes' theorem (1), but knowing only $\mathbb{P}[D \geq D^* | H_0]$ is not enough to determine $\mathbb{P}[H_0 | D \geq D^*]$;⁶ in fact, Sect. 2.1's example of the static analyzer (Ex. 2.1) showed a case where one conditional probability is 99% while the other is only 50%.

Rejecting the null hypothesis in response to a small p -value looks like a sound inference, but in reality it is not, because it follows an unsound probabilistic extension of a reasoning rule that is perfectly sound in Boolean logic. The sound rule is *modus tollens*:

$$\begin{array}{l} \text{if } X \text{ implies that } Y \text{ is false} \\ \text{and we observe } Y \\ \text{then } X \text{ is false} \end{array} \quad \frac{X \longrightarrow \neg Y \quad Y}{\neg X} \quad (3)$$

For example, if X means “a person lives in Switzerland” and Y means “a person is the King of Sweden”, if we observe a person that is the Swedish king we can conclude that he does not live in Switzerland. The probabilistic extension of modus tollens, which is unsound, is:

$$\begin{array}{l} \text{if } X \text{ implies that } Y \text{ is } \textit{probably} \text{ false} \\ \text{(equivalently: } Y \text{ is } \textit{improbably} \text{ true)} \\ \text{and we observe } Y \\ \text{then } X \text{ is } \textit{probably} \text{ false} \end{array} \quad \frac{\mathbb{P}[Y | X] < \epsilon \quad Y}{\mathbb{P}[X] < \epsilon} \quad (4)$$

5. Depending on the nature of the data, the p -value may also be defined as a left-tail event $\mathbb{P}[D \leq D^* | H_0]$ or as a 2-tailed event $2 \cdot \min(\mathbb{P}[D \leq D^* | H_0], \mathbb{P}[D \geq D^* | H_0])$.

6. Assuming that they are equal is the “confusion of the inverse” [24].

To see that rule (4) is unsound let's consider the case where Y means "a person is the Queen of England" and X means "a person lives in London". Clearly, $\mathbb{P}[Y \mid X]$ is small because only one person out of million Londoners is indeed the Queen of England. However, if the observed Y happens to be the Queen of England, we would be wrong to infer that she's unlikely to live in London; indeed, the sound conclusion is that she *definitely* lives in London. Whenever we reject $X = H_0$ after observing $Y = D$ following a calculation that $\mathbb{P}[Y \mid X] = p$ is small, we are applying this unsound rule.

This unsound inference is not the only issue with using p -values. Methodological problems come from how hypothesis testing pits the null hypothesis against the alternative hypothesis: as the number of observations grows, it becomes increasingly likely that *some* effect is detectable (or, equivalently, it becomes increasingly unlikely that *no* effects are), which leads to rejecting the null hypothesis, independent of the alternative hypothesis, just because it is unreasonably restrictive (that is, it's generally unlikely that there are genuinely no effects). This problem may result both in suggesting that some negligible effect is significant just because we reject the null hypothesis, and in discarding some interesting experimental results just because they fail to trigger the arbitrary $p < 0.05$ threshold of significance. This is part of the more general problem with insisting on a dichotomous view between two alternatives: a better approach would be based on richer statistics than one or few summary values (such as the p -value) and would combine quantitative and qualitative data (for example, about the sign and magnitude of an effect) to get richer pictures.

To sum up, null hypothesis testing has both technical problems (p -values are insufficient statistics to conclude anything about the null hypothesis) and methodological ones (null hypotheses often are practically useless models).

2.2.2 Types of Statistical Inference Errors

Frequentist statistics often analyze the probability of so-called *type 1* and *type 2* errors:

Type 1: a type 1 error is a *false positive*: rejecting a hypothesis when it is actually true;

Type 2: a type 2 error is a *false negative*: accepting a hypothesis when it is actually false.

In the context of hypothesis testing, the probability of a type 1 error is related to the p -value, but the p -value alone does *not* "denote the probability of a Type I error" [5]: the p -value is a probability conditional on H_0 being true, and hence it cannot be used to conclude whether H_0 itself is true—an assessment that is required to estimate the probability of a type 1 error, which occurs, by definition, only if H_0 is true. In the context of experimental design, the p -value's *rejection threshold* α is an *upper bound* on the probability of a type 1 error [45]. In this case, a type 1 error occurs when H_0 is true but $p < \alpha$; in order to assign a probability to this scenario, we need to assign a *prior probability* that the null hypothesis is true:

$$\mathbb{P}[\text{type 1 error}] = \mathbb{P}[p < \alpha \mid H_0] \cdot \mathbb{P}[H_0] = \alpha \cdot \mathbb{P}[H_0] \quad (5)$$

Therefore, the probability of a type 1 error in an experiment with rejection threshold set to α lies between zero and α .⁷

Note, however, that even the technically correct interpretation of the significance level α as an upper bound on the probability of a type 1 error breaks down when we have *multiple testing*, as we illustrate in Sect. 4.2.2. In all cases, the p -value is a statistics of the data alone, and hence it is insufficient to estimate the probability of a hypothesis being true.

Bayesian analysis prefers to avoid the whole business of hypothesis testing and instead computing actual posterior probability distributions based on data and priors. The Bayesian viewpoint also suggests that the usual frequentist preoccupation with type 1 and 2 errors is often misplaced: the null hypothesis captures a narrow notion of "no effects", which is rarely the case. As Gelman remarks:⁸

A type 1 error occurs only if the null hypothesis is true (typically if a certain parameter, or difference in parameters, equals zero). In the applications I've worked on [...] I've never come across a null hypothesis that could actually be true. A type 2 error occurs only if I claim that the null hypothesis is true, and I would certainly not do that, given my statement above!

Instead, Gelman recommends [40] focusing on type S and type M errors. A type S error occurs when we mistakenly infer the *sign* of an effect: for example that a language A is consistently faster than language B , when in fact it is B that is faster than A . A type M error occurs when we overestimate the *magnitude* of an effect: that language A is 3 times faster than language B , when in fact it is at most 5% faster. Sect. 4.3.2 discusses type S and type M errors in Nanz and Furia [71]'s case study.

2.2.3 Scalar Summaries vs. Posterior Distributions

Even though hypothesis testing normally focuses on two complementary hypotheses (H_0 and its negation H_1), it is often necessary to simultaneously compare several alternative hypotheses $h \in H$ grouped in a discrete or even dense set H . For example, in Sect. 4's reanalysis, H is the continuous interval of all possible *speedups* of one language relative to another. A distinctive advantage of full-fledged Bayesian statistics is that it supports deriving a complete *distribution* of posterior probabilities, by applying (1) for all hypotheses $h \in H$, rather than just scalar summaries (such as estimators of mean, median, and standard deviation, standardized measures of effect size, or p -values on discrete partitions of H). Given a distribution we can still compute scalar summaries (including confidence/credible intervals), but we retain additional advantages such as being able to visualize the distribution and to derive other distributions by iterative application of Bayes' theorem. This supports decisions based on a variety of criteria and on a richer understanding of the experimental data, as we demonstrate in the case studies of Sect. 3 and Sect. 4.

7. You will still find the erroneous statement that α is the probability of a type 1 error in numerous publications including some bona fide statistics textbooks [82], [96].

8. Andrew Gelman's blog *Type 1, type 2, type S, and type M errors*, 29 December 2004.

2.2.4 The Role of Prior Information

The other distinguishing feature of Bayesian analysis is that it *starts from a prior probability* which models the initial knowledge about the hypotheses. The prior can record previous results in a way that is congenial to the way science is supposed to work—not as completely independent experiments in an epistemic vacuum, but by constantly scrutinizing previous results and updating our models based on new evidence.

A canned criticism of Bayesian statistics observes that using a prior is a potential source of bias. However, explicitly taking into account this very fact helps analyses be more rigorous and more comprehensive. In particular, we often consider several different alternative priors to perform Bayesian analysis. Priors that do not reflect any strong a priori bias are traditionally called *uninformative*; a uniform distribution over hypotheses is the most common example. Note, however, that the term “uninformative” is misleading: a uniform prior encodes as much information as a non-uniform one—it is just a different kind of information.⁹ In a different context, this misunderstanding is illustrated metaphorically by this classic *koan* [81, Appendix A] involving artificial intelligence pioneers Marvin Minsky and Gerald Sussman:

In the days when Sussman was a novice, Minsky once came to him as he sat hacking at the PDP-6. “What are you doing?”, asked Minsky. “I am training a randomly wired neural net to play Tic-Tac-Toe” Sussman replied. “Why is the net wired randomly?”, asked Minsky. “I do not want it to have any preconceptions of how to play”, Sussman said. Minsky then shut his eyes. “Why do you close your eyes?”, Sussman asked his teacher. “So that the room will be empty.” At that moment, Sussman was enlightened.

A uniform prior is the equivalent a randomly wired neural network—with unbiased assumptions but still with specific assumptions.

Whenever the posterior distribution is largely independent of the chosen prior, we say that the data *swamps* the prior, and hence the experimental evidence is decidedly strong. If, conversely, choosing a suitable, biased prior is necessary to get sensible results, it means that the evidence is not overwhelming, and hence any additional reliable source of information should be vetted and used to sharpen the analysis results. In these cases, characteristics of the biased priors may suggest what kinds of assumptions about the data distribution we should investigate further. For example, suppose that sharp results emerge only if we use a prior that assigns vanishingly low probabilities to a certain parameter x taking negative values. In this case, the analysis should try to ascertain whether x can or cannot meaningfully take such values: if we find theoretical or practical reasons why x should be mostly positive, the biased prior acts as a safeguard against spurious data; conversely, if x could take negative values in other experiments, we conclude that the results obtained with the biased prior

⁹ In fact, in a statistical inference context, a uniform prior often leads to overfitting because we learn too much from the data without correction [62].

are only valid in a restricted setting, and new experiments should focus on the operational conditions where x may be negative.

Assessing the impact of priors with markedly different characteristics leads to a process called (*prior sensitivity analysis*), whose goals are clearly described by McElreath [62, Ch. 9]:

A sensitivity analysis explores how changes in assumptions influence inference. If none of the alternative assumptions you consider have much impact on inference, that’s worth reporting. Likewise, if the alternatives you consider do have an important impact on inference, that’s also worth reporting. [...] In sensitivity analysis, many justifiable analyses are tried, and all of them are described.

Thus, the goal of sensitivity analysis is *not* choosing a prior, but rather better understanding the statistical model and its features. Sect. 4.3.3 presents a prior sensitivity analysis for some of the language comparisons of Nanz and Furia [71]’s data.

To sum up, Bayesian analysis stresses the importance of careful *modeling* of assumptions and hypotheses, which is more conducive to accurate analyses than the formulaic application of ready-made statistics. In the rest of the paper, we illustrate how frequentist and Bayesian data analysis work in practice on two case studies based on software engineering empirical data. We will see that the strength of Bayesian statistics translates into analyses that are more robust, sound, and whose results are natural to interpret in practical terms.

2.3 Techniques for Bayesian Inference

No matter how complex the model, every Bayesian analysis ultimately boils down to computing a *posterior probability distribution*—according to Bayes’ theorem (1)—given likelihood, priors, and data. In most cases the posterior cannot be computed *analytically*, but numerical approximation algorithms exist that work as well for all practical purposes. Therefore, in the remainder of this article we will always mention the “posterior distribution” as if it were computed exactly; provided we use reliable tools to compute the approximations, the fact that we are actually dealing with numerical approximations does not make much practical difference.

The most popular numerical approximation algorithms for Bayesian inference work by *sampling* the posterior distribution at different points; for each sampled point \bar{h} , we calculate the posterior probability $\mathcal{P}_d[\bar{h}] = \mathbb{P}[\bar{h} \mid d]$ by evaluating the right-hand side of (1) at \bar{h} for the given likelihood and data d . Sampling must be neither too sparse (resulting in an inaccurate approximation) nor too dense (taking too much time). The most widely used techniques achieve a simultaneously effective and efficient sampling by using Markov Chain Monte Carlo (MCMC) [14] randomized algorithms. In a nutshell, the sampling process corresponds to visiting a Markov chain [33, Ch. 6], whose transition probabilities reflect the posterior probabilities in a way that ensures that they are sampled densely enough (after the

chain reaches its steady state). Tools such as Stan [16] provide scalable generic implementations of these techniques.

Users of Bayesian data analysis tools normally need not understand the details of how Bayesian sampling algorithms work. As long as they use tried-and-true statistical models and follow recommended guidelines (such as those we outline in Sect. 5), validity of analysis results should not depend on the details of how the numerical approximations are computed.

3 LINEAR REGRESSION MODELING: SIGNIFICANCE AND PREDICTION

Being able to automatically generate test cases can greatly help discover bugs quickly and effectively: the unit tests generated by tools such as Randoop (using random search [74]) are likely to be complementary to those a human programmer would write, and hence have a high chance of exercising behavior that may reveal flaws in software. Conversely, the very features that distinguish automatically-generated tests from programmer-written ones—such as uninformative identifiers and a random-looking nature [25]—may make the former less effective than the latter when programmers use tests to *debug* faulty programs. Ceccato et al. [18] designed a series of experiments to study this issue: *do automatically generated tests provide good support for debugging?*

3.1 Data: Automated vs. Manual Testing

Ceccato et al. [18] compare automatically-generated and manually-written tests—henceforth called ‘autogen’ and ‘manual’ tests—according to different features. The core experiments are empirical studies involving, as subjects, students and researchers, who debugged Java systems based on the information coming from autogen or manual tests.

Our reanalysis focuses on Ceccato et al. [18]’s first study, where autogen tests were produced running *Randoop* and the success of debugging was evaluated according to the experimental subjects’ *effectiveness*—namely, how many of the given bugs they could successfully fix given a certain time budget.

Each experimental data point assigns a value to six variables:

- fixed:** the dependent variable measuring the (nonnegative) number of bugs correctly fixed by the subject;
- treatment:** *auto* (the subject used autogen tests) or *manual* (the subject used manual tests);
- system:** *J* (the subject debugged Java application JTopas) or *X* (the subject debugged Java library XML-Security);
- lab:** 1 or 2 according to whether it was the first or second of two debugging sessions each subject worked in;
- experience:** *B* (the subject is a bachelor’s student) or *M* (the subject is a master’s student);
- ability:** *low*, *medium*, or *high*, according to a discretized assessment of ability based on the subject’s GPA and performance in a training exercise.

For further details about the experimental protocol see Ceccato et al. [18].

3.2 Frequentist Analysis

3.2.1 Linear Regression Model

Ceccato et al. [18]’s analysis fits a linear regression model that expresses outcome variable *fixed* as a linear combination of five predictor variables:

$$\begin{aligned} \text{fixed} = c_1 + c_T \cdot \text{treatment} + c_S \cdot \text{system} + c_L \cdot \text{lab} \\ + c_E \cdot \text{experience} + c_A \cdot \text{ability} + \epsilon \end{aligned} \quad (6)$$

where each variable’s value is encoded on an integer interval scale, and ϵ models the error as a normally distributed random variable with mean zero and unknown variance. We can rewrite (6) as:

$$\text{fixed} \sim \mathcal{N} \left(\begin{array}{c} c_1 + c_T \cdot \text{treatment} + c_S \cdot \text{system} + c_L \cdot \text{lab} \\ + c_E \cdot \text{experience} + c_A \cdot \text{ability} \end{array}, \sigma \right) \quad (7)$$

where $\mathcal{N}(\mu, \sigma)$ is a normal distribution with mean μ and (unknown) variance σ^2 . Form (7) is equivalent to (6), but better lends itself to the generalizations we will explore later on in Sect. 3.3.

