Improving statistical inferences: Can a MOOC reduce statistical misconceptions about p -values, confidence intervals, and Bayes factors?

Arianne Herrera-Bennett¹, Moritz Heene¹, Daniël Lakens², Stefan Ufer³

¹Department of Psychology, Ludwig Maximilian University of Munich, Germany

²Department of Industrial Engineering & Innovation Sciences, Human Technology Interaction, Eindhoven University of Technology, Netherlands

³Department of Mathematics, Ludwig Maximilian University of Munich, Germany

OSF project page: <https://osf.io/5x6k2/>

Supplementary Information: <https://osf.io/87emq/>

Submitted to *Meta-Psychology*. Click here to follow the fully transparent editorial process of this submission. Participate in open peer review by commenting through hypothes.is directly on this preprint.

IMPROVING STATISTICAL INFERENCES

Abstract

Statistical reforms in psychology and surrounding domains have been largely motivated by doubts surrounding the methodological and statistical rigor of researchers. In light of high rates of p -value misunderstandings, against the backdrop systematically raised criticisms of NHST, some authors have deemed misconceptions about p -value “impervious to correction” (Haller & Krauss, 2002, p. 1), while others have advocated for the entire abandonment of statistical significance in place of alternative methods of inference (e.g., Amrhein, Greenland, & McShane, 2019). Surprisingly little work, however, has empirically investigated the extent to which statistical misconceptions can be improved reliably within individuals, nor whether these improvements are transient or maintained. The current study ($N = 2,320$) evaluated baseline misconception rates of p -value, confidence interval, and Bayes factor interpretations among online learners, as well as rates of improvement in accuracy across an 8-week massive open online course (MOOC). Results demonstrated statistically significant improvements in accuracy rates, across all three concepts, for immediate learning (first post-test), as well as support for retained learning until week 8 (second post-test). The current work challenges the idea that statistical misconceptions are impervious to correction, and highlights the importance of pinpointing those more problematic misunderstandings that may warrant more efforts to support improvement. Before we abandon one of the most widely used approaches to statistical inferences in place of an alternative, it is worth exploring whether it is first possible to improve existing misunderstandings amid current methods used.

Keywords: MOOC, p -values, confidence intervals, Bayes factors, statistical misconceptions

IMPROVING STATISTICAL INFERENCES

Improving Statistical Inferences: Can a MOOC Reduce Statistical Misconceptions about *P*-values, Confidence Intervals and Bayes Factors?

Introduction

Statisticians have argued for decades that the misuse of statistics and the pervasiveness of statistical misconceptions among researchers stems from fundamental misunderstandings about how to draw statistical inferences (e.g., Craig, Elson, & Metze, 1976; Carver, 1978; Falk & Greenbaum, 1995). This criticism has focussed especially on the widespread use of null hypothesis significance testing (NHST; for an overview of criticisms, see Nickerson, 2000) and the overreliance on *p*-values, which continue to dominate the published literature in psychology (approx. 97% of articles across 10 leading journals; Cumming et al., 2007). In light of periodically observed high rates of *p*-value misconceptions (e.g., Oakes, 1986; Haller & Krauss, 2002; Badenes-Ribera, Frías-Navarro, Monerde-i-Bort, & Pascual-Soler, 2015; Lyu, Peng, & Hu, 2018; Lyu, Xu, Zhao, Zuo, & Hu, 2020), and against the backdrop of NHST criticisms systematically raised over the last eighty years (e.g., Berkson, 1938; Wasserstein & Lazar, 2016), some authors have deemed *p*-value misinterpretations “impervious to correction” (Haller & Krauss, 2002, p. 1). Others have advocated for the abandonment of statistical significance as a method of statistical inference altogether (e.g., Amrhein, Greenland, & McShane, 2019).

The trouble with such absolute claims and extreme solutions is that, even though misconception rates might appear stable, studies up till now have been based on surveys of existing knowledge, without examining whether these misconception rates can be reduced through education. Previous studies typically involved cross-sectional rather than longitudinal designs (i.e. Oakes, 1986; Haller & Krauss, 2002; Badenes-Ribera et al., 2015; Lyu, et al., 2018; Lyu et al., 2020). As

IMPROVING STATISTICAL INFERENCES

such, the tendency to accept that NHST misconceptions are “impervious to correction” appears to be largely a product of single time-point assessments measuring misconceptions in a sample that had received traditional statistics education. Therefore, surveys like Oakes’s (1986) mainly inform us that misconceptions remain relatively high after traditional statistics education, rather than that these surveys provide evidence that instructional interventions are ineffective.

Abandoning p -values, and moving toward the use of alternative or complementary statistical techniques, such as confidence intervals, Bayesian statistics (Wagenmakers, Wetzels, Borsboom, & van der Maas, 2011), and effect sizes (e.g., Bakan, 1966; Falk & Greenbaum, 1995; Cohen, 1990; Rozeboom, 1997; Wasserstein & Lazar, 2016) is not necessarily a solution, since “mindless statistics are not limited to p -values” (Lakens, 2019, p. 9): Alternative approaches such as confidence intervals and Bayesian inference carry their own possible misinterpretations (for overviews, see Greenland et al., 2016; Morey, Hoekstra, Rouder, Lee, & Wagenmakers, 2016). As such, simply replacing p -values with other statistical tools is unlikely to resolve the problem whereby, in the process of drawing statistical inferences, researchers fall prey to misconceptions. Before we abandon one of the most widely used approaches to statistical inferences due to misuse, it is worth exploring whether it is possible to improve the correct understanding of p -values (Lakens, 2019; Shafer, 2019).

To this end, the current study investigated misconception rates among learners in the context of an 8-week Coursera massive open online course (MOOC) geared broadly toward the instruction of core statistical concepts relevant to psychological research. To extend previous work, our sample was not limited to psychologists, but included learners across various disciplines. Moreover, while previous research has focused on inferences about statistically significant p -values (i.e. interpreting the meaning of $p = .01$ or $.001$), our work also examines interpretations of non-significant p -values

IMPROVING STATISTICAL INFERENCES

(i.e. $p = .30$), as well as confidence intervals and Bayes factors. The current paper asks: Are statistical misconceptions really impervious to correction, or can they be systematically improved within individuals through education, as indicated by repeated assessments over time?

Statistical Misconceptions: Frequencies and Common Fallacies

Since the seminal work of Oakes (1986), who found that 97% of academic psychologists fell prey to at least one erroneous p -value interpretation, misconception rates have appeared to remain relatively stable. Similar frequencies were observed among German ($N = 113$; Haller & Krauss, 2002), Spanish ($N = 418$; Badenes-Ribera et al., 2015), and Chinese ($N = 246$ and $N = 1,479$; Lyu et al., 2018 and Lyu et al., 2020, resp.) samples, and were not limited to the field of psychology ($N = 221$ communication researchers; Rinke & Schneider, 2018). Although accuracy rates varied across academic status (Haller & Krauss, 2002), and qualification (Lyu et al., 2018), differences were only marginal. In fact, even methodology researchers and instructors, who made comparatively fewer incorrect interpretations (e.g., 80% making at least one mistake; Haller & Krauss, 2002), were not immune to general misunderstandings (see also Lecoutre, Poitevineau, & Lecoutre, 2003; Badenes-Ribera et al., 2015). In response to such observations, some authors have advocated for a shift toward more statistical approaches deemed more intuitive, such as Bayesian inference, criticizing “the fundamental inadequacy of NHST to the needs of the users” (Lecoutre et al., 2003, p. 42).

It is worth pointing out that most of the misconception rate literature reported above, including Oakes’s (1986) initial work, pertain specifically to the percentage of individuals who made one or more mistakes; in other words, misconception rates are coarsely dichotomized into the proportion of individuals who did (or did not) obtain a perfect sum score. This means that, at face value, findings are inherently ambiguous, and may misrepresent the extent to which academics grasp

IMPROVING STATISTICAL INFERENCES

the concept of significance testing. Take for instance the percentage of psychology students (100%) or methodology instructors (80%), in Haller & Krauss's (2002) study, who fell prey to at least one misconception. Such findings, without a question, serve as ample evidence for the need to improve statistical education; that said, mean number of wrongly endorsed items, respectively 2.5 and 1.9 (out of 6.0), arguably indicate that individuals are not wholly misguided when it comes to interpreting the correct meaning of p -values. In other words, while it is clear that at the group level the vast majority of individuals tend toward imperfect scores, at the individual level participants on average tend to get more items right than they get wrong. It is possible that specific clarification of these more problematic items or non-intuitive concepts might be what is especially warranted in order to improve overall misconception rates.

The current research strives to tackle this goal, specifically drawing upon the work by Badenes-Ribera and colleagues (2015) who identified the rates of four specific p -value fallacies, namely the *inverse probability* (Shaver, 1993; Kirk, 1996), *replication* (Carver, 1978; Fidler, 2005), *effect size* (Gliner, Vaske, & Morgan, 2001), and *clinical or practical significance* (Kirk, 1996) fallacies. The *inverse probability* misconception derives fundamentally from falsely assuming that one can draw conclusions about the probability of a theory or hypothesis, given sample data. Instead, the correct interpretation entails the converse: The p -value determines the probability of observing data (or more extreme results), *contingent* on having assumed the null-hypothesis is true. The formal definition can be expressed as follows: $P(X \geq x|H_0)$ or $P(X \leq x|H_0)$, for right- versus left-tailed events, where X represents a random variable and x the observed event. The *replication*, *effect size* and *clinical or practical significance* fallacies are misconceptions in which implications of a significant p -value are misinterpreted. Specifically, in the case of the *replication* fallacy, individuals falsely take the p -value probability (e.g., $p = .03$, or 3%) as the complement of the replication

IMPROVING STATISTICAL INFERENCES

probability (i.e. $1 - .03 = .97$, or 97%). For the *effect size* fallacy, a significant p -value is falsely assumed to necessarily entail a large effect size (whereas a non-significant p -value necessarily implies a small effect size); and in the case of the *clinical or practical significance* fallacy, a finding that is statistically significance is conflated with the idea that it is practically important or clinically meaningful (and vice-versa).