Tab. 1 shows the results of fitting model (7) using a standard least-squares algorithm. Each predictor variable v ’s coefficient c_v gets a maximum-likelihood *estimate* \bar{c}_v and a *standard error* e_v of the estimate. The numbers are a bit different from those in Table III of Ceccato et al. [18] even if we used their replication package; after discussion¹⁰ with the authors of [18], we attribute any remaining differences to a small mismatch in how the data has been reported in the replication package. Importantly, the difference does not affect the overall findings, and hence Tab. 1 is a sound substitute of Ceccato et al. [18]’s Table III for our purposes.

Null hypothesis. To understand whether any of the predictor variables v in (7) makes a significant contribution to the number of *fixed* bugs, it is customary to define a *null hypothesis* H_0^v , which expresses the fact that debugging effectiveness does *not* depend on the value of v . For example, for predictor *treatment*:

$H_0^{\text{treatment}}$: there is no difference in the effectiveness of debugging when debugging is supported by autogen or manual tests.

Under the linear model (7), if H_0^v is true then regression coefficient c_v should be equal to zero. However, error term ϵ makes model (7) stochastic; hence, we need to determine the *probability* that some coefficients are zero. If the probability that c_v is close to zero is *small* (typically, less than 5%) we customarily say that the corresponding variable v is a *statistically significant predictor* of *fixed*; and, correspondingly, we reject H_0^v .

p -values. The other columns in Tab. 1 present standard frequentist statistics used in hypothesis testing. Column t is the t -statistic, which is simply the coefficient’s estimate divided by its standard error. One could show [89] that, if the null hypothesis H_0^v were *true*, its t statistic t_v over repeated experiments would follow a t distribution $t(n_o - n_\beta)$ with $n_o - n_\beta$ degrees of freedom, where n_o is the number of observations (the sample size), and n_β is the number of coefficients in (7).

The p -value is the probability of observing data that is, under the null hypothesis, at least as extreme as the one

¹⁰ Personal communication between Torkar and Ceccato, December 2017.

TABLE 1: Regression coefficients in the linear model with Gaussian error model (7), computed with frequentist analysis. For each coefficient c_v , the table reports maximum likelihood *estimate* \bar{c}_v , standard *error* e_v of the estimate, *t statistic* of the estimate, *p-value* (the probability of observing the *t* statistic under the null hypothesis), and *lower* and *upper* endpoints of the 95% confidence interval $C_{95\%}^v$ of the coefficient.

COEFFICIENT c_v	ESTIMATE \bar{c}_v	ERROR e_v	<i>t</i> STATISTIC	<i>p</i> -VALUE	LOWER	UPPER
c_I (intercept)	-1.756	0.725	-2.422	0.020	-3.206	-0.306
c_T (treatment)	0.651	0.330	1.971	0.055	-0.010	1.311
c_S (system)	-0.108	0.332	-0.325	0.747	-0.772	0.556
c_L (lab)	0.413	0.338	1.224	0.228	-0.262	1.089
c_E (experience)	0.941	0.361	2.607	0.013	0.219	1.663
c_A (ability)	0.863	0.297	2.904	0.006	0.269	1.457

that was actually observed. Thus, the *p*-value for v is the probability that t_v takes a value smaller than $-|\bar{t}_v|$ or greater than $+|\bar{t}_v|$, where \bar{t}_v is the value of t_v observed in the data; in formula:

$$p_v = \int_{-\infty}^{-|\bar{t}_v|} \mathcal{T}(x) dx + \int_{+|\bar{t}_v|}^{+\infty} \mathcal{T}(x) dx, \quad (8)$$

where \mathcal{T} is the density function of distribution $t(n_o - n_\beta)$. Column *p*-value in Tab. 1 reports *p*-values for every predictor variable computed according to (8).

3.2.2 Significance Analysis

Significant predictors. Using a standard confidence level $\alpha = 0.05$, Ceccato et al. [18]’s analysis concludes that:

- *system* and *lab* are *not statistically significant* because their *p*-values are much larger than α —in other words, we cannot reject H_0^{system} or H_0^{lab} ;
- *experience* and *ability* are *statistically significant* because their *p*-values are strictly smaller than α —in other words, we reject $H_0^{\text{experience}}$ and H_0^{ability} ;
- *treatment* is *borderline significant* because its *p*-value is close to α —in other words, we tend to reject $H_0^{\text{treatment}}$.

Unfortunately, as we discussed in Sect. 2, *p*-values cannot reliably assess statistical significance. The problem is that the *p*-value is a probability *conditional on the null hypothesis being true*: $p_v = \mathbb{P}[|t_v| > |\bar{t}_v| \mid H_0^v]$; if $p_v \ll \alpha$ is small, the data is unlikely to happen under the null hypothesis, but this is not enough to soundly conclude whether the null hypothesis itself is likely true.

Another, more fundamental, problem is that null hypotheses are simply too restrictive, and hence unlikely to be literally true. In our case, a coefficient may be statistically significant from zero but still very small in value; from a practical viewpoint, this would be as if it were zero. In other words, statistical significance based on null hypothesis testing may be irrelevant to assess *practical significance* [42].

Confidence intervals. Confidence intervals are a better way of assessing statistical significance based on the output of a fitted linear regression model. The most common 95% confidence interval is based on the normal distribution, i.e.,

$$C_{95\%}^v = (\bar{c}_v - 2e_v, \bar{c}_v + 2e_v)$$

is the 95% confidence interval for a variable v whose coefficient c_v has estimate \bar{c}_v and standard error e_v given by the linear regression fit.

If confidence interval $C_{95\%}^v$ includes zero, it means that there is a significant chance that c_v is zero; hence, we conclude that v is not significant with 95% confidence. Conversely, if $C_{95\%}^v$ for variable v does not include zero, it includes either only positive or only negative values; hence, there is a significant chance that c_v is not zero, and we conclude that v is significant with 95% confidence.

Even though, with Ceccato et al. [18]’s data, it gives the same qualitative results as the analysis based on *p*-values, grounding significance on confidence intervals has the distinct advantage of providing a *range of values* the coefficients of a linear regression model may take based on the data—instead of the *p*-value’s conditional probability that really answers not quite the right question.

Interpreting confidence intervals. It is tempting to interpret a frequentist confidence intervals $C_{95\%}^v$ in terms of probability of the possible values of the regression coefficient c_v of variable v .¹¹ Nonetheless the correct frequentist interpretation is much less serviceable [67]. First, the 95% probability does not refer to the values of c_v relative to the specific interval $C_{95\%}^v$ we computed, but rather to the *process* used to compute the confidence interval: in an infinite series of repeated experiments—in each of which we compute a different interval $C_{95\%}^v$ —the fraction of repetitions where the *true* value of c_v falls in $C_{95\%}^v$ will tend to 95% in the limit. Second, confidence intervals have no distributional information: “there is no direct sense by which parameter values in the middle of the confidence interval are more probable than values at the ends of the confidence interval” [58].

To be able to soundly interpret $C_{95\%}^v$ in terms of probabilities we need to assign a *prior* probability to the regression coefficients and apply Bayes’s theorem. A natural assumption is to use a uniform prior; then, under the normal model for the error term ϵ in (6), v behaves like a random variable with a normal distribution with mean equal to v ’s linear regression estimate \bar{c}_v and standard deviation equal to v ’s standard error e_v . Therefore, $C_{95\%}^v = (\ell, h)$ is the (centered) interval such that

$$\int_{\ell}^h \mathcal{N}(x) dx = 0.95,$$

where \mathcal{N} is the normal density function of distribution $\mathcal{N}(\bar{c}_v, e_v)$ —with mean equal to c_v ’s estimate and standard deviation equal to c_v ’s standard error.

11. This interpretation is widespread even though it is incorrect [47].

However, this entails that if $C_{95\%}^V$ lies to the right of zero ($\ell > 0$) then the probability that the true value of c_V is positive is

$$\begin{aligned} \int_0^{+\infty} \mathcal{N}(x) dx &\geq \int_{\ell}^{+\infty} \mathcal{N}(x) dx \\ &= \int_{\ell}^h \mathcal{N}(x) dx + \int_h^{+\infty} \mathcal{N}(x) dx \\ &= 0.95 + 0.025, \end{aligned}$$

that is at least 97.5%.

3.2.3 Summary and Criticism of Frequentist Statistics

In general, assessing statistical significance using p -values is hard to justify methodologically, since the conditional probability that p -values measure is neither sufficient nor necessary to decide whether the data supports or contradicts the null hypothesis. Confidence intervals might be a better way to assess significance; however, they are surprisingly tricky to interpret in a purely frequentist setting as perusing Ceccato et al. [18]’s analysis has shown. More precisely, we have seen that any sensible, intuitive interpretation of confidence intervals requires the introduction of *priors*. We might as well go all in and perform a full-fledged Bayesian analysis instead; as we will demonstrate in Sect. 3.3, this generally brings greater flexibility and other advantages.

3.3 Bayesian Reanalysis

3.3.1 Linear Regression Model

As a first step, let us apply Bayesian techniques to fit the very same linear model (7) used in Ceccato et al. [18]’s analysis. Tab. 2 summarizes the Bayesian fit with statistics similar to those used in the frequentist setting: for each predictor variable v , we estimate the value and standard error of multiplicative coefficient c_V —as well as the endpoints of the 95% credible interval (last two columns in Tab. 2).

The 95% *uncertainty interval* for variable v is the centered¹² interval (ℓ, h) such that

$$\int_{\ell}^h \mathcal{P}_V(x) dx = 0.95,$$

where \mathcal{P}_V is the posterior distribution of coefficient c_V . Since \mathcal{P}_V is the actual distribution observed on the data, the uncertainty interval can be interpreted in a natural way as the range of values such that $c_V \in (\ell, h)$ with probability 95%.

Uncertainty intervals are traditionally called *credible intervals*, but we agree with Gelman’s suggestion¹³ that *uncertainty interval* reflects more clearly what they measure.¹⁴ We could use the same name for the frequentist confidence intervals, but we prefer to use different terms to be able to distinguish them easily in writing.

The numbers in Tab. 2 summarize a multivariate distribution on the regression coefficients. The distribution—called the *posterior distribution*—is obtained, as described in Sect. 2.3, by repeatedly sampling values from a *prior*

distribution, and applying Bayes’ theorem using a likelihood function derived from (7) (in other words, based on the normal error term). The posterior is thus a sampled, numerical distribution, which provides rich information about the possible range of values for the coefficients, and lends itself to all sorts of derived computations—as we will demonstrate in later parts of the reanalysis.

Statistical significance. Using the same conventional threshold used by the frequentist analysis, we stipulate that a predictor v is *statistically significant* if the 95% uncertainty interval of its regression coefficient c_V does not include zero. According to this definition, *treatment*, *experience*, and *ability* are all significant—with *treatment* more weakly so. While the precise estimates differ, the big picture is consistent with the frequentist analysis—but this time the results’ interpretation is unproblematic.

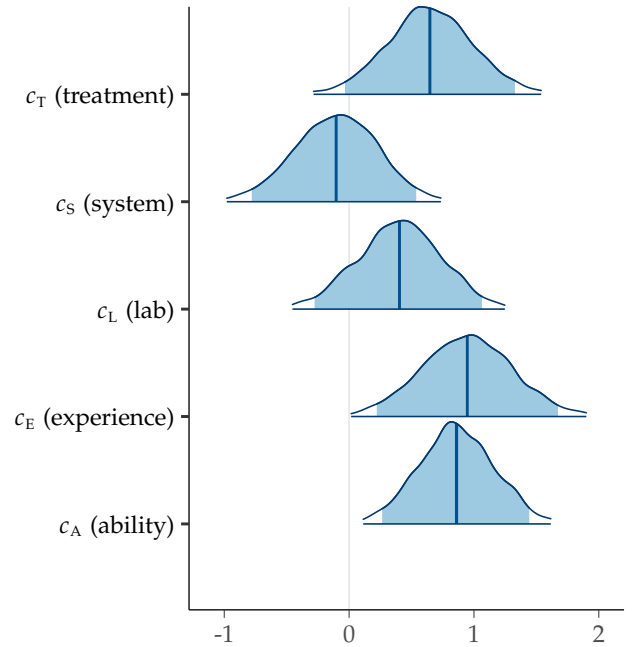


Fig. 2: Plots of the posterior probability distributions of each regression coefficient in the linear model with Gaussian error model (7), computed with Bayesian analysis and weak unbiased priors. Each plot is a density plot that covers an area corresponding to 99% probability; the inner shaded area corresponds to 95% probability; and the thick vertical line marks the distribution’s median.

Since the coefficients’ distributions are fully available we can even plot them. In Fig. 2, the thick vertical lines mark the estimate (the distribution’s median); the shaded areas correspond to 95% probability; and the distribution curves extend to cover the 99% probability overall. Visualization supports a more nuanced view of statistical significance, because the plots include probability distribution information:¹⁵ it is clear that *experience* and *ability* are strongly significant; even though *treatment* is borderline significant, there remains a high probability that it is significant (since

12. That is, centered around the distribution’s mean.

13. Andrew Gelman’s blog *Instead of “confidence interval,” let’s say “uncertainty interval”*, 21 December 2010

14. More recently, other statisticians suggested yet another term: *compatibility interval* [3].

15. Remember that confidence intervals do not possess this kinds of information (see Sect. 3.2.2).

TABLE 2: Regression coefficients in the linear model with Gaussian error model (7), computed with Bayesian analysis. For each coefficient c_V , the table reports posterior mean *estimate* \bar{c}_V , posterior standard deviation (or *error*) e_V of the estimate, and *lower* and *upper* endpoints of the 95% uncertainty interval of the coefficient (corresponding to 2.5% and 97.5% cumulative posterior probability).

COEFFICIENT c_V	ESTIMATE \bar{c}_V	ERROR e_V	LOWER	UPPER
c_I (intercept)	-1.753	0.719	-3.159	-0.319
c_T (treatment)	0.659	0.337	0.011	1.332
c_S (system)	-0.112	0.347	-0.808	0.566
c_L (lab)	0.408	0.344	-0.265	1.069
c_E (experience)	0.947	0.363	0.245	1.670
c_A (ability)	0.863	0.298	0.290	1.443

most of the curve is to the right of the origin); while neither *system* nor *lab* is significant at the 95% value, we can discern a weak tendency towards significance for *lab*—whereas *system* has no consistent effect.

3.3.2 Generalized Models

Priors. The posterior distributions look approximately normal; indeed, we have already discussed how they are exactly normal if we take uniform priors. A uniform prior assumes that any possible value for a coefficient is as likely as any other. In some restricted conditions (see Sect. 3.3.4), and with completely uniform priors, frequentist and Bayesian analysis tend to lead to the same overall numerical results; however, Bayesian analysis is much more flexible because it can compute posterior distributions with priors that are not completely uniform.