Badenes-Ribera et al. (2015) demonstrated that academic psychologists were particularly prone to the *inverse probability* fallacy (93.8% error rate), a fallacy which has been deemed “the most common, and arguably the most damaging, misinterpretation of p value (Oakes, 1986)” (Kalinowski, Fidler, & Cumming, 2008). In contrast, the *replication*, *effect size* and *clinical or practical significance* fallacies incurred relatively lower error rates (resp. 34.7%, 13.2%, & 35.2%). Interestingly, nearly 50 years prior, Bakan (1966) referred to such misinterpretations as a form of researcher bias; that is, an intrinsic misattribution about p -value characteristics stemming from placing too much weight on the role of significance testing. When statistical tests or outcomes assume the “burden of scientific inference” (Bakan, 1966, p. 423), researchers have the tendency in “attributing to the test of significance characteristics which it does not have, and overlook characteristics that it does have” (Bakan, 1966, p. 428). Similar impressions were shared by other critics of NHST, who considered the approach harmful “because such tests do not provide the information that many researchers assume they do” (Shaver, 1993, p. 294), and in turn allow a p -value to be “interpreted to mean something it is not” (Carver, 1978, p. 392). Whether this bias or tendency to inappropriately overattribute meaning to p -values stems from honest error (e.g., as a product of general ignorance), ritualized practice (or what Bakan referred to as “automaticity of inference”; 1966, p. 423), or a misuse of the flexibility of interpretation because it fits with one’s goals (i.e. researcher degrees of freedom; Simmons, Nelson, & Simonsohn, 2011, p. 1), cannot be

IMPROVING STATISTICAL INFERENCES

said for certain. Nevertheless, what is clear is that when it comes the process of quantifying the outcome or implications of a study, it is common for the measure of statistical significance (i.e. p -value) to be confounded with measures of replicability (i.e. *replication* fallacy), size of impact (i.e., *effect size* fallacy), and meaningfulness (i.e. *clinical or practical significance* fallacy).

Misconceptions have also been identified with regard to confidence intervals: “CIs [confidence intervals] do not have the properties that are often claimed on their behalf” (Morey et al., 2016, p. 118). Some examples include the false belief that overlapping confidence intervals necessarily imply a statistically non-significant mean difference (Belia, Fidler, Williams, & Cumming, 2005), that confidence intervals reflect the probability of containing the true population value (deemed the “fundamental confidence fallacy” Morey et al., 2016; Hoekstra, Morey, Rouder, & Wagenmakers, 2014), as well as the confidence-level misconception, i.e. the (erroneous) belief “that a 95% CI [confidence interval] will on average capture 95% of replication means” (Cumming, Williams, & Fidler, 2004, p. 299; Cumming & Maillardet, 2006). Prevalent misconception rates about confidence intervals have been observed among undergraduates and graduate students, as well as researchers (see Hoekstra et al., 2014, Lyu et al., 2018; Lyu et al., 2020; Rinke & Schneider, 2018), and have lead certain authors to also recommend improved statistical education and empirical work on the use and interpretations of confidence intervals (Fidler, Thomason, Cumming, Finch, & Leeman, 2004).

Another statistic that has been proposed as a replacement to p -values are Bayes factors. Currently, no studies have examined Bayes factor misconceptions, neither with respect to baseline levels of understanding nor improvement rates across individuals. Would confidence intervals or Bayes factors fare better than p -values? And to what extent can we improve their understanding through teaching? An important first step in answering these questions is to evaluate the extent to

IMPROVING STATISTICAL INFERENCES

which individuals are familiar with each inferential method and the relative misconception rates across these three statistics.

Current Study

The current study aimed at investigating baseline misconception rates and improvements across an 8-week MOOC geared broadly toward the instruction of core statistical concepts relevant to psychological research (e.g., frequentist concepts, including p -values, confidence intervals, and multiple comparisons, as well as introductions to Bayesian statistics, effect sizes, etc.). Specifically, items measuring misconception rates focused on inferences about the meaning of p -values, confidence intervals, and Bayes factors. To date, previous studies have focused on how individuals draw inferences about statistically significant outcomes (i.e. $p = .01$ or $.001$). A comprehensive understanding of p -values, however, should not be limited to only grasping the correct meaning of significant outcomes but also non-significant outcomes, especially given the observed tendency among researchers to falsely interpret null findings as evidence for non-existent effects (Schatz, Jay, McComb, & McLaughlin, 2005; Fidler, Burgman, Cumming, Buttrose, & Thomason, 2006; Hoekstra Finch, Kiers, & Johnson, 2006; Bernardi, Chakhaia, & Leopold, 2017). Accordingly, in the sixth week of the MOOC, the concept of equivalence tests were explained, which provide one approach to interpret the absence of a meaningful effect. The study therefore made use of two scale versions in order to examine the same types of p -value fallacies for both significant (i.e. $p = .001$) as well as non-significant outcomes (i.e. $p = .30$).