A *weak unbiased* prior is the most appropriate choice in most cases. Even when we do not know much about the possible values that the regression coefficients might take, we often have a rough idea of their variability range. Take variable *treatment*: discovering whether it is definitely positive or negative is one of the main goals of the analysis; but we can exclude that *treatment* (indeed, any variable) has a huge effect on the number of fixed bugs (such as hundred or even thousands of bugs difference). Since such a huge effect is ruled out by experience and common sense, we used a normal distribution with mean 0 and standard deviation 20 as prior. This means that we still have no bias on whether *treatment* is significant (mean 0), but we posit that its value is most likely between $-60 = 0 - 3 \cdot 20$ and $+60 = 0 + 3 \cdot 20$ (three standard deviations off the mean).¹⁶ In this case, using a completely uniform prior would not lead to qualitatively different results, but would make the estimate a bit less precise without adding any relevant information.

If we had evidence—from similar empirical studies or other sources of information—that sharpen our initial estimate of an effect, we could use it to pick a more biased prior. We will explore this direction more fully in Sect. 4.3 where we present the other case study.

Poisson data model. To further demonstrate how Bayesian modeling is flexible and supports quantitative

16. Since a normal distribution has infinitely long tails, using it as prior does not rule out any posterior value completely, but makes extreme values exceedingly unlikely without exceptionally strong data to support them (i.e., we make infinity a non-option).

predictions, let us modify (7) to build a *generalized* linear model of the Ceccato et al. [18]’s experiments.

Since *system* and *lab* turned out to be not significant, we exclude them from the new model. Excluding insignificant variables has the advantage of simplifying the model without losing accuracy—thus supporting better generalizations [62]. Then, since *fixed* can only be a nonnegative integer, we model it as drawn from a Poisson distribution (suitable for counting variables) rather than a normal distribution as in (7) (which allows for negative real values as well). Finally, we use a stronger prior for c_T corresponding to a normal distribution $\mathcal{N}(0.5, 0.8)$, which is a bit more biased as it nudges c_T towards positive values of smallish size:

$$fixed \sim \mathcal{P} \left(\exp \left(\begin{array}{l} c_I + c_T \cdot \text{treatment} \\ + c_E \cdot \text{experience} \\ + c_A \cdot \text{ability} \end{array} \right) \right) \quad (9)$$

where $\mathcal{P}(\lambda)$ is the Poisson distribution with rate λ . Since λ has to be positive, it is customary to take the *exponential* of the linear combination of predictor variables as a parameter of the Poisson.

Fitting model (9) gives the estimates in Tab. 3, graphically displayed in Fig. 3. The results are still qualitatively similar to previous analyses: *treatment*, *experience*, and *ability* are significant at the 95% level—*treatment* more weakly so. However, a simpler model is likely, all else being equal, to perform better predictions; let us demonstrate this feature in the next section.

3.3.3 Quantitative Prediction

Simplicity is a definite advantage of linear regression models: simple models are easy to fit and, most important, easy to interpret. After estimating the coefficients in (6), we can use those estimates to perform quantitative predictions. For example, in the linear model (6), $\bar{c}_T = 0.66$ is the average difference in the number of bugs that a programmer working with autogen tests could fix compared to working with manual tests—all other factors staying the same:

$$\begin{aligned} & \left(\begin{array}{l} \text{treatment} = \text{auto} \\ \bar{c}_I + \overbrace{\bar{c}_T \cdot 1}^{\text{treatment} = \text{auto}} + \bar{c}_S \cdot \text{system} + \bar{c}_L \cdot \text{lab} \\ + \bar{c}_E \cdot \text{experience} + \bar{c}_A \cdot \text{ability} + \epsilon \end{array} \right) \\ - & \left(\begin{array}{l} \text{treatment} = \text{manual} \\ \bar{c}_I + \overbrace{\bar{c}_T \cdot 0}^{\text{treatment} = \text{manual}} + \bar{c}_S \cdot \text{system} + \bar{c}_L \cdot \text{lab} \\ + \bar{c}_E \cdot \text{experience} + \bar{c}_A \cdot \text{ability} + \epsilon \end{array} \right) \\ = & \bar{c}_T \end{aligned}$$

TABLE 3: Regression coefficients in the generalized linear model with Poisson error model (9), computed with Bayesian analysis. For each coefficient c_V , the table reports posterior mean *estimate* $\overline{c_V}$, posterior standard deviation (or *error*) e_V of the estimate, and *lower* and *upper* endpoints of the 95% uncertainty interval of the coefficient (corresponding to 2.5% and 97.5% cumulative posterior probability).

COEFFICIENT c_V	ESTIMATE $\overline{c_V}$	ERROR e_V	LOWER	UPPER
c_I (intercept)	-1.953	0.510	-3.006	-0.954
c_T (treatment)	0.493	0.254	0.012	0.991
c_E (experience)	0.800	0.309	0.201	1.400
c_A (ability)	0.642	0.226	0.205	1.096

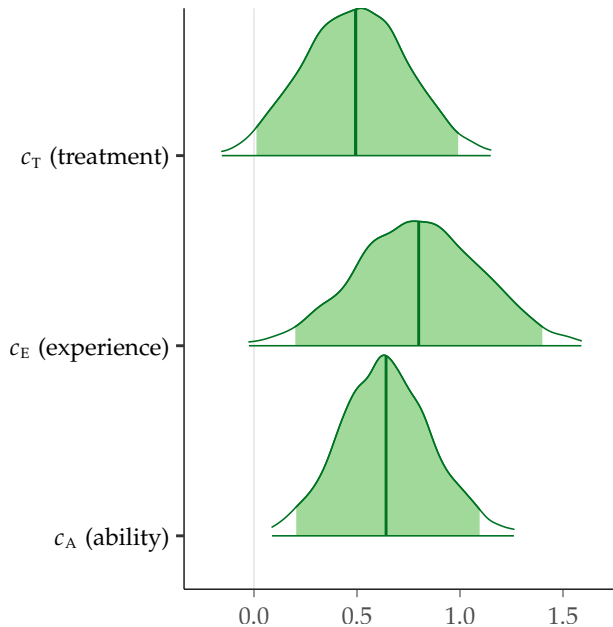


Fig. 3: Plots of the posterior probability distributions of each regression coefficient in the generalized linear model with Poisson error model (9), computed with Bayesian analysis. Each plot is a density plot that covers an area corresponding to 99% probability; the inner shaded area corresponds to 95% probability; and the thick vertical line marks the distribution's median.

Even though the transparency of the linear model is a clear plus, performing estimates and predictions only based on the *estimates* of the regression coefficients has its limits. On the one hand, we ignore the richer information—in the form of posterior distributions—that Bayesian analyses provide. On the other hand, analytic predictions on the fitted model may become cumbersome on more complex models, such as the Poisson regression model (9), where the information carried by regression coefficients is less intuitively understandable. Ideally, we would like to select a model based on how well it can characterize the data without having to trade this off against mainly computational issues.

Since Bayesian techniques are based on numerical methods, they shine at estimating and predicting derived statistics. We can straightforwardly *sample* the predictor input

space in a way that reflects the situation under study, and *simulate* the outcomes numerically.¹⁷ Provided the sample is sufficiently large, the average outcome reflects the fitted model's predictions. Let us show one such simulation to try to answer a relevant and practical software engineering question based on the data at hand.

Our analysis has shown that autogen tests are better—lead to more fixed bugs—than manual tests; and that programmers with higher ability are also better at fixing bugs. Since high-ability programmers are a clearly more expensive resource than automatically generated tests, can we use autogen tests extensively to partially make up for the lack of high-ability programmers? We instantiate the question in quantitative terms:

How does the bug-fixing performance of these two teams of programmers compare?

Team low/auto: made of 80% low-ability, 10% medium-ability, and 10% high-ability programmers; using 10% manual tests and 90% autogen tests

Team high/manual: made of 40% low-ability, 40% medium-ability, and 20% high-ability programmers; using 50% manual tests and 50% autogen tests

Simulating the performance of these two teams using model (9) fitted as in Tab. 3, we get that team high/manual fixes 20% more bugs than team low/auto in the same conditions. This analysis tells us that ability is considerably more decisive than the nature of tests, but autogen tests can provide significant help for a modest cost. It seems likely that industrial practitioners could be interested in simulating several such scenarios to get evidence-based support for decisions and improvement efforts. In fact, the posterior distributions obtained from a Bayesian statistical analysis could be used to optimize, by repeated simulation, different variables—the number of high-ability team members and the proportion of manual versus autogen test cases in our example—to reach specific goals identified as important by practitioners.

3.3.4 Summary of Bayesian Statistics

Using Bayesian techniques brings advantages over using frequentist techniques even when the two approaches lead to the same qualitative results. Bayesian models have an unambiguous probabilistic interpretation, which relates the

¹⁷ Since frequentist statistics do not have post-data distributional information, this can be done only within a Bayesian framework.

variables of interest to their probabilities of taking certain values. A fitted model consists of *probability distributions* of the estimated parameters, rather than less informative point estimates. In turn, this supports quantitative estimates and predictions of arbitrary scenarios based on numerical simulations, and fosters nuanced data analyses that are grounded in practically significant measures.

Numerical simulations are also possible in a frequentist setting, but their soundness often depends on a number of subtle assumptions that are easy to overlook or misinterpret. In the best case, frequentist statistics is just an inflexible hard-to-interpret version of Bayesian statistics; in the worst case, it may even lead to unsound conclusions.

Frequentist \neq Bayesian + Uniform Priors. In this case study, the Bayesian analysis’s main results turned out to be numerically quite close to the frequentist’s. We also encountered several cases where the “intuitive” interpretation of measures such as confidence intervals is incorrect under a strict frequentist interpretation but can be justified by assuming a Bayesian interpretation and uniform (or otherwise “uninformative”) priors. It is tempting to generalize this and think that frequentist statistics are mostly Bayesian statistics with uniform priors by other means. Unfortunately, this is in general *not* true: on the contrary, the correspondence between frequentist statistics and Bayesian statistics with uninformative priors holds only in few simple cases.

Specifically, frequentist confidence intervals and Bayesian uncertainty intervals calculations coincide only in the simple case of estimating the mean of a normal distribution—what we did with the simple linear regression model (7)—when the Bayesian calculation uses a particular uninformative prior [52]. In many other applications [67], frequentist confidence intervals markedly differ from Bayesian uncertainty intervals, and the procedures to calculate the former do not normally come with any guarantees that they can be validly interpreted as post-data statistics.

4 DIRECT BAYESIAN MODELING: THE ROLE OF PRIORS

Comparing programming languages to establish which is better at a certain task is an evergreen topic in computer science research; empirical studies have made such comparisons more rigorous and their outcomes more general.

In this line of work, Nanz and Furia [71] analyzed programs in the Rosetta Code wiki [83]. Rosetta Code is a collection of programming tasks, each implemented in various programming languages. The variety of tasks, and the expertise and number of contributors to the Rosetta Code wiki, buttresses a comparison of programming languages under conditions that are more natural than few performance benchmarks, yet still more controlled than empirical studies of large code repositories in the wild.

4.1 Data: Performance of Programming Languages

Nanz and Furia [71] compare eight programming languages for features such as conciseness and performance, based on experiments with a curated selection of programs from the Rosetta Code repository [83]. Our reanalysis focuses on Nanz and Furia [71]’s analysis of running time *performance*.

For each language ℓ among C, C#, F#, Go, Haskell, Java, Python, and Ruby, Nanz and Furia [71]’s experiments involve an ordered set $T(\ell)$ of programming *tasks*, such as sorting algorithms, combinatorial puzzles, and NP-complete problems; each task comes with an explicit sample input. For a task $t \in T(\ell)$, $R(\ell, t)$ denotes the set of running times of any implementations of task t in language ℓ , among those available in Rosetta Code, that ran without errors or timeout on the same sample input. $R(\ell, t)$ is a *set* because Rosetta Code often includes several implementation of the same task (algorithm) in the same language—variants that may explore different programming styles or were written by different contributors.

Take a *pair* ℓ_1, ℓ_2 of languages; $T_{1,2}$ is the set $T(\ell_1) \cap T(\ell_2)$ of tasks that have implementations in both languages. Thus, we can directly compare the performance of ℓ_1 and ℓ_2 on any task $t \in T_{1,2}$ by juxtaposing each running time in $R(\ell_1, t)$ to each in $R(\ell_2, t)$. To this end, define $R_{1,2}(t)$ as the Cartesian product $R(\ell_1, t) \times R(\ell_2, t)$: a component (r_1, r_2) of $R_{1,2}(t)$ measures the running time r_1 of some implementation of a task in language ℓ_1 against the running time r_2 of some implementation of the same task in the other language ℓ_2 .

It is useful to summarize each such comparison in terms of the *speedup* of one implementation in one language relative to one in the other language. To this end, $S_{1,2}(t)$ is a vector of the same length as $R_{1,2}(t)$, such that each component (r_1, r_2) in the latter determines component $\iota(r_1, r_2)$ in the former, where

$$\iota(a, b) = \text{sgn}(a - b) \cdot \left(1 - \frac{\min(a, b)}{\max(a, b)} \right). \quad (10)$$

The signed ratio $\iota(r_1, r_2)$ varies over the interval $(-1, +1)$: it is zero iff $r_1 = r_2$ (the two implementations ran in the same time exactly), it is negative iff $r_1 < r_2$ (the implementation in language ℓ_1 was faster than the implementation in language ℓ_2), and it is positive iff $r_1 > r_2$ (the implementation in language ℓ_2 was faster than the implementation in language ℓ_1). The absolute value of the signed ratio is proportional to how much faster the faster language is relative to the slower language; precisely, $1/(1 - |\iota(r_1, r_2)|)$ is the faster-to-slower speedup ratio. Nanz and Furia [71] directly compared such speedup ratios, but here we prefer to use the *inverse speedups* $\iota(r_1, r_2)$ because they encode the same information but using a smooth function that ranges over a bounded interval, which is easier to model and to compare to statistics, such as effect sizes, that are also normalized.

Among the various implementations of a chosen task in each language, it makes sense to focus on the *best* one relative to the measured features; that is, to consider the *fastest* variant. We use identifiers with a hat to denote the fastest measures, corresponding to the smallest running times: $\hat{R}(\ell, t)$ is thus $\min R(\ell, t)$, and $\hat{R}_{1,2}(t)$ and $\hat{S}_{1,2}(t)$ are defined like their unhatted counterparts but refer to $\hat{R}(\ell_1, t)$ and $\hat{R}(\ell_2, t)$. In the following, we refer to the hatted data as the *optimal data*.

Analysis goals. The main goal of Nanz and Furia [71]’s analysis is to determine which languages are faster and which are slower based on the empirical performance data described above. The performance comparison is pairwise: for each language pair ℓ_1, ℓ_2 , we want to determine:

- 1) whether the performance difference between ℓ_1 and ℓ_2 is significant;
- 2) if it is significant, how much ℓ_1 is faster than ℓ_2 .

4.2 Frequentist Analysis

Sect. 4.2.1 follows closely Nanz and Furia [71]’s analysis using null hypothesis testing and effect sizes. Then, Sect. 4.2.2 refines the frequentist analysis by taking into account the multiple comparisons problem [66].