Procedure

Building on previous work, and through extensive pilot tests, a 14-item True/False scale was developed to target possible p -value, Bayes factor, and confidence interval misconceptions:

IMPROVING STATISTICAL INFERENCES

Specifically, the scale was composed of eight p -value items (seven targeting the four aforementioned fallacies, and one item for the correct p -value interpretation), three Bayes factor items, and three confidence interval items (see Table 1). While we did not assume the scale to be unidimensional, and in fact found stronger support for a 2-factor model (distinguishing between frequentist and Bayesian subdimensions), we also assessed fit indices for the one-factor model of the scale at *pretest* ($n = 1,915$) and found reasonably acceptable fit (CFI = .90, RMSEA = .05, SRMR = .04). Because composite scores (i.e. sum scores across all 14 items) were used to infer overall improvements in learning, this warranted consideration of the reliability estimate of the full scale, which was found to be acceptable (omega coefficient = .78; for further details on piloting, scale dimensionality, and reliabilities, see Supplementary Information).

Regarding scale implementation, we followed the approach of Haller & Krauss (2002) and prefaced items with the instruction that “Several or none of the statements may be correct”. We also included an “I don’t know” response option as we expected that several concepts would be completely new to some learners and wanted to discourage guessing. Additionally, the two versions of the scale provided alternate wordings for the same misconceptions: Specifically, for p -value items, version 1 and version 2 respectively framed items in terms of a statistically significant ($p = .001$) versus statistically non-significant ($p = .30$) outcome (see Appendices Q1 and Q2). Both versions were implemented (item order fixed-randomized, and scale versions counterbalanced across participants) in the form of six “Pop Quizzes” within the 8-week MOOC “*Improving your statistical inferences*” taught by Daniël Lakens on the Coursera platform. A pre-/post-test design with three measurement periods (see Figure 1) was used to measure: i) prior knowledge (*pretest*, i.e. Pop Quiz 1 administered in week 1), ii) immediate improvement (*post-test 1*, i.e. Pop Quizzes 2-5 assessed across weeks 1-5), and iii) retained learning (*post-test 2*, i.e. Pop Quiz 6 in week 8). *Post-*

IMPROVING STATISTICAL INFERENCES

test 1 items were staggered across weeks 1 to 5 in order to occur immediately after the relevant module whose content pertained to the concepts in question. Questions were incidentally clustered into four subsets (see Table 1): subsets 1 and 2 (p -value items), subset 3 (Bayes factor items), and subset 4 (confidence interval items). Accordingly, separate analyses for each of the item subsets are also reported in order to account for differences in item administration (i.e. where items occur within the progression of the course) and item completion (i.e. differences in the amount of time elapsing between any two measurements, referred to below as ‘lag’). Finally, confidence ratings of responses (7-point Likert scale ranging from 1 “*Not at all confident*” to 7 “*Very confident*”) were requested, once per Pop Quiz, in order to measure whether clarifications of concepts also reduced uncertainty in responding. Demographics were also voluntarily requested (e.g., self-rated statistical expertise level, level of education obtained) at the start of the course.

IMPROVING STATISTICAL INFERENCE

Figure 1

Figure 1 – Item implementation across the 8-week MOOC. Subsets of items were administered in form of Pop Quizzes (PQs) across the 8 weeks, and within three measurement periods: All 14 items (i.e. all 4 subsets) were first tested altogether in week 1 (*pretest* or PQ1), followed by a staggered first follow-up test (*post-test 1*, broken up into PQ2 to PQ5), and finally a second post-test in week 8 (*post-test 2* or PQ6). Each subset measured a separate set of items: subset 1 (five *p*-value items, including *inverse probability* (IP) and *replication* (R) fallacies), subset 2 (three *p*-value items, including *effect size* (ES) and *clinical or practical significance* (CPS) fallacies, as well as the correct *p*-value interpretation), subset 3 (three Bayes factor items), and subset 4 (three confidence interval items). ‘Lag’ represents the time elapsed between any 2 measurements.

IMPROVING STATISTICAL INFERENCES

Table 1

Scale items (version 1). Items from scale version 1. Each pop quiz included the following directions: “**Instructions:** Please read through the following statements, and mark each as True or False. Note that several or none of the statements may be correct.” Additionally, each of the p -value items were prefaced by “Let’s suppose that a research article indicates a value of $p = .001$ in the results section ($\alpha = .05$)”. Item options included True, False, and “I don’t know”. For scale version 2, see Appendix Q2.

Subset	Item	Fallacy
1	PV1 You have absolutely proven your alternative hypothesis (that is, you have proven that there is a difference between the population means).	Inverse probability
	PV2 You have found the probability of the null hypothesis being true ($p = .001$).	Inverse probability
	PV3 The null hypothesis has been shown to be false.	Inverse probability
	PV4 The p -value gives the probability of obtaining a significant result whenever a given experiment is replicated.	Replication
	PV5 The probability that the results of the given study are replicable is not equal to $1-p$.	Replication
2	PV6 The value $p = .001$ does not directly confirm that the effect size was large.	Effect size
	PV7 Obtaining a statistically significant result implies that the effect detected is important.	Clinical or practical significance
	PV8 The p -value of a statistical test is the probability of the observed result or a more extreme result, assuming the null hypothesis is true.	Correct interpretation
3	BF1 When a Bayesian t-test yields a $BF = 0.1$, it is ten times more likely that there is no effect than that there is an effect.	N/A
	BF2 A Bayes Factor that provides strong evidence for the null model does not mean the null hypothesis is true.	N/A
	BF3 A Bayes Factor close to 1 (inconclusive evidence) means that the effect size is small.	N/A
4	CI1 The specific 95% confidence interval observed in a study has a 95% chance of containing the true effect size.	N/A
	CI2 If two 95% confidence intervals around the means overlap, then the difference between the two estimates is necessarily non-significant ($\alpha = .05$).	N/A
	CI3 An observed 95% confidence interval does not predict that 95% of the estimates from future studies will fall inside the observed interval.	N/A

Note. Correct answers: Items that should be answered True (PV5, PV6, PV8, BF2, CI3); the rest are False.

IMPROVING STATISTICAL INFERENCES

Exclusion Criteria

Due to the flexible nature of the online course, participants could complete the pop quizzes more than once, opt out from responding altogether, and move between modules at their own pace. Our analyses only include subjects' first response attempt for each quiz, and excluded any users that did not complete measurements in the expected order (e.g., participants were excluded if they completed the *post-test* before the *pretest*). Response latencies (i.e. lag between any 2 measurement occasions, measured in hours) were corrected for skewedness, excluding outliers (± 3 median absolute deviations; Leys, Ley, Klein, Bernard, & Licata, 2013). The resulting asymmetry after outlier exclusion fell within acceptable limits (i.e. between -2 and +2; George & Mallery, 2010; see Appendix R1 in Supplementary Information).

Results

Sample

Data collection ran for 1 year (Aug 2017 – Aug 2018). Total number of MOOC learners at *pretest* was $N = 2,320$; of those learners who responded voluntary to the demographic questionnaire ($n = 611$), 57.8% were male (39% female; 3.3% opting not to specify), with a mean age of 37.93 years ($SD = 10.77$). Thirty-four percent of users reported English as their native language, and 82.0% indicated to have previously taken a statistics course before participating in the MOOC. When asked to rate their level of statistical knowledge and understanding, 38.9% rated themselves as beginners, 54.2% intermediate, and 6.9% advanced. Regarding academic experience, about a third of individuals held a bachelor's degree or lower (29.7% with high school diploma or bachelor's degree), roughly half had completed graduate-level training (51.6% with Master's or PhD degrees), and the rest held post-graduate degrees (14.4% post-doctoral degree, 4.3% professorship).

IMPROVING STATISTICAL INFERENCES

Response Trends

In keeping with typically high MOOC dropout rates (average completion rates approx. 5%; for review, see Feng, Tang, & Liu, 2019), marked response attrition was observed across the six quizzes, with total number of learners at Pop Quiz1 ($n_1 = 1,915$) dropping to $n_2 = 1,045$ (Pop Quiz 2), $n_3 = 621$ (Pop Quiz 3), $n_4 = 421$ (Pop Quiz 4), $n_5 = 371$ (Pop Quiz 5), and finally $n_6 = 276$ (Pop Quiz 6) respondents. Sample sizes for pop quizzes and reported mean scores exclude cases with missing data on any of the respective quiz items. A larger proportion of “I don’t know” responses were observed at *pretest* (27% at Pop Quiz 1) as compared to *post-tests* (approx. 1% - 3% of responses across pop quizzes 2 – 6). Mean confidence ratings increased slightly relative to baseline (Pop Quiz 1), with ratings ranging between 4.04 to 5.80 (see Table 2).

Table 2

Response trends & confidence levels. Mean proportion of correct, incorrect, and “I don’t know” responses, as well as mean confidence ratings, across pop quizzes 1 to 6. Items covered in each quiz were clustered into four subsets: subset 1 (five *p*-value items), subset 2 (three *p*-value items), subset 3 (three Bayes factor items), and subset 4 (three confidence interval items).

	Pop Quiz 1	Pop Quiz 2	Pop Quiz 3	Pop Quiz 4	Pop Quiz 5	Pop Quiz 6
Items covered	Subsets 1 – 4 (all 14 items)	Subset 1 (5 PV items)	Subset 3 (3 BF items)	Subset 2 (3 PV items)	Subset 4 (3 CI items)	Subsets 1 – 4 (all 14 items)
<i>n</i>	1,915	1,045	621	421	371	276
Correct (%)	48.06	75.26	64.93	89.73	77.43	80.84
Incorrect (%)	24.67	22.22	31.70	9.33	21.53	17.90
I don’t know (%)	27.26	2.54	3.33	0.97	0.97	1.26
mean confidence rating (7-pt Likert)	4.04	5.06	4.74	5.80	5.35	5.17