4.2.1 Pairwise Comparisons

Null hypothesis. As customary in statistical hypothesis testing, Nanz and Furia [71] use null hypotheses to express whether differences are significant or not. For each language pair ℓ_1, ℓ_2 , null hypothesis $H_0^{1,2}$ denotes lack of significant difference between the two languages:

$H_0^{1,2}$: there is no difference in the performance of languages ℓ_1 and ℓ_2 when used for the same computational tasks

Nanz and Furia [71] use the Wilcoxon signed-rank test to compute the p -value expressing the probability of observing performance data ($\widehat{R}(\ell_1, t)$ against $\widehat{R}(\ell_2, t)$ for every $t \in T_{1,2}$) at least as extreme as the one observed, assuming the null hypothesis $H_0^{1,2}$. The Wilcoxon signed-rank test was chosen because it is a paired, unstandardized test (it does not require normality of the data). Thus, given a probability α , if $p < \alpha$ we reject the null hypothesis and say that the difference between ℓ_1 and ℓ_2 is *statistically significant at level α* .

Effect size. Whenever the analysis based on p -values indicates that the difference between ℓ_1 and ℓ_2 is significant, we still do not know whether ℓ_1 or ℓ_2 is the faster language of the two, nor how much faster it is. To address this question, we compute Cliff’s δ effect size, which is a measure of how often the measures of one language are smaller than the measures of the other language;¹⁸ precisely, $\delta < 0$ means that ℓ_1 is faster than ℓ_2 on average—the closer δ is to -1 the higher ℓ_1 ’s relative speed; conversely, $\delta > 0$ means that ℓ_2 is faster than ℓ_1 on average.

Frequentist analysis results. Tab. 4 shows the results of Nanz and Furia [71]’s frequentist analysis using colors to mark p -values that are **statistically significant at level 0.01** ($p < 0.01$) and **statistically significant at level 0.05** ($0.01 \leq p < 0.05$); for significant comparisons, the effects δ are colored according to which language (the one in the **column** header, or in the **row** header) is faster, that is whether δ is negative or positive. Tab. 4 also reports summary statistics of the same data: the median m and mean μ inverse speedup $\widehat{S}_{1,2}(t)$ across all tasks $t \in T_{1,2}$.

Take for example the comparison between **F#** and **Go** (column #5, row #3). The p -value is colored in dark blue because $0.01 < p = 0.025 < 0.05$, and hence the performance difference between the two languages is statistically significant at level 0.05. The effect size δ is colored in green because it is positive $\delta = 0.408 > 0$ meaning that the language on the **row header** (**Go**, also in green) was faster on average.

18. Nanz and Furia [71]’s analysis used Cohen’s d as effect size; here we switch to Cliff’s δ both because it is more appropriate for nonparametric data and because it varies between -1 and $+1$ like the *inverse speedup ratio* t we use in the reanalysis.

The *language relationship graph* in Fig. 4 summarizes all pairwise comparisons: nodes are languages; an arrow from ℓ_1 to ℓ_2 denotes that the difference between the two languages is significant ($p < 0.05$), and goes from the slower to the faster language according to the sign of δ ; dotted arrows denote significant differences at level 0.05 but not at level 0.01; the thickness of the arrows is proportional to the absolute value of the effect δ . To make the graph less cluttered, after checking that the induced relation is transitive, we remove the arrows that are subsumed by transitivity.

4.2.2 Multiple Comparisons

Conspicuously absent from Nanz and Furia [71]’s analysis is how to deal with the *multiple comparisons* problem [66] (also known as *multiple tests* problem). As the name suggests, the problem occurs when several statistical tests are applied to the same dataset. If each test rejects the null hypothesis whenever $p < \alpha$, the probability of committing a type 1 error (see Sect. 2.2.2) in *at least one* test can grow as high as $1 - (1 - \alpha)^N$ (where N is the number of independent tests), which grows to 1 as N increases.

In a frequentist settings there are two main approaches to address the multiple comparisons problem [5], [66]:

- adjust p -values to compensate for the increased probability of error; or
- use a so-called “omnibus” test, which can deal with multiple comparisons at the same time.

In this section we explore both alternatives on the Rosetta Code data.

P-value adjustment. The intuition behind p -value adjustment (also: correction) is to reduce the p -value threshold to compensate for the increased probability of “spurious” null hypothesis rejections which are simply a result of multiple tests. For example, the Bonferroni adjustment works as follows: if we are performing N tests on the same data and aim for a α significance level, we will reject the null hypothesis in a test yielding p -value p only if $p \cdot N < \alpha$ —thus effectively tightening the significance level.

There is no uniform view about when and how to apply p -value adjustments [1], [5], [35], [70], [72], [78]; and there are different procedures for performing such adjustments, which differ in how strictly they correct the p -values (with Bonferroni’s being the strictest adjustment as it assumes a “worst-case” scenario). To provide as broad a view as possible, we applied three widely used p -value corrections, in decreasing strictness: Bonferroni’s [1], Holm’s [1], and Benjamini-Hochberg’s [12].

Tab. 5 shows the results compared to the unadjusted p -values; and Fig. 4 shows how the language relationship graph changes with different significance levels. Even the mildest correction reduces the number of significant language performance differences that we can claim, and strict corrections drastically reduce it: the 17 language pairs such that $p < 0.05$ without adjustment become 15 pairs using Benjamini-Hochberg’s, and only 7 pairs using Holm’s or Bonferroni’s.

Multiple comparison testing. Frequentist statistics also offers tests to simultaneously compare samples from several groups in a way that gets around the multiple comparisons

LANGUAGE	MEASURE	C	C#	F#	Go	Haskell	Java	Python
C#	p	0.000						
	δ	-0.440						
	m	-0.793						
	μ	-0.600						
F#	p	0.013	0.182					
	δ	-0.603	-0.223					
	m	-0.875	-0.429					
	μ	-0.693	-0.370					
Go	p	0.000	0.006	0.030				
	δ	-0.298	0.270	0.408				
	m	-0.387	0.717	0.727				
	μ	-0.474	0.384	0.534				
Haskell	p	0.000	0.142	0.929	0.025			
	δ	-0.591	-0.186	0.124	-0.356			
	m	-0.964	-0.655	-0.401	-0.909			
	μ	-0.688	-0.315	0.003	-0.440			
Java	p	0.000	0.140	0.008	0.451	0.064		
	δ	-0.446	0.062	0.434	-0.147	0.234		
	m	-0.869	0.466	0.605	-0.337	0.557		
	μ	-0.612	0.188	0.519	-0.194	0.220		
Python	p	0.000	0.042	0.875	0.006	0.334	0.007	
	δ	-0.604	-0.265	-0.021	-0.423	0.048	-0.314	
	m	-0.960	-0.733	-0.016	-0.859	-0.100	-0.714	
	μ	-0.730	-0.378	-0.018	-0.446	0.003	-0.319	
Ruby	p	0.000	0.016	0.695	0.003	0.233	0.001	0.090
	δ	-0.609	-0.360	-0.042	-0.532	-0.059	-0.396	-0.100
	m	-0.973	-0.855	0.076	-0.961	-0.343	-0.831	-0.156
	μ	-0.659	-0.545	-0.030	-0.543	-0.115	-0.468	-0.173

TABLE 4: Frequentist reanalysis of Nanz and Furia [71]’s Rosetta Code data about the performance of 8 programming languages. For each comparison of language ℓ_1 (column header) with language ℓ_2 (row header), the table reports: the p -value of a Wilcoxon signed-rank test comparing the running times of programs in ℓ_1 and in ℓ_2 (colored differently if $p < 0.01$ or $0.01 \leq p < 0.05$); Cliff’s δ effect size, which is negative if ℓ_1 tends to be faster, and positive if ℓ_2 tends to be faster; and the median m and mean μ of the inverse speedup of ℓ_1 vs. ℓ_2 (again, negative values indicate that ℓ_1 is faster on average, and positive values that ℓ_2 is).

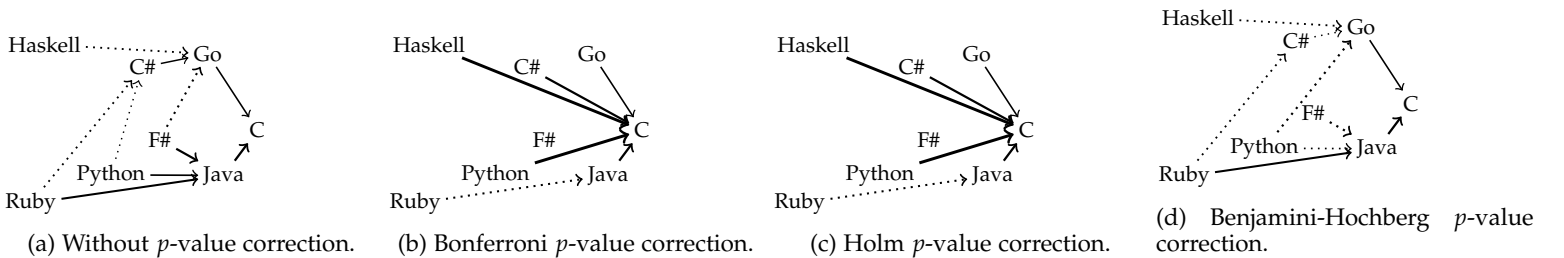


Fig. 4: Language relationship graphs summarizing the frequentist reanalysis of Nanz and Furia [71], with different p -value correction methods. In every graph, an arrow from node ℓ_1 to node ℓ_2 indicates that the p -value comparing performance data between the two languages is $p < 0.01$ (solid line) or $0.01 \leq p < 0.05$ (dotted line), and Cliff’s δ effect size indicates that ℓ_2 tends to be faster; the thickness of the arrows is proportional to the absolute value of Cliff’s δ .

problem. The Kruskal-Wallis test [48] is one such test that is applicable to nonparametric data, since it is an extension of the Mann-Whitney test to more than two groups, or a nonparametric version of ANOVA on ranks.

To apply Kruskal-Wallis, the samples from the various groups must be directly comparable. The frequentist analysis done so far has focused on *pairwise comparisons* resulting in inverse speedup ratios; in contrast, we have to consider absolute running times to be able to run the Kruskal-Wallis test. Unfortunately, this means that only the Rosetta Code tasks for which *all 8 languages* have an implementation can be included in the analysis (using the notation of Sect. 3.1: $\cap_{\ell} T(\ell)$); otherwise, we would be including running times

on unrelated tasks, which are not comparable in a meaningful way. This restriction leaves us with a mere 6 data points per language—for a total of $24 = 6 \cdot 8$ running times.

On this dataset, the Kruskal-Wallis test indicates that there are statistically significant differences ($p = 0.002 < 0.05$); but the test cannot specify *which* language pairs differ significantly. To find these out, it is customary to run a post-hoc analysis. We tried several that are appropriate for this setting:

MC: a multiple-comparison between treatments [85, Sec. 8.3.3]

Nemenyi: Nemenyi’s test of multiple comparisons [101]

Dunn: Dunn’s test of multiple comparisons [101] with p -

LANGUAGE	CORRECTION	C	C#	F#	Go	Haskell	Java	Python
C#	-	0.000						
	B	0.007						
	H	0.006						
	B-H	0.001						
F#	-	0.013	0.182					
	B	0.358	1.000					
	H	0.205	1.000					
	B-H	0.028	0.232					
Go	-	0.000	0.006	0.030				
	B	0.000	0.169	0.849				
	H	0.000	0.121	0.394				
	B-H	0.000	0.018	0.053				
Haskell	-	0.000	0.142	0.929	0.025			
	B	0.001	1.000	1.000	0.706			
	H	0.001	1.000	1.000	0.353			
	B-H	0.000	0.189	0.929	0.047			
Java	-	0.000	0.140	0.008	0.451	0.064		
	B	0.002	1.000	0.214	1.000	1.000		
	H	0.002	1.000	0.130	1.000	0.709		
	B-H	0.000	0.189	0.018	0.505	0.100		
Python	-	0.000	0.042	0.875	0.006	0.334	0.007	
	B	0.000	1.000	1.000	0.180	1.000	0.198	
	H	0.000	0.502	1.000	0.122	1.000	0.128	
	B-H	0.000	0.069	0.908	0.018	0.389	0.018	
Ruby	-	0.000	0.016	0.695	0.003	0.233	0.001	0.090
	B	0.001	0.441	1.000	0.083	1.000	0.025	1.000
	H	0.001	0.236	1.000	0.062	1.000	0.020	0.900
	B-H	0.000	0.031	0.748	0.010	0.283	0.004	0.133

TABLE 5: Frequentist reanalysis of Nanz and Furia [71]’s Rosetta Code data about the performance of 8 programming languages. For each comparison of language ℓ_1 (column header) with language ℓ_2 (row header), the table reports the p -value of a Wilcoxon signed-rank test—comparing the running times of programs in ℓ_1 and in ℓ_2 —corrected with different methods: no correction (–, identical to the p -values in Tab. 4), Bonferroni correction (B), Holm correction (H), Benjamini-Hochberg correction (B-H). Values are colored if $p < 0.01$ or $0.01 \leq p < 0.05$.

values adjusted using Benjamini-Hochberg

PW: a pairwise version of the Wilcoxon test with p -values adjusted using Benjamini-Hochberg [48]

The number of significantly different language pairs are as follows:

POST HOC	95% level	99% level
MC	2	0
Nemenyi	2	0
Dunn	3	0
PW	5	0

A closer look into the results shows that all the significant comparisons involve C and another language, with the exception of the comparisons between Go and Python, and Go and Ruby, which the pairwise Wilcoxon classifies as significant at the 95% level; no comparison is significant at the 99% level.

4.2.3 Summary and Criticism of Frequentist Statistics

Analyzing significance using hypothesis testing is problematic for a number of general reasons. First, a null hypothesis is an oversimplification since it forces a binary choice between the extreme case of *no effects*—which is unlikely to happen in practice—versus *some effects*. Second, p -values are probabilities conditional on the null hypothesis being true; using them to infer whether the null hypothesis itself is likely true is unsound in general.

Even if we brush these general methodological problems aside, null-hypothesis testing becomes even murkier when dealing with multiple tests. If we choose to ignore the multiple tests problem, we may end up with incorrect results that overestimate the weight of evidence; if we decide to add corrections it is debatable which one is the most appropriate, with a conflicting choice between “strict” corrections—safer from a methodological viewpoint, but also likely to strike down any significant findings—and “weak” corrections—salvaging significance to a larger degree, but possibly arbitrarily chosen.

Using effect sizes is a definite improvement over p -values: effect sizes are continuous measures rather than binary alternatives; and are unconditional summary statistics of the experimental data. Nonetheless, they remain elementary statistics that provide limited information. In our case study, especially in the cases where the results of a language comparison are somewhat borderline or inconclusive, it is unclear whether an effect size provides a more reliable measure than statistics such as mean or median (even when one takes into consideration common language effect size statistics [4], normally referred to as Vargha-Delaney effect size [95]). Besides, merely summarizing the experimental data is a very primitive way of modeling it, as it cannot easily accommodate different assumptions to study how they would change our quantitative interpretation of the data. Finally, frequentist best practices normally recommend to use effect size *in addition to* null-hypothesis testing;

which brings additional problems whenever a test does not reach significance, but the corresponding effect size is non-negligible.

As some examples that expose these general weaknesses, consider the four language comparisons C# vs. F#, F# vs. Python, F# vs. Ruby, and Haskell vs. Python. According to the p -values the four comparisons are inconclusive. The effect sizes tell a different story: they suggest a sizable speed advantage of C# over F# ($\delta = -0.223$), and a small to negligible difference in favor of F# over Python ($\delta = -0.021$), of F# over Ruby ($\delta = -0.042$), and of Python over Haskell ($\delta = 0.048$). Other summary statistics make the picture even more confusing: while median and mean agree with the effect size in the F# vs. C# and F# vs. Python comparisons, the median suggests that Ruby is slightly faster than F# ($m = 0.076$), and that Haskell is faster than Python ($m = -0.1$); whereas mean and effect size suggest that F# ($\mu = -0.03$) and Python ($\mu = 0.003$) are slightly faster. We might argue about the merits of one statistic over the other, but there is no clear-cut way to fine-tune the analysis without arbitrarily complicating it. In contrast, with Bayesian statistics we can obtain nuanced, yet natural-to-interpret, analyses of the very same experimental data—as we demonstrate next.

4.3 Bayesian Reanalysis

We analyze Nanz and Furia [71]’s performance data using an explicit Bayesian model. For each language pair ℓ_1, ℓ_2 , the model expresses the probability that one language achieves a certain speedup s over the other language. Speedup s is measured by the inverse speedup ratio (10) ranging over the interval $(-1, +1)$.

4.3.1 Statistical Model

To apply Bayes’ theorem, we need a likelihood—linking observed data to inferred data—and a prior—expressing our initial assumption about the speedup.

Likelihood. The likelihood $\mathcal{L}[d, s]$ weighs how likely we are to observe an inverse speedup d assuming that the “real” inverse speedup is s . Using the terminology of Sect. 2, d is the *data* and s is the *hypothesis* of probabilistic inference. In this case, the likelihood essentially models possible imprecisions in the experimental data; thus, we express it in terms of the difference $d - s$ between observed and the “real” speedup:

$$d - s \sim \mathcal{N}(0, \sigma),$$

that is the difference is normally distributed with mean zero (no systematic error) and unknown standard deviation. This just models experimental noise that might affect the measured speedups (and corresponds to all sorts of approximations and effects that we cannot control for) so that they do not coincide with the “real” speedups.

Priors. The prior $\pi[s]$ models initial assumptions on what a plausible speedup for the pair of languages might be. We consider three different priors, from uniform to biased:

- 1) a *uniform* prior $\mathcal{U}(-1, 1)$ —a uniform distribution over $(-1, 1)$;
- 2) a weak unbiased *centered normal* prior $\mathcal{N}(0, \sigma_1)$ —a normal distribution with mean 0 and standard deviation σ_1 , truncated to have support $(-1, 1)$;

- 3) a biased *shifted normal* prior $\mathcal{N}(\mu_2, \sigma_2)$ —a normal distribution with mean μ_2 and standard deviation σ_2 , truncated to have support $(-1, 1)$;

Picking different priors corresponds to assessing the impact of different assumptions on the data under analysis. Unlike the uniform prior, the normal priors have to be based on some data other than Nanz and Furia [71]’s. To this end we use the Computer Language Benchmarks Game [91] (for brevity, *Bench*). The data from Bench are comparable to those from Nanz and Furia [71] as they also consist of curated selections of collectively written solutions to well-defined programming tasks running on the same input and refined over a significant stretch of time; however, Bench was developed independently of Rosetta Code, which makes it a complementary source of data.

In the centered normal prior, we set σ_1 to the largest absolute inverse speedup value observed in Bench between ℓ_1 and ℓ_2 . This represents the assumption that the largest speedups observed in Bench give a range of plausible maximum values for speedups observed in Rosetta Code. This is still a very weak prior since maxima will mostly be much larger than typical values; and there is no bias because the distribution mean is zero.

In the shifted normal prior, we set μ_2 and σ_2 to the mean and standard deviation of Bench’s optimal data. The resulting distribution is biased towards Bench’s data, as it uses Bench’s performance data to shape the prior assumptions; if $\mu_2 \gg 0$ or $\mu_2 \ll 0$, the priors are biased in favor of one of the two languages being faster. In most cases, however, the biased priors still allow for a wide range of speedups; and, in any case, they can still be swamped by strong evidence in Nanz and Furia [71]’s data.

Posteriors and uncertainty intervals. For each pair of languages ℓ_1 and ℓ_2 , for each prior, applying Bayes’ theorem gives a posterior distribution $\mathcal{P}_{1,2}(s)$, which assigns a probability to any possible value s of inverse speedup. We summarize the posterior using descriptive statistics:

- 1) 95% and 99% uncertainty intervals;
- 2) median m and mean μ of the posterior.

Since we are trying to establish which language, ℓ_1 or ℓ_2 , is faster according to our experiments, it is useful to check whether the uncertainty intervals include the origin or not. For 95% probability, consider the leftmost uncertainty interval $(c_{0\%}, c_{95\%})$ and the rightmost uncertainty interval $(c_{5\%}, c_{100\%})$; that is, c_x is the value such that the posterior probability of drawing an inverse speedup less than c_x is x : $\int_{-1}^{c_x} \mathcal{P}_{1,2}(s) ds < x$. If $c_{95\%} < 0$ there is a 95% chance that ℓ_1 is at least $c_{95\%}$ faster than ℓ_2 according to the posterior; conversely, if $0 < c_{5\%}$ there is a 95% chance that ℓ_2 is at least $c_{5\%}$ faster;¹⁹ if neither is the case, the comparison between the two languages is inconclusive at a 95% certainty level. A similar reading applies to the 99% uncertainty intervals.

4.3.2 Bayesian Reanalysis Results

Tab. 6 shows the results of the Bayesian analysis of Nanz and Furia [71] data, for each of three priors. If two languages differ at a 95% certainty level, the table reports the value $c_{95\%}$ if the language in the [column](#) header is faster, and the

19. Since $c_{5\%} \leq c_{95\%}$, the two conditions are mutually exclusive.

LANGUAGE MEASURE	C			C#			F#			Go			Haskell			Java			Python					
	\mathcal{U}	\mathcal{N}	\mathcal{S}	\mathcal{U}	\mathcal{N}	\mathcal{S}	\mathcal{U}	\mathcal{N}	\mathcal{S}	\mathcal{U}	\mathcal{N}	\mathcal{S}	\mathcal{U}	\mathcal{N}	\mathcal{S}	\mathcal{U}	\mathcal{N}	\mathcal{S}	\mathcal{U}	\mathcal{N}	\mathcal{S}			
C#	95%	-0.398	-0.392	-0.631																				
	99%	-0.307	-0.306	-0.610																				
	m	-0.601	-0.594	-0.678																				
	μ	-0.598	-0.592	-0.677																				
F#	95%	-0.355	-0.347	-0.411	-0.082	-0.073	-0.080																	
	99%	-0.178	-0.172	-0.301	+0.000	+0.000	+0.000																	
	m	-0.682	-0.661	-0.651	-0.367	-0.360	-0.311																	
	μ	-0.667	-0.649	-0.647	-0.367	-0.358	-0.308																	
Go	95%	-0.358	-0.355	-0.363	0.099	0.091	0.090	0.187	0.169	0.189														
	99%	-0.304	-0.300	-0.316	+0.001	+0.000	+0.009	+0.053	+0.035	+0.090														
	m	-0.473	-0.473	-0.470	+0.386	+0.375	+0.321	+0.535	+0.513	+0.458														
	μ	-0.473	-0.471	-0.470	0.387	0.376	0.321	0.533	0.513	0.456														
Haskell	95%	-0.512	-0.502	-0.494	-0.011	0.000	0.000	0.000	0.000	0.000	-0.211	-0.201	-0.057											
	99%	-0.419	-0.423	-0.419	+0.000	+0.000	+0.000	+0.000	+0.000	+0.000	-0.105	-0.099	+0.000											
	m	-0.687	-0.681	-0.653	-0.318	-0.306	+0.055	+0.009	+0.001	+0.229	-0.443	-0.429	-0.243											
	μ	-0.686	-0.680	-0.653	-0.317	-0.305	0.057	0.007	0.001	0.233	-0.441	-0.427	-0.242											
Java	95%	-0.461	-0.457	-0.575	0.000	0.000	0.000	0.280	0.278	0.240	0.000	0.000	-0.011	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	
	99%	-0.397	-0.389	-0.552	+0.000	+0.000	+0.000	+0.186	+0.182	+0.173	+0.000	+0.000	+0.000	+0.000	+0.000	+0.000	+0.000	+0.000	+0.000	+0.000	+0.000	+0.000	+0.000	
	m	-0.610	-0.606	-0.632	+0.187	+0.177	+0.121	+0.520	+0.511	+0.447	-0.193	-0.190	-0.187	+0.217	+0.216	+0.123								
	μ	-0.610	-0.606	-0.632	0.189	0.177	0.120	0.520	0.511	0.443	-0.193	-0.190	-0.186	0.219	0.216	0.120								
Python	95%	-0.581	-0.569	-0.791	-0.102	-0.103	-0.164	0.000	0.000	0.000	0.000	-0.232	-0.224	-0.308	0.000	0.000	-0.083	-0.080	-0.161					
	99%	-0.515	-0.502	-0.765	+0.000	+0.000	+0.062	+0.000	+0.000	+0.000	-0.138	-0.125	-0.224	+0.000	+0.000	+0.000	+0.000	+0.000	-0.077					
	m	-0.728	-0.724	-0.851	-0.375	-0.371	-0.404	-0.021	-0.017	-0.150	-0.445	-0.436	-0.495	+0.003	+0.004	-0.117	-0.319	-0.316	-0.369					
	μ	-0.728	-0.723	-0.851	-0.376	-0.371	-0.405	-0.023	-0.016	-0.158	-0.445	-0.437	-0.497	0.004	0.004	-0.119	-0.319	-0.315	-0.372					
Ruby	95%	-0.452	-0.443	-0.950	-0.288	-0.274	-0.548	0.000	0.000	-0.568	-0.321	-0.312	-0.631	0.000	0.000	-0.858	-0.255	-0.244	-0.807	0.000	0.000	-0.008	-0.008	
	99%	-0.366	-0.347	-0.947	-0.163	-0.150	-0.488	+0.000	+0.000	-0.505	-0.222	-0.211	-0.575	+0.000	+0.000	-0.843	-0.161	-0.154	-0.786	+0.000	+0.000	+0.000	+0.000	
	m	-0.656	-0.646	-0.957	-0.545	-0.529	-0.699	-0.026	-0.026	-0.723	-0.542	-0.538	-0.749	-0.114	-0.115	-0.897	-0.470	-0.460	-0.854	-0.173	-0.172	-0.175	-0.175	
	μ	-0.655	-0.644	-0.957	-0.543	-0.527	-0.698	-0.029	-0.024	-0.723	-0.542	-0.536	-0.749	-0.113	-0.114	-0.897	-0.469	-0.460	-0.854	-0.172	-0.171	-0.175	-0.175	

TABLE 6: Bayesian reanalysis of Nanz and Furia [71]’s Rosetta Code data about the performance of 8 programming languages. For each comparison of language ℓ_1 (column header) with language ℓ_2 (row header), for each choice of prior distribution among uniform \mathcal{U} , centered normal \mathcal{N} , and shifted normal \mathcal{S} , the table reports the endpoint of the 95% and 99% uncertainty intervals of the posterior inverse speedup of ℓ_1 vs. ℓ_2 that is closest to the origin, if such interval does *not* include the origin; in this case, the endpoint’s absolute value is a lower bound on the inverse speedup of ℓ_1 vs. ℓ_2 , and indicates that ℓ_1 tends to be faster if it is **negative** and that ℓ_2 tends to be faster if it is **positive**. If the uncertainty interval includes the origin, the table reports a value of 0.0, which indicates the performance comparison is inconclusive. The table also reports median m and mean μ of the *posterior* inverse speedup of ℓ_1 vs. ℓ_2 (again, negative values indicate that ℓ_1 is faster on average, and positive values that ℓ_2 is).

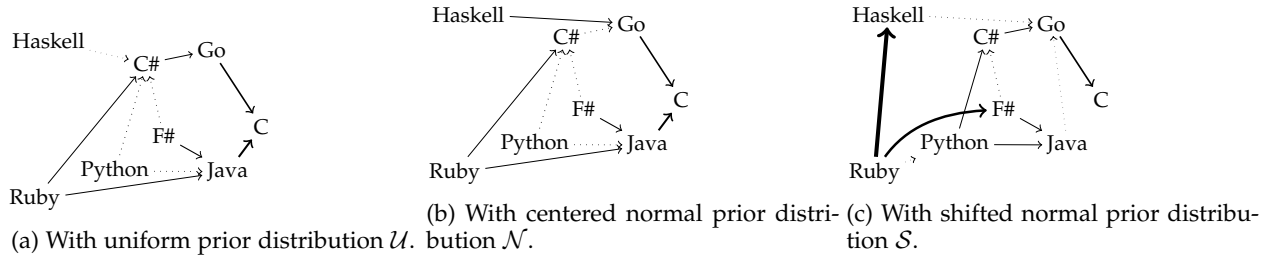


Fig. 5: Language relationship graphs summarizing the Bayesian reanalysis of Nanz and Furia [71], with different prior distributions of speedups. In every graph, an arrow from node ℓ_1 to node ℓ_2 indicates that, according to the posterior probability distribution of inverse speedups, ℓ_1 is faster than ℓ_2 with 99% probability (solid line) or 95% probability (dotted line); the thickness of the arrows is proportional to the absolute value of the minimum inverse speedup.

ANALYSIS	\mathcal{U}								\mathcal{N}		\mathcal{S}	
	-		B		H		BH		-		-	
	95%	99%	95%	99%	95%	99%	95%	99%	95%	99%	95%	99%
frequentist	17	12	7	6	7	6	15	7				
Bayesian	19	15							18	14	22	18

TABLE 7: The number of language comparisons, out of a total 28 language pairs, where each ANALYSIS found a significant difference at the 95% and 99% levels. The two rows correspond to *frequentist* and *Bayesian* analyses; the frequentist analyses differ in what kind of p -value correction they employ (– for none, B for Bonferroni, H for Holm, and BH for Benjamini-Hochberg); the Bayesian analyses differ in what prior they employ (\mathcal{U} for uniform, \mathcal{N} for centered normal, and \mathcal{S} for shifted normal distributions).

value $c_{5\%}$ if the language in the row header is faster; if the comparison is inconclusive, the table reports the value 0.0. The table displays the results of the comparison at a 99% certainty level in a similar way, using the values $c_{99\%}$, $c_{1\%}$, or 0.0.

The *language relationship graphs* in Fig. 5 summarize all comparisons similarly to the frequentist analysis: nodes are languages; an arrow from ℓ_1 to ℓ_2 denotes that the difference between the two languages is significant (at 95% certainty level), and goes from the slower to the faster language; solid arrows denote differences that are significant at 99% certainty level (otherwise arrows are dotted); the thickness of the arrows is proportional to the absolute value of the endpoint of the 95% or 99% certainty level closer to the origin (as explained above).

Prediction errors and multiple tests. Multiple testing is generally not an issue in Bayesian analysis [39]. In fact, there is no null-hypothesis testing in Bayesian analysis: we are not performing multiple tests, each with a binary outcome, but are computing posterior distribution probabilities based on data and priors. We can estimate mean and standard deviation of the posterior, and get a clear, quantitative picture of how likely there is an effect and of what size—instead of having to guess a correction for how the errors in binary tests may compound. As we discussed in Sect. 2.2, it is also advisable to focus on type S and M errors, which are directly meaningful in terms of practical significance.