Note: PV = *p*-value, BF = Bayes factor, CI = confidence interval

IMPROVING STATISTICAL INFERENCES

Baseline Misconception Rates

Baseline accuracy rates were computed in two ways: once as the proportion of correct responses across all responses (with “I don’t know” responses coded as incorrect), and a second time as the proportion of correct responses after “I don’t know” responses were omitted (Figure 2). As respondents in the Badenes-Ribera et al. (2015) study were only provided with “True” or “False” options, whereas our study also allowed for “I don’t know” responses, the observed misconception rates are not able to be directly compared. That said, our findings were broadly consistent with past research regarding the relative difficulty levels between different p -value fallacies: Namely, the *inverse probability* fallacy (especially item PV3) was relatively problematic, whereas the item testing the *effect size* fallacy (PV6) yielded the greatest proportion of correct responses (71.8%). In our sample, the *replication* fallacy (especially PV5) also proved to be particularly problematic among respondents (with the lowest mean proportion of correct responses across all p -value items, 37%), as well as confidence intervals (approximately 30% - 51% incorrect; see Table 3). Bayes factors were less familiar as a concept across users, as demonstrated by the relatively larger proportions of responses classified as “I don’t know” (approx. 45% - 67% of respondents) as compared to the rest of the scale items. What is worth noting is relative difficulty levels across the different statistical concepts was fairly consistent regardless of the different scoring computations (i.e. whether “I don’t know” responses were coded as incorrect vs. omitted), except for the Bayes factor items: Rates showed that among those who attempted to answer the Bayes factor items (i.e. did not opt for the “I don’t know” option), accuracy was markedly higher. Specifically, when “I don’t know” responses were coded as incorrect, the 3 Bayes factor items yielded accuracy rates of 14% (BF1), 47% (BF2), and 27% (BF3); in contrast, when only True and False options were compared, accuracy was 43%, 86%, and 64%, respectively (see Figure 2).

IMPROVING STATISTICAL INFERENCES

Table 3

Baseline proportions of correct, incorrect, and “I don’t know” responses. Percentages reflect the proportion of respondents at *pretest* who obtained a correct response, an incorrect response, or selected the “I don’t know” option, for each of the individual items in the 14-item scale. Due to some missing responses, *ns* vary slightly across items and are accordingly reported. PV = *p*-value item, BF = Bayes factor item, CI = confidence interval item. Baseline accuracy rates (i.e. percentage of correct responses) are visually depicted in Figure 2.

Item	<i>n</i>	Correct (%) ^a	Incorrect (%) ^b	“I don’t know” (%) ^c
PV1	1,973	67.6	16.8	15.7
PV2	1,978	55.9	27.3	16.8
PV3	1,984	48.8	36.0	15.2
PV4	1,980	62.3	22.8	14.9
PV5	1,977	37.0	27.2	35.8
PV6	1,977	71.8	7.1	21.1
PV7	1,979	53.9	30.6	15.6
PV8	1,983	62.3	21.2	16.5
BF1	1,986	14.1	18.8	67.1
BF2	1,986	46.7	7.8	45.4
BF3	1,985	27.2	15.3	57.5
CI1	1,980	34.9	51.0	14.0
CI2	1,971	35.9	30.4	33.7
CI3	1,966	49.0	34.0	17.0

^a Bolded values reflect top 4 items with greatest proportion of correct responses: PV6, PV1, PV4 & PV8.

^b Bolded values reflect top 3 items with greatest proportion of incorrect responses: CI1, PV3, CI3.

^c Bolded values reflect top 3 items with greatest proportion of “I don’t know” responses: BF1, BF3, BF2.

IMPROVING STATISTICAL INFERENCES

Figure 2 – Baseline accuracy rates. Rates at pretest are plotted for each of the 14 individual items (**top**), as well as averaged across statistical concepts and *p*-value fallacies (**bottom**). Accuracy is computed in two manners: as the proportion of correct responses with “I don’t know” responses coded as incorrect (**striped bars**), and as the proportion of correct responses after “I don’t know” responses were omitted (**solid bars**). **Top:** Pretest items consist of 8 *p*-value (PV1 to PV8), 3 confidence interval (CI1 to CI3), and 3 Bayes factor (BF1 to BF3) items. **Bottom:** Statistical concepts are grouped into *p*-value fallacies (*inverse probability* fallacy; *replication* fallacy; *effect size* fallacy; *clinical or practical significance* fallacy; *correct interpretation*), confidence intervals, and Bayes factors.

IMPROVING STATISTICAL INFERENCES

Effects of scale versions. To date, the majority of studies investigating p -value misconceptions have only surveyed individuals' statements about significant p -values (i.e. given $p = .01$ or $.001$), with the exception of Lyu et al.'s recent (2020) study who also looked at interpretations of non-significant p -values for the *inverse probability* and *replication* fallacies. As we also administered a version of the scale dealing with statements about non-significant outcomes, (i.e. $p = .30$), for all four p -value fallacies, binary logistic regression analyses were run to compare baseline accuracy rates for both sets of the p -value items ("I don't know" responses coded as incorrect; Holm-Bonferroni correction applied for multiple testing). Odds ratios, computed for each p -value item, served to denote the relative likelihood of yielding a correct response when the item dealt with non-significant as compared to significant outcomes (for comparison of baseline rates and full overview of logistic regression results, see Figure F2 and Table T3 in Supplementary Information).