While it is possible to reason in terms of type S and M errors also using frequentist techniques such as effect

sizes, Bayesian analysis provide a natural framework to precisely and soundly assess these probabilities of errors. In our Bayesian reanalysis of Nanz and Furia [71]’s data, we can readily assess the probability of type S and M errors:

- if a 95% uncertainty interval does not include the origin, there is at most a 5% chance of committing a type S error in that comparison;
- the width W of a 95% uncertainty interval indicates that there is at most a 5% chance of committing a type M error larger than W .

Significance analysis. There are no major inconsistencies between the overall pictures drawn by the frequentist (with effect sizes) and by the Bayesian analysis. Bayesian statistics, however, produced results that are more natural to interpret in terms of practically significant measures, and lend richer information that supports a detailed analysis even when the frequentist analysis was utterly inconclusive.

In quantitative terms, every “significant” difference according to frequentist analysis was confirmed by Bayesian analysis. Conversely, the Bayesian analysis provides conclusive comparisons (that is, one language being faster than another) for more language pairs than the frequentist analysis (see Tab. 7): while the latter finds significant differences in 17 language pairs (out of a total of 28), Bayesian analysis finds 18 with centered normal priors, 19 with uniform priors, and 22 with shifted normal priors.

We consider this a distinct advantage from a software engineering perspective: we could draw more conclusions from the same data. The fact that priors of different strengths lead to increasingly precise outcomes also sharpens the overall conclusions: the Rosetta dataset and the Bench dataset (used to model the stronger prior) seem to often corroborate each other. More generally, such techniques support analyses where each new scientific study increases the clarity and robustness of the overall understanding of what is really going on. Even in situations where a new study does not corroborate but contradicts a previous one, we would enrich our understanding of the issues by making it more nuanced in a quantitative way.

4.3.3 Prior Sensitivity Analysis

Differences due to using different priors point to what kinds of assumptions are necessary to narrow down the analysis in certain cases. To disentangle the cases where weak data leads to inconclusive or inconsistent results, we can always visually and numerically inspect the *posterior distributions*

provided by Bayesian analysis. This is a kind of *sensitivity analysis* of the priors, whose goal is reporting the effects of choosing different priors, in order to better understand the strength of the data under analysis and the ramifications of assumptions that we may introduce.

Rather than going through an exhaustive sensitivity analysis, let us single out four interesting cases. These are the four language comparisons that we mentioned at the end of Sect. 4.3.4: C# vs. F#, F# vs. Python, F# vs. Ruby, and Haskell vs. Python. These comparisons were inconclusive in the frequentist analysis, sometimes with contradictory evidence given by different summary statistics. In the Bayesian analysis, we get nuanced yet clearer results; and we can go beyond a binary choice made on the basis of summary statistics, looking at the posterior distributions plot for these three cases (shown in Fig. 6):

C# vs. F#: regardless of the prior, the posterior estimated (inverse) speedup is negative with 95% probability. This confirms that C# tends to be faster in the experiments, largely independent of prior assumptions: the data *swamps* the priors.

F# vs. Python: the analysis is genuinely inconclusive if we use an unbiased prior (uniform or centered normal), as both posteriors are basically centered around the origin. If we use a biased prior (the shifted normal prior), instead, the posterior tends to negative values. This may indicate some evidence that F# is faster, but only if we are willing to use Bench’s data as the starting point to process Rosetta Code’s data (which, alone, remains too weak to suggest a significant effect).

Precisely, we can compute on the posterior draws—built using the shifted prior—that the probability that the speedup is negative is approximately 73%; and the probability that it is less than -0.09 (corresponding to a 1.1-or-greater speedup of F# over Python) is about 60%. The shifted normal prior is $\mathcal{N}(-0.55, 0.69)$ in this case, which corresponds to a marked bias in favor of F# being substantially faster. Nonetheless, we can find both tasks where F# is faster and tasks where Python is faster both in Rosetta Code and in Bench, indicating no clear winner with the current data even if F# has some advantage on average.

F# vs. Ruby: here the posterior changes drastically according to the prior. With uniform and centered normal priors, we cannot claim any noticeable speedup of one language over the other. But with the shifted normal prior, the posterior lies decidedly to the left of the origin and is quite narrow, indicating a high probability that F# is much faster than Ruby; precisely, the speedup is less than -0.4 with near certainty.

In this comparison, the shifted prior is $\mathcal{N}(-0.79, 0.31)$ —strongly biased in favor of F#. The dependence on such a strong prior means that Rosetta’s data about this language comparison is not conclusive by its own: both datasets may indicate that F# tends to be faster, but Bench’s data is overwhelming whereas Rosetta’s is very mixed. Regardless of whether we are willing to take Bench as a starting point of our analysis, we understand that Rosetta’s data about this language comparison is inadequate and should be extended and sharpened

with new experiments or with a deeper analysis of Rosetta’s experimental conditions that determined the outlier results.

Haskell vs. Python: the language comparison is inconclusive if we use an unbiased prior (uniform or centered normal). In contrast, starting from a shifted normal prior, Haskell emerges as the faster language with 82% probability.

In this case the biased prior is $\mathcal{N}(-0.67, 0.55)$, which is clearly biased in favor of Haskell but would not prevent strong data evidence from changing the outcome. If we are willing to see this biased prior as encoding a reasonable assumption on the preferential advantage that a compiled language such as Haskell ought to have over an interpreted language, Rosetta’s data tempers the speed advantage Haskell can claim but does not nullify it. Regardless of whether we find this interpretation reasonable, the analysis clearly shows what the outcome of this language comparison hinges on, and suggests that further experiments involving these two languages are needed to get to a more general conclusion.

Repeating the sensitivity analysis for every language comparison, extending the models, and possibly collecting more data, are beyond the scope of this paper. The goal of this subsection was illustrating concrete examples of how sensitivity analysis can support a sophisticated data analysis while remaining grounded in results that have a direct connection to the quantities of interest that determine practical significance. Key to this flexibility are the very features of Bayesian statistics that we presented throughout the paper.

4.3.4 Summary of Bayesian Statistics

The availability of full posterior distributions makes Bayesian analyses more transparent and based on richer, multi-dimensional information than frequentist analyses. Priors are another tool unique to Bayesian analysis that fully supports “what-if” analyses: by trying out different priors—from completely uniform to informative/biased—we can understand the impact of assumptions on interpreting the experimental data, and naturally connect our analysis to others and, more generally, to the body of knowledge of the domain under empirical investigation.

5 DOING BAYESIAN STATISTICS IN PRACTICE: GUIDELINES AND TOOLS

The main goal of this paper is demonstrating the advantages of Bayesian statistics over frequentist statistics. Once we become convinced of this point, how do we *concretely* adopt Bayesian statistics in the kind of empirical research that is common in software engineering? While we leave a detailed presentation of a framework to future work, here we summarize some basic guidelines organized around a few broadly applicable steps.

Research questions. If we formulate the research questions following the usual approach based on null hypothesis (no effect) vs. alternative hypothesis (some effect), then using statistical hypothesis testing may appear as the only natural analysis approach. To avoid this, we should try

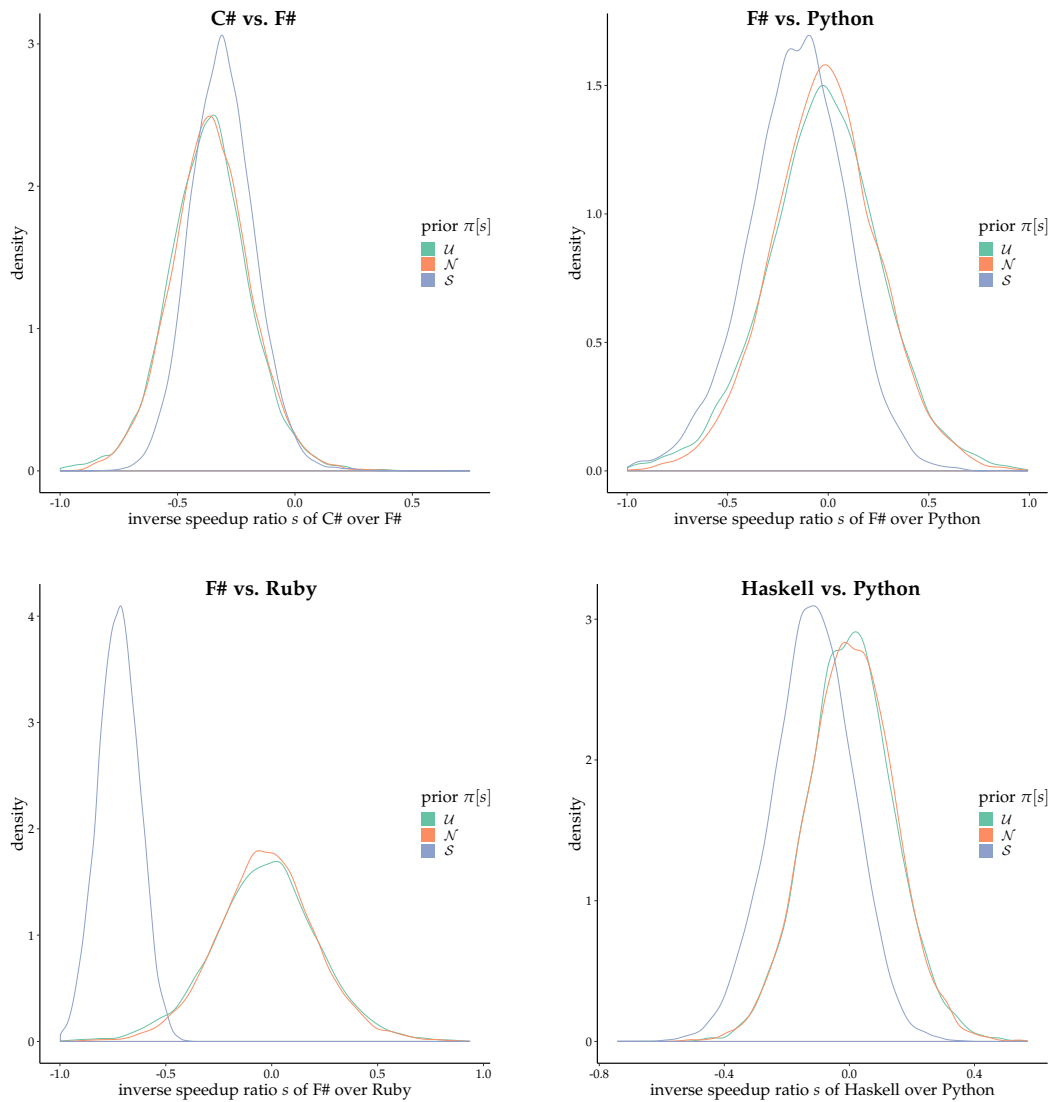


Fig. 6: Posterior distributions of inverse speedup ratios comparing C# to F#, F# to Python, F# to Ruby, and Haskell to Python. Each comparison comprises three posterior distributions built using Bayesian analysis with uniform \mathcal{U} , centered normal \mathcal{N} , and shifted normal \mathcal{S} prior distributions.

to phrase the research questions in a way that: 1) they have a clear connection to variables in the *analysis domain*; 2) they *avoid* dichotomous alternative answers; and 3) they *quantify*—if approximately—the expected ranges of values that constitute a *practically* significant effect.

For example, we could rephrase Nanz and Furia [71]’s generic research question about two languages ℓ_1 and ℓ_2 :

Is ℓ_1 significantly faster than ℓ_2 ?

as:

What is the average speedup of ℓ_1 over ℓ_2 , and what is the probability that it is greater than 1%?²⁰

The new formulation requires that every answer to the question provide an estimate of speedup, and that small (less than 1%) differences are practically irrelevant.

20. In Sect. 4.3, we answered a different question—about the probability that it is greater than zero—so as to allow a direct comparison to the results of the original analysis based on hypothesis testing.

Regressive models. Regressive models—from simple linear regressions to more flexible generalized linear models—are very powerful statistical models which are applicable with remarkable flexibility to a wide variety of problems and data domains [38].²¹ Since regression can also be done by means of frequentist techniques, regressive models are an approachable entry point to Bayesian statistics for researcher used to frequentist techniques. There exist statistical packages (such as `rstanarm` and `brms` for R) that provide functions with an interface very similar to the classic `lm` (which does regression with frequentist techniques), with optional additional inputs (for example, to change the priors) to take full advantage of Bayesian statistics’ greater flexibility. We used such packages in the reanalysis of Ceccato et al. [18] in Sect. 3.

21. In particular, regressive models are not limited to continuous variables but can also model discrete variables such as those on a Likert scale [68].

Even when a statistical model other than standard linear regression is a better choice, it often is useful to start the analysis using simple regressive models, and then tweak them with the additional modeling features that may be required. This incrementality may not be possible—or as easy—in a frequentist setting, where changing statistical model normally requires switching to a new custom algorithm. In contrast, Bayesian statistics are all based on a uniform framework that applies Bayes’ theorem to update the posterior based on the data; the model is all in the choice of priors and likelihoods. Thus, performing software engineering studies based on Bayesian techniques would support a flexible use of different models of varying complexity within a single, coherent statistical framework.

Choosing priors. After following this paper’s reanalysis, it should be clear that “choosing the prior” need not be a complex exercise. In fact, in Bayesian analysis you don’t *choose the prior*, but rather experiment with a variety of priors as a way of evaluating different modeling choices. This process, presented in Sect. 2.2.4, is called *sensitivity analysis*.

As a minimum, priors should indicate reasonable ranges of values, or, conversely, rule out “impossible” ranges of values as highly unlikely. This is what the uniform and centered normal priors did in Sect. 4.3. If we have more precise information about the plausible variable ranges and we want to explore their effect on the model, we can try out more biased priors—such as the shifted normal priors in Sect. 4.3. The goal of sensitivity analysis is not to choose a prior over others—not even after the analysis—but rather to explore the same data from different angles, painting a rich, nuanced picture.

When the goal of sensitivity analysis is to identify the model with the greatest *predictive power*—that is, that strikes a good balance between underfitting and overfitting—there exist quantitative comparison criteria based on measures of information entropy [62]. In most common cases, however, a qualitative sensitivity analysis—such as the one in Sect. 4.3.3—goes a long way in describing the analysis results while remaining grounded in the problem domain.

As more empirical evidence is gathered, and studies are performed, about a certain software engineering area of interest, priors can be chosen based on the evidence that is already known and is relevant to new studies within the same area. This can give empirical software engineering a concrete way in which to build chains of evidence that span multiple studies, and ultimately present the current state-of-the-art knowledge to practitioners in a quantitative form. Doing so could also make research results more actionable, since different scenarios could be evaluated and simulated.

Visualization. At every step during the analysis, visualizing data and intermediate results acts as a sanity check on the data analysis. Of course visualization is not helpful only in a Bayesian setting, but one of the key strengths of Bayesian analysis is that it often better lends itself to visualization since it is based on *distributions of data*, and how they change to better reflect the data. For example, the posterior plots of Fig. 6 support the sensitivity analysis of

three language comparisons that closes Sect. 4.3.²²

Dichotomous decisions. Even if we avoid using the problematic notion of “null hypothesis”, there exist research questions that ultimately require a dichotomous (yes/no) answer. How can we approach such questions without giving up the benefits of a Bayesian analysis?

Since Bayesian analysis’s primary source of information is in the form of posterior distributions, the best way to approach dichotomous decisions is to ground them in an *analysis of the posteriors*. This also helps ensure that the decisions are taken based on *practically significant* measures rather than artificial statistics such as the p -values.

This general guideline is very flexible. Then, there exist more specific protocols that can be applied to formalize the notion of how to answer dichotomous questions based on an analysis of the posteriors. In clinical studies, for instance, researchers have used the notion of *region of practical equivalence* (ROPE) to express “a range of parameter values that are equivalent to the null value for current practical purpose” [58]. ROPEs are typically derived based on historical data and meta analyses in each data domain. Then a decision is taken based on whether the ROPE is included in or excluded from a nominal uncertainty interval (typically, with 95% or 99% probability).

Bayes factors [43], [52], [56] are another Bayesian technique that may be used, under certain conditions, to choose between models in a dichotomous fashion. A Bayes factor is a *ratio of likelihoods* $B_{1,2} = \mathcal{L}[d;h_1]/\mathcal{L}[d;h_2]$ of the same data d under different hypotheses (that is, statistical models) h_1 and h_2 . In a nutshell, the Bayes factor $B_{1,2}$ expresses how much the data would shift the posterior probability in favor of h_1 over h_2 : $B_{1,2} < 1$ indicates that h_1 ’s probability decreases, whereas $B_{1,2} > 1$ indicates that it increases. Therefore, if h_1 ’s and h_2 ’s prior probabilities are about the same, $B_{1,2}$ can be used to decide whether the data supports more strongly h_1 or h_2 . When the prior probabilities are not similar in value, their exact ratio must be known in order to meaningfully interpret a Bayes factor.

We mentioned ROPE analysis and Bayes factor as examples of techniques that may be appropriate in certain specific contexts. However, we stress again that there is no substitute for a careful, bespoke posterior analysis—performed with domain knowledge—if we want to get to sound conclusions grounded in practical significance.

Tools. A number of powerful probabilistic languages, geared towards analyzing Bayesian statistical models, have been developed in the last decade. The best known are JAGS,²³ BUGS,²⁴ Turing,²⁵ and Stan²⁶; Stan is probably the most mature and actively maintained at the time of writing. All of these tools input a statistical model—such as the one in (7)—including information about priors, as well as the analyzed data; they output a posterior distribution computed numerically as we described in Sect. 2.3, which can be visualized, analyzed, and used to build derived models. This process is central to every Bayesian data analysis.

22. Gabry et al. [34] suggest guidelines on visualization in a Bayesian analysis workflow.

23. <http://mcmc-jags.sourceforge.net/>

24. <https://www.mrc-bsu.cam.ac.uk/software/bugs/>

25. <https://github.com/TuringLang>

26. <https://mc-stan.org/>

These tools are normally used as back ends through convenient interfaces in the most common statistical analysis toolsets—such as R,²⁷ Python’s Scipy,²⁸ and Julia.²⁹ The interface functions take care of importing and exporting data with the back ends, and provide functionality that makes building and using common models (such as regressive ones) straightforward. Another interesting tool worth mentioning is JASP³⁰ which supports both frequentist and Bayesian analyses under a single, user-friendly interface. This way, it can provide an accessible way to design Bayesian analyses that replace the use of better-known frequentist techniques.

6 THREATS TO VALIDITY AND LIMITATIONS

Like every empirical study, this paper’s reanalyses (Sect. 3 and Sect. 4) may incur threats to validity. Rather than just discussing them as traditionally done [99], we illustrate how using Bayesian statistics can generally help to mitigate them (Sect. 6.1) and to avoid other common problems in empirical software engineering (Sect. 6.2). This section concludes with a mention of the limitations of Bayesian statistics (Sect. 6.3).

6.1 Threats to Validity

Using Bayesian statistics is unlikely to affect *construct* validity, which has to do with whether we measured what the study was supposed to measure. Our reanalyses reused the same data that was produced in the original studies by Ceccato et al. [18] and by Nanz and Furia [71]; therefore, the same construct validity threats—and threat mitigation strategies—apply to the reanalyses in Sect. 3.3 and Sect. 4.3.

Conclusion validity depends on the application of appropriate statistical tests. As we discussed throughout the paper, frequentist statistics are often hard to apply correctly in a way that is also functional to answering a study’s research questions; Bayesian statistics, which are more transparent and straightforward to interpret, can certainly help in this respect. Thus, conclusion validity threats are generally lower in our reanalyses than in the original analyses.

Internal validity is mainly concerned with whether causality is correctly evaluated. This depends on several details of experimental design that are generally independent of whether frequentist or Bayesian statistics are used. However, one important aspect of internal validity pertains to the avoidance of *bias*; this is where Bayesian statistics can help, thanks to their ability of weighting out many different competing models rather than restricting the analysis to two rigid hypotheses (null vs. alternative hypothesis). This aspect is particularly relevant for Sect. 4’s reanalysis of Rosetta Code data, where Bayesian modeling replaced the inflexible null-hypothesis testing.

Since they can integrate previous, or otherwise independently obtained, information in the form of priors, and are capable of modeling practically relevant information in a rich way, Bayesian statistics can help mitigate threats to *external validity*, which concern the generalizability of

findings. In particular, prior sensitivity analysis (described in Sect. 2.2.4) assessed the (in)dependence of the statistical conclusions on the choice of priors in both case studies (Sect. 3.3.2 and Sect. 4.3.3), thus giving a clearer picture of the generalizability of their findings.

6.2 Analytics Bad Smells

Menzies and Shepperd [64] recently proposed a list of 12 “bad smells”—indicators of poor practices in analytics research. Let’s outline how using Bayesian statistics can help avoid most of these pitfalls.

- 1) **Not interesting.** While whether a study is interesting mainly depends on its design and topic, Bayesian statistics make it easier to model quantities that are of direct practical relevance.
- 2) **Not using related work.** Sect. 1.1 indicates that Bayesian statistics is hardly ever used in empirical software engineering research, suggesting that there is ample room for progress in this direction.
- 3) **Using deprecated or suspect data.** This poor practice is independent of the statistical techniques that are used; in our reanalyses, the original studies’ data is of good quality to the best of our knowledge.
- 4) **Inadequate reporting.** We make available all raw data, as well as the analysis scripts, we used in our reanalyses to help reproducibility and further analyses. Note that the raw data was already released by the original studies [18], [71].
- 5) **Underpowered studies.** The Bayesian focus on type *M* and type *S* errors (Sect. 2.2.2), and its prior sensitivity analysis (Sect. 2.2.4), help assess the limitations of a study that really matter to achieve practical significance.
- 6) **$p < 0.05$ and all that!** It should be clear that one of the main recommendations of this paper is to dispense with *p*-values and statistical hypothesis testing.
- 7) **Assumptions (e.g., of normality) in statistical analysis.** Bayesian statistical modeling makes assumptions visible, and hence harder to apply without justification. Sensitivity analysis is a very effective way of assessing the impact of different assumptions.
- 8) **No data visualization.** Visualization is a valuable practice regardless of the kind of statistical techniques that are being used. Since Bayesian analysis provides full posterior distributions (as opposed to scalar summaries), visualization is a natural way to inspect details of an analysis’s outcome—as shown, for example, in Fig. 2, Fig. 3, and Fig. 6.
- 9) **Not exploring stability.** The practice of sensitivity analysis (Sect. 2.2.4) is precisely a form of robustness assessment.

The three remaining “smells” (*not tuning*, *not exploring simplicity*, and *not justifying choice of learners*) mainly pertain to details of how statistical inference is carried out. Since all of our models are simple (conceptually, and to analyze using standard tools), these smells are not really relevant, but they might be if we used significantly more complex models. In these cases too, frequentist statistics—with its rigid procedures and unintuitive results—does not seem to be at an advantage.

27. <https://www.r-project.org/>

28. <https://www.scipy.org/>

29. <https://julialang.org/>

30. <https://jasp-stats.org/>

6.3 Limitations of Bayesian Statistics

This paper discussed at length the limitations of frequentist statistics. Let us now outline some limitations of Bayesian statistics.

Experimental design and data quality are primary determinants of the value of every data analysis's results. This is true when applying Bayesian statistics as much as it is when applying frequentist statistics. In particular, Bayesian analysis techniques are not a silver bullet, and cannot fix a botched study design or data that is irrelevant or spurious.

In contrast to frequentist techniques such as null hypothesis testing—which can be used largely as a black box—applying Bayesian statistical techniques requires some *modeling* effort. Modeling means more than just following a ready-made recipe, and should be based on an understanding of the data domain. This does not mean that Bayesian statistics cannot be used for exploratory studies; in fact, the transparency of Bayesian statistical models helps understand the role and impact of every assumption in our analysis, which involves more effort but is ultimately beneficial to achieving clear results.

Like with every powerful analysis technique, it takes some practice to become familiar with Bayesian techniques and tools, and to fully appreciate how they can be applied to the kinds of problems a researcher encounters most frequently. To help this process, Sect. 5 discussed some basic guidelines that should be applicable to a variety of common analyses. We plan to come up with a more comprehensive set of guidelines in future work.

Very practical limitations of Bayesian statistics in empirical software engineering may derive from the mere fact that frequentist statistics are currently dominating the research practice (see Sect. 1.1). This implies that it may be harder to publish analyses based on Bayesian statistics simply because the research community expects frequentist statistics (e.g., “statistical significance”) and is unfamiliar with Bayesian techniques as best practices. We are convinced these challenges will be only temporary; and hopefully the research community will be receiving favorably the usage of Bayesian statistics. Some of Bayesian data analysis's fundamental features can be key to its practical acceptance—in particular, its focus on practical significance, and the flexibility and understandability of its models and techniques.

7 DISCUSSION

Frequentist statistics remain the standard choice for data analysis in empirical software engineering, as it offers frameworks that are easy to apply and widely accepted as good practices. This situation is unfortunate, because frequentist techniques, as applied in practice, are plagued by a number of issues that undermine their effectiveness in practice:

- statistical hypothesis testing over-simplifies complex analyses by phrasing them as binary choices between a straw-man null hypothesis and an alternative hypothesis;
- *p*-values are, in general, unsound to assess statistical significance;

- multiple tests further challenge null-hypothesis testing, and there is no consensus on how to deal with them in a way that is safe and consistent;
- effect sizes and confidence intervals are useful measures, but can be surprisingly tricky to interpret correctly in terms of practically significant quantities.

In this paper, we illustrated concrete instances of the above issues from an empirical software engineering perspective. By using Bayesian statistics instead, we demonstrated gains in the rigor, clarity, and generality of the empirical analyses of Ceccato et al. [18] and Nanz and Furia [71]:

- Bayesian analysis stresses precise modeling of assumptions, uncertainty, and domain features, which help gain a deep understanding of the data;
- applying Bayes' theorem gives a full *posterior distribution* of each estimated parameter, which quantifies the imprecision of the estimate and supports both visual inspections and numerical analyses;
- the problem of multiple testing is immaterial, both because there is no “hypothesis testing” as such and because Bayesian analysis's quantitative posteriors support reasoning about the actual sign and magnitude of an effect rather than about possible compounding of errors that may or may not matter in practice;
- *priors* naturally incorporate assumptions and expectations (including from previous studies), and help understand to what extent analysis results depend on specific assumptions;
- the numeric simulation-based tools implementing Bayesian analysis techniques offer a great flexibility and support *predictions* in a variety of complex scenarios.

Based on these results, we recommend using Bayesian statistics in empirical software engineering data analysis. Besides supporting nuanced analyses in a variety of domains, they have the potential to step up generality and predictive ability, thus making progress towards the ultimate goal of improving software engineering practice in terms of effectiveness and efficiency. By using informed priors based on previous, related work they can further help link several empirical software engineering studies together, and show what we currently can conclude based on them taken as a whole. This could help develop fruitful scientific practices where each research contribution builds on the others, and they collectively make progress towards settling important problems. By simulating real-world scenarios that are relevant for practitioners such practices could also help better connect software engineering research to practice—thus fostering a greater impact of the former on the latter.

Acknowledgments

We would like to thank Jonah Gabry (Columbia University, USA), for discussions regarding some of this paper's content. The first author also thanks for interesting discussions the participants to the EAQSE (Empirical Answers to Questions in Software Engineering) workshop held in Villebrumier, France in November 2018, where he presented a preliminary subset of this article's content.