Analyses revealed that individuals systematically scored better on items about non-significant p -values, and were 2.19 times more likely ($p < .001$) to correctly recognize the *clinical or practical significance* fallacy (PV7) when interpreting a false statement about a non-significant outcome (i.e. 67% correct on "Obtaining a statistically non-significant result implies that the effect detected is unimportant") as compared to a significant outcome (i.e. 48% correct on "Obtaining a statistically significant result implies that the effect detected is important"). Additionally, *inverse probability* item PV1 was 1.76 times more likely ($p < .001$) to be correctly recognized as false when in the context of interpreting a non-significant p -value (i.e. 76% correct when given $p = .30$ ($\alpha = .05$), "You have absolutely proven the null hypothesis (that is, you have proven that there is no difference between the population means)") than in the context of interpreting a significant p -value (i.e. 64% correct when given $p = .001$ ($\alpha = .05$), "You have absolutely proven your alternative hypothesis (that is, you have proven that there is a difference between the population means)").

IMPROVING STATISTICAL INFERENCES

It should be noted, however, that observed misconception rates were also significantly different for two further items of which the item phrasings were identical in both versions. Specifically, individuals were respectively 1.44 and 1.61 times more likely ($p < .001$) to correctly interpret the *replication* fallacy item (PV5) and the correct definition item (PV8) when each of these items were considered in the context of a non-significant outcome. While we did not expect to find such differences, this pattern in the data might be explained by the existence of potential item dependencies contributing to consistency in item responses. In other words, perhaps the non-significant context forced learners to think more carefully about the correct interpretation of other items, such as PV1 and PV7, which in turn led to improved responses on other items administered within the same pretest assessment. On the other hand, observed differences in accuracy between scale versions might also be due to random sampling error, and simply reflect differences between subsamples despite random assignment (i.e. counterbalancing) across scale versions.

Remaining p -value items did not yield significant differences in accuracy between versions. Taken together, there is some evidence to suggest that the context in which a p -value is being interpreted (i.e. significant vs. non-significant outcome) may influence the propensity for individuals to fall prey to some common misconceptions, whereby the odds of endorsing a misinterpretation is greater when concerning significant outcomes. Such findings align with Lyu et al.'s (2020) observation that error rates, when averaged across all p -value items, were significantly lower in the context of a non-significant p -value interpretation. One, however, should be cautious when speculating about these results as they could reflect spurious differences.

Rates of Improvement

Rates of improvement were operationalized as any increase in performance, thereby constituting a shift from either a lack of knowledge (i.e. "I don't know" response) or lack in

IMPROVING STATISTICAL INFERENCES

understanding (i.e. incorrect response) to obtaining a correct response. Therefore, for analyses of improvement rates, we computed accuracy scores as the proportion of correct answers given all responses (i.e. “I don’t know” responses were coded as incorrect).

Overall learning. Linear mixed model (LMM) analysis¹ first investigated overall learning effects, across the subset of individuals who completed all six pop quizzes ($n = 162$, after response attrition and exclusions). LMM with random intercepts was used to regress quiz scores (i.e. composite total scores at *pretest*, *post-test 1*, and *post-test 2*) on *time* (dummy coded categorical predictor with 3 time points (*pretest*, *post-test 1*, *post-test 2*); *pretest* as reference category). Random intercepts variance ($\tilde{\tau}$) demonstrated substantial individual variability in *pretest* scores across participants (see Table 4). As effects of time on learning might be influenced by individual differences in course duration (i.e. number of days/weeks required to complete course; median course duration = 47.93 days or 6.85 weeks), the model also included an interaction term between *time* and *lag* (where here, *lag* (continuous, mean-centered) represented total time of completion from the *pretest* to *post-test 2*, measured in hours). As the model contained three effects of interest, i.e. effect of *time* (at *post-test 1* & at *post-test 2*) and interaction effect *time*lag*, significance levels for non-directional tests were corrected for multiple comparisons using a Bonferroni correction ($\alpha = .0167$). Post-hoc Tukey’s HSD analysis was included as a follow-up test, accounting for multiple comparisons (family-wise Type-1 error rate = 5%). LMM analysis (model $R_{\beta}^2 = .28^2$, ICC = .19) revealed a significant effect of *time*, with scores significantly increasing from a baseline mean score of 8.26 ($SD = 2.99$) to 11.13 ($SD = 2.13$) out of 14 at *post-test 1*, indicating a mean increase from *pretest* to *post-test 1* of 2.87 correct items, $p < .001$). At *post-test 2*, final mean scores reached 11.60

¹ LMM analyses run using the lmer function (lmerTest R package; Kuznetsova, Brockhoff, & Christensen, 2017).

² model R_{β}^2 was used, i.e. a standardized measure of multivariate association between the fixed predictors and the observed outcome (Edwards et al., 2008).

IMPROVING STATISTICAL INFERENCES

points ($SD = 1.85$), displaying an overall mean increase from *pretest* to *post-test 2* of 3.35 correct items ($p < .001$), which we consider a meaningful improvement. The interaction of *time* with *lag* was non-significant ($\hat{\beta} = 1.01\text{e-}04$, $p = .777$), indicating that the effect of total time of course completion (which varied between individuals) on improvement rates, was too small to be detected as significant.

Table 4

Overall learning effects ($n = 162$). Improvements in learning for the 14-item quiz scores, across the 3 assessment time points (*pretest*, *post-test 1*, *post-test 2*); *pretest* as reference category. LMM is summarized for fixed effects parameter estimates, as well as random effects (random intercepts variance, $\tilde{\tau}$).