REFERENCES

- [1] M. Aickin and H. Gensler. Adjusting for multiple testing when reporting research results: the Bonferroni vs Holm methods. *American journal of public health*, 86(5):726–728, 1996.
- [2] Valentin Amrhein, Sander Greenland, and Blake McShane. Scientists rise up against statistical significance. *Nature*, 567:305–307, 2019.
- [3] Valentin Amrhein, David Trafimow, and Sander Greenland. Inferential statistics as descriptive statistics: there is no replication crisis if we don't expect replication. *PeerJ Preprints*, 6:e26857v4, October 2018.
- [4] Andrea Arcuri and Lionel C. Briand. A practical guide for using statistical tests to assess randomized algorithms in software engineering. In *Proceedings of the 33rd International Conference on Software Engineering (ICSE 2011)*, pages 1–10. ACM, 2011.
- [5] Andrea Arcuri and Lionel C. Briand. A hitchhiker's guide to statistical tests for assessing randomized algorithms in software engineering. *Softw. Test., Verif. Reliab.*, 24(3):219–250, 2014.
- [6] Alberto Bacchelli and Christian Bird. Expectations, outcomes, and challenges of modern code review. In Notkin et al. [73], pages 712–721.
- [7] Monya Baker. Statisticians issue warning over misuse of P values. *Nature*, 531(151), 2016.
- [8] David Barber. *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012.
- [9] Gabriele Bavota, Bogdan Dit, Rocco Oliveto, Massimiliano Di Penta, Denys Poshyvanyk, and Andrea De Lucia. An empirical study on the developers' perception of software coupling. In Notkin et al. [73], pages 692–701.
- [10] A. Bener and A. Tosun. If it is software engineering, it is (probably) a Bayesian factor. In Menzies et al. [65].
- [11] D. J. Benjamin, J. O. Berger, M. Johannesson, B. A. Nosek, and E. et al. Wagenmakers. Redefine statistical significance. *Human Nature Behavior*, 2017.
- [12] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society*, 57(1):125–133, 1995.
- [13] Antonia Bertolino, Gerardo Canfora, and Sebastian G. Elbaum, editors. *37th IEEE/ACM International Conference on Software Engineering, ICSE 2015, Volume 1*. IEEE Computer Society, 2015.
- [14] Steve Brooks, Andrew Gelman, Galin L. Jones, and Xiao-Li Meng, editors. *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC, 2011.
- [15] Nélio Cacho, Thiago César, Thomas Filipe, Eliezio Soares, Arthur Cassio, Rafael Souza, Israel García, Eiji Adachi Barbosa, and Alessandro Garcia. Trading robustness for maintainability: an empirical study of evolving C# programs. In Jalote et al. [50], pages 584–595.
- [16] Bob Carpenter, Andrew Gelman, Matthew Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Gu, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1):1–32, 2017.
- [17] Jeffrey C. Carver, Natalia Juristo Juzgado, Maria Teresa Baldassarre, and Sira Vegas. Replications of software engineering experiments. *Empirical Software Engineering*, 19(2):267–276, 2014.
- [18] Mariano Ceccato, Alessandro Marchetto, Leonardo Mariani, Cu D. Nguyen, and Paolo Tonella. Do automatically generated test cases make debugging easier? an experimental assessment of debugging effectiveness and efficiency. *ACM Transactions on Software Engineering and Methodology*, 25(1), 2015. Article 5.
- [19] Thierry Titcheu Chekam, Mike Papadakis, Yves Le Traon, and Mark Harman. An empirical study on mutation, statement and branch coverage fault revelation that avoids the unreliable clean program assumption. In Uchitel et al. [92], pages 597–608.
- [20] Junjie Chen, Wenxiang Hu, Dan Hao, Yingfei Xiong, Hongyu Zhang, Lu Zhang, and Bing Xie. An empirical comparison of compiler testing techniques. In Dillon et al. [27], pages 180–190.
- [21] Shauvik Roy Choudhary, Mukul R. Prasad, and Alessandro Orso. X-PERT: accurate identification of cross-browser issues in web applications. In Notkin et al. [73], pages 702–711.
- [22] Jacob Cohen. The earth is round ($p < .05$). *American Psychologist*, 49(12):997–1003, 1994.
- [23] L. Jonathan Cohen. *An introduction to the philosophy of induction and probability*. Oxford University Press, 1989.
- [24] Confusion of the inverse. http://rationalwiki.org/wiki/Confusion_of_the_inverse, February 2016.
- [25] Ermira Daka, José Campos, Jonathan Dorn, Gordon Fraser, and Westley Weimer. Generating readable unit tests for Guava. In *Proceedings of SSBSE*, volume 9275 of *Lecture Notes in Computer Science*, pages 235–241. Springer, 2015.
- [26] Premkumar T. Devanbu, Thomas Zimmermann, and Christian Bird. Belief & evidence in empirical software engineering. In Dillon et al. [27], pages 108–119.
- [27] Laura K. Dillon, Willem Visser, and Laurie Williams, editors. *Proceedings of the 38th International Conference on Software Engineering, ICSE 2016*. ACM, 2016.
- [28] Allen B. Downey. *Think Bayes*. O'Reilly Media, 2013.
- [29] Neil A. Ernst. Bayesian hierarchical modelling for tailoring metric thresholds. In *Proceedings of the 15th International Conference on Mining Software Repositories (MSR)*, pages 587–591. ACM, 2018.
- [30] Stephen E. Fienberg. When did Bayesian inference become 'Bayesian'? *Bayesian Anal.*, 1(1):1–40, 03 2006.
- [31] Carlo A. Furia. Bayesian statistics in software engineering: Practical guide and case studies. <http://arxiv.org/abs/1608.06865>, August 2016.
- [32] Carlo A. Furia. What good is Bayesian data analysis for software engineering? In *Proceedings of the 39th International Conference on Software Engineering (ICSE) Companion (Posters Track)*, pages 374–376. ACM, 2017.
- [33] Carlo A. Furia, Dino Mandrioli, Angelo Morzenti, and Matteo Rossi. *Modeling Time in Computing*. Monographs in Theoretical Computer Science. An EATCS series. Springer, 2012.
- [34] Jonah Gabry, Daniel Simpson, Aki Vehtari, Michael Betancourt, and Andrew Gelman. Visualization in bayesian workflow. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 182(2):389–402, 2019.
- [35] L. García. Escaping the Bonferroni iron claw in ecological studies. *Oikos*, 105(3):657–663, 2004.
- [36] Andrew Gelman. The problems with p -values are not just with p -values. *The American Statistician*, 2016. Online discussion: http://www.stat.columbia.edu/~gelman/research/published/asa_pvalues.pdf.
- [37] Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis*. CRC Texts in Statistical Science. Chapman & Hall, 3rd edition, 2013.
- [38] Andrew Gelman and Jennifer Hill. *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press, 2007.
- [39] Andrew Gelman, Jennifer Hill, and Masanao Yajima. Why we (usually) don't have to worry about multiple comparisons. *Journal of Research in Educational Effectiveness*, 5:189–211, 2012.
- [40] Andrew Gelman and Francis Tuerlinckx. Type s error rates for classical and bayesian single and multiple comparison procedures. *Computational Statistics*, 15(3):373–390, 2000.
- [41] Jeff Gill. The insignificance of null hypothesis significance testing. *Political Research Quarterly*, 52(3):647–674, 1999.
- [42] Francisco Gomes de Oliveira Neto, Richard Torkar, Robert Feldt, Lucas Gren, Carlo A. Furia, and Ziwei Huang. The evolution of statistical analysis in empirical software engineering research. <https://arxiv.org/abs/1706.00933>, June 2017.
- [43] Irving John Good. Explicativity, corroboration, and the relative odds of hypotheses. *Synthese*, 30(1/2):39–73, 1975.
- [44] Steven N. Goodman. Toward evidence-based medical statistics. 1: The p value fallacy. *Annals of Internal Medicine*, 130(12):995–1004, 1999.
- [45] Steven N. Goodman. A dirty dozen: Twelve p -value misconceptions. *Seminars in Hematology*, 45(3):135–140, 2008.
- [46] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2nd edition, 2009.
- [47] R. Hoekstra, R. D. Morey, J. N. Rouder, and E.-J. Wagenmakers. Robust misinterpretation of confidence intervals. *Psychonomic Bulletin & Review*, 13:1033–1037, 2014.
- [48] Myles Hollander and Douglas A. Wolfe. *Nonparametric Statistical Methods*. John Wiley & Sons, 2nd edition, 1999.
- [49] John P. A. Ioannidis. Why most published research findings are false. *PLoS Med*, 2(8), 2005.
- [50] Pankaj Jalote, Lionel C. Briand, and André van der Hoek, editors. *36th International Conference on Software Engineering, ICSE '14*. ACM, 2014.
- [51] Andreas Jedlitschka, Natalia Juristo Juzgado, and H. Dieter Rombach. Reporting experiments to satisfy professionals' information needs. *Empirical Software Engineering*, 19(6):1921–1955, 2014.

- [52] Harold Jeffreys. *Theory of Probability*. Oxford Classic Texts in the Physical Sciences. Oxford University Press, 3rd edition, 1998.
- [53] Mitchell Joblin, Sven Apel, Claus Hunsen, and Wolfgang Mauerer. Classifying developers into core and peripheral: an empirical study on count and network metrics. In Uchitel et al. [92], pages 164–174.
- [54] Miryung Kim, Thomas Zimmermann, Robert DeLine, and Andrew Begel. The emerging role of data scientists on software development teams. In Dillon et al. [27], pages 96–107.
- [55] Miryung Kim, Thomas Zimmermann, Robert DeLine, and Andrew Begel. The emerging role of data scientists on software development teams. In Dillon et al. [27], pages 96–107.
- [56] John K. Kruschke. Bayesian assessment of null values via parameter estimation and model comparison. *Perspect. Psychol. Sci.*, 6(3):299–312, 2011.
- [57] John K. Kruschke. *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and Stan*. Academic Press, 2nd edition, 2014.
- [58] John K. Kruschke and Torrin M. Liddell. The bayesian new statistics: Hypothesis testing, estimation, meta-analysis, and power analysis from a bayesian perspective. *Psychon. Bull Rev.*, 25:178–206, 2018.
- [59] Davy Landman, Alexander Serebrenik, and Jurgen J. Vinju. Challenges for static analysis of Java reflection: literature review and empirical study. In Uchitel et al. [92], pages 507–518.
- [60] Irene Manotas, Christian Bird, Rui Zhang, David C. Shepherd, Ciera Jaspán, Caitlin Sadowski, Lori L. Pollock, and James Clause. An empirical study of practitioners’ perspectives on green software engineering. In Dillon et al. [27], pages 237–248.
- [61] Daniel Matichuk, Toby C. Murray, June Andronick, D. Ross Jeffery, Gerwin Klein, and Mark Staples. Empirical study towards a leading indicator for cost of formal software verification. In Bertolino et al. [13], pages 722–732.
- [62] Richard McElreath. *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. CRC Texts in Statistical Science. Chapman & Hall, 2015.
- [63] Blakeley B. McShane, David Gal, Andrew Gelman, Christian Robert, and Jennifer L. Tackett. Abandon statistical significance. *The American Statistician*, 73(S1):235–245, 2019.
- [64] Tim Menzies and Martin Shepperd. Bad smells in software analytics papers. *Information and Software Technology*, 112:35–47, 2019.
- [65] Tim Menzies, Laurie Williams, and Thomas Zimmermann, editors. *Perspectives on Data Science for Software Engineering*. Morgan Kaufmann, 2016.
- [66] R. G. Miller. *Simultaneous statistical inference*. Springer-Verlag, 2nd edition, 1981.
- [67] Richard D. Morey, Rink Hoekstra, Jeffrey N. Rouder, Michael D. Lee, and Eric-Jan Wagenmakers. The fallacy of placing confidence in confidence intervals. *Psychonomic Bulletin & Review*, 23(1):103–123, 2016.
- [68] Saman Muthukumarana and Tim B. Swartz. Bayesian analysis of ordinal survey data using the Dirichlet process to account for respondent personality traits. *Communications in Statistics – Simulation and Computation*, 43(1):82–98, 2014.
- [69] Sarah Nadi, Thorsten Berger, Christian Kästner, and Krzysztof Czarnecki. Mining configuration constraints: static analyses and empirical results. In Jalote et al. [50], pages 140–151.
- [70] S. Nakagawa. A farewell to Bonferroni: the problems of low statistical power and publication bias. *Behavioral Ecology*, 15(6):1044–1045, 2004.
- [71] Sebastian Nanz and Carlo A. Furia. A comparative study of programming languages in Rosetta Code. In *Proceedings of the 37th International Conference on Software Engineering (ICSE)*, pages 778–788. ACM, 2015.
- [72] William S. Noble. How does multiple testing correction work? *Nature Biotechnology*, 27:1135–1137, 2009.
- [73] David Notkin, Betty H. C. Cheng, and Klaus Pohl, editors. *35th International Conference on Software Engineering, ICSE ’13*. IEEE Computer Society, 2013.
- [74] Carlos Pacheco and Michael D. Ernst. Randoop: Feedback-directed random testing for Java. In *Companion to the 22Nd ACM SIGPLAN Conference on Object-oriented Programming Systems and Applications Companion, OOPSLA ’07*, pages 815–816. New York, NY, USA, 2007. ACM.
- [75] Sebastiano Panichella, Annibale Panichella, Moritz Beller, Andy Zaidman, and Harald C. Gall. The impact of test case summaries on bug fixing performance: an empirical investigation. In Dillon et al. [27], pages 547–558.
- [76] Mike Papadakis, Yue Jia, Mark Harman, and Yves Le Traon. Trivial compiler equivalence: A large scale empirical study of a simple, fast and effective equivalent mutant detection technique. In Bertolino et al. [13], pages 936–946.
- [77] Harold Pashler and Eric-Jan Wagenmakers. Editors’ introduction to the special section on replicability in psychological science: A crisis of confidence? *Perspectives on Psychological Science*, 7(6):528–530, 2012.
- [78] T. Perneger. What’s wrong with Bonferroni adjustments. *British Medical Journal*, 316:1236–1238, 1998.
- [79] Marian Petre. UML in practice. In Notkin et al. [73], pages 722–731.
- [80] C. Radhakrishna Rao. R. a. fisher: The founder of modern statistics. *Statist. Sci.*, 7(1):34–48, 02 1992.
- [81] Eric S. Raymond. *The New Hacker’s Dictionary*. MIT Press, 3rd edition, 1996.
- [82] H.M. Reid. *Introduction to Statistics: Fundamental Concepts and Procedures of Data Analysis*. SAGE Publications, 2013.
- [83] Rosetta code. <http://rosettacode.org/>, Aug 2016.
- [84] Marija Selakovic and Michael Pradel. Performance issues and optimizations in JavaScript: an empirical study. In Dillon et al. [27], pages 61–72.
- [85] S. Siegel and N. J. Castellan. *Nonparametric statistics for the behavioral sciences*. McGraw-Hill, 2nd edition, 1988. Implemented in R’s package `pgirmess` as function `kruskalmc`.
- [86] Janet Siegmund, Norbert Siegmund, and Sven Apel. Views on internal and external validity in empirical software engineering. In Bertolino et al. [13], pages 9–19.
- [87] Joseph P. Simmons, Leif D. Nelson, and Uri Simonsohn. False-positive psychology. *Psychological Science*, 22(11):1359–1366, 2011.
- [88] Klaas-Jan Stol, Paul Ralph, and Brian Fitzgerald. Grounded theory in software engineering research: a critical review and guidelines. In Dillon et al. [27], pages 120–131.
- [89] Student. The probable error of a mean. *Bibliometrika*, 6(1):1–25, 1908.
- [90] Guoxin Su and David S. Rosenblum. Perturbation analysis of stochastic systems with empirical distribution parameters. In Jalote et al. [50], pages 311–321.
- [91] The computer language benchmarks game. <http://benchmarksgame.alioth.debian.org/>, Aug 2016.
- [92] Sebastián Uchitel, Alessandro Orso, and Martin P. Robillard, editors. *Proceedings of the 39th International Conference on Software Engineering, ICSE 2017, Buenos Aires, Argentina, May 20-28, 2017*. IEEE / ACM, 2017.
- [93] Phillip Merlin Uesbeck, Andreas Stefik, Stefan Hanenberg, Jan Pedersen, and Patrick Daleiden. An empirical study on the impact of C++ lambdas and programmer experience. In Dillon et al. [27], pages 760–771.
- [94] Noah N N van Dongen, Johnny van Doorn, Quentin F Gronau, Don van Ravenzwaaij, Rink Hoekstra, Matthias Haucke, Daniel Lakens, Christian Hennig, Richard D Morey, Saskia Homer, and et al. Multiple perspectives on inference for two simple statistical scenarios. *PsyArXiv*, Dec 2018. <https://psyarxiv.com/ue5wb>.
- [95] A. Vargha and H. D. Delaney. A critique and improvement of the CL common language effect size statistics of McGraw and Wong. *Journal of Educational and Behavioral Statistics*, 25(2):101–132, 2000.
- [96] R. M. Warner. *Applied Statistics: From Bivariate Through Multivariate Techniques: From Bivariate Through Multivariate Techniques*. SAGE Publications, 2012.
- [97] Ronald L. Wasserstein and Nicole A. Lazar. The ASA statement on *p*-values: Context, process, and purpose. *The American Statistician*, 70(2):129–133, 2016. <https://www.amstat.org/asa/files/pdfs/P-ValueStatement.pdf>.
- [98] Ronald L. Wasserstein and Nicole A. Lazar. The ASA’s statement on *p*-values: Context, process, and purpose. *The American Statistician*, 70(2):129–133, 2016.
- [99] Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, and Björn Regnell. *Experimentation in Software Engineering*. Springer, 2012.
- [100] Aiko Fallas Yamashita and Leon Moonen. Exploring the impact of inter-smell relations on software maintainability: an empirical study. In Notkin et al. [73], pages 682–691.
- [101] J. H. Zar. *Biostatistical analysis*. Pearson Prentice Hall, 5th edition, 2010.

- [102] Hongyu Zhang, Liang Gong, and Steven Versteeg. Predicting bug-fixing time: an empirical study of commercial software projects. In Notkin et al. [73], pages 1042–1051.
- [103] Hao Zhong and Zhendong Su. An empirical study on real bug fixes. In Bertolino et al. [13], pages 913–923.