Fixed effects						
Parameter	Estimate $\hat{\beta}$	SE	df	τ	p-value	semi-partial R^2^a
(Intercept)	8.26e+00	1.87e-01	339.30	44.28	< .001	
Time (<i>post-test 1</i>)	2.87e+00	1.95e-01	319.50	14.75	< .001	.25
Time (<i>post-test 2</i>)	3.35e+00	1.95e-01	319.50	17.20	< .001	.20
Time*Lag	1.01e-04	3.58e-04	401.80	0.28	.777	.00
Random effects						
Parameter	$\tilde{\tau}$					
ID (Intercept)	2.57					
Residual	3.07					

^a Suggested R-squared values of .02, .15, and .35 as small, medium, and large effect sizes (Cohen, 1988).

Quiz-level effects. Next, effects of learning were investigated at the quiz-level, entailing separate analyses for each of the subset of items detailed above (see Procedure above). This was to account simultaneously for differences in item administration (i.e. the fact that items occurred at different stages of the course progression), as well as difference in item completion (i.e. amount of time or ‘lag’ elapsing between any two measurement occasions, which could vary between

IMPROVING STATISTICAL INFERENCES

individuals). We applied once again a random intercepts LMM approach to assess the effect of *time*, while accounting for possible interactions between *time* and *lag* (mean-centered). Specifically, eight separate LMMs were run to assess improvements across each of the four subsets of items, first for immediate learning (i.e. four LMMs with pre-registered directional hypotheses), and then retained learning (i.e. another four LMMs with pre-registered non-directional hypotheses). For each set of analyses, Bonferroni correction for multiple comparisons was applied for 4 analyses with 2 effects of interest each (i.e. for immediate learning, $\alpha = .0125$; and for retained learning, $\alpha = .006$). Due to response attrition, improvements in learning for concepts covered later in the course are based on smaller *ns*, and reported accordingly. Effects for immediate learning, and retained learning, are now described in more detail below (for a summary of the eight LMMs, i.e. model R-squared, intra-class correlation, and random effects, see Table T2 in Supplementary Information).

Immediate learning. Results revealed significant improvements in immediate learning (*pretest* to *post-test 1*) across all sets of items: Specifically, main effect of *time* on quiz scores was most notable for least familiar items (i.e. the Bayes factors items which incurred the highest proportion of “I don’t know” responses) and more problematic items (i.e. the confidence intervals which incurred high proportions of incorrect responses) concepts, respectively subset 3 ($n = 478$, $\hat{\beta} = 1.02$, $p < .001$, $R^2_{\text{semi-partial}} = .27$) and subset 4 ($n = 271$, $\hat{\beta} = 0.93$, $p < .001$, $R^2_{\text{semi-partial}} = .23$), and less pronounced for the two *p*-value quizzes, i.e. subset 1 ($n = 712$, $\hat{\beta} = 0.72$, $p < .001$, $R^2_{\text{semi-partial}} = .07$) and subset 2 ($n = 325$, $\hat{\beta} = 0.38$, $p < .001$, $R^2_{\text{semi-partial}} = .07$). In general, baseline rates for Bayes factor and confidence interval items started off lower (resp. at 0.98 and 1.40 correct, out of max 3.0), with scores improving on average by approximately 1 point for both sets of items (i.e. increasing to 2.00 and 2.32 correct items at *post-test 1*, respectively). In contrast, *p*-value scores, which were initially higher at baseline (subset 1 $M = 3.12$ ($SD = 1.49$) out of 5; subset 2 $M = 2.34$ ($SD = 0.85$))

IMPROVING STATISTICAL INFERENCES

out of 3), incurred relatively smaller improvements, resulting in mean scores of 3.85 ($SD = 1.15$, subset 1) and 2.72 ($SD = 0.54$, subset 2) at *post-test 1*. Across all analyses, interactions between *lag* and *time* were non-significant ($ps \geq .564$), indicating that the amount of time that elapsed between pretest and the first post-test did not have a statistically significant impact on degree of improvement observed across individuals. Overall, when learning of statistical concepts was assessed directly after the relevant course module was taught, performance across sets of items tended to improve. That said, some subsets of items displayed less substantial improvements compared to others, such as the group of items assessing the *inverse probability* and *replication* fallacies (i.e. subset 1, which increased from 3.12 to 3.85, out of a max. score of 5). It is possible that the reason for these more minimal improvements reflects a general resistance to learning for some specific concepts, or at the very least, the need for more instructional support or explicit clarification for these concepts.

Retained learning. Regarding retained learning (*post-test 1* to *post-test 2*), LMMs demonstrated further positive effects of *time* on *p*-value items, which were statistically significant for subset 1 ($n = 207$, $\hat{\beta} = 0.38$, $p < .001$, $R^2_{\text{semi-partial}} = .03$), but non-significant for subset 2 ($n = 216$, $\hat{\beta} = 0.10$, $p = .007$, $R^2_{\text{semi-partial}} = .01$); both analyses yielding non-significant interaction effects of *lag*time* ($ps > .472$). In other words, though learning gains were relatively small, individuals continued to improve on *p*-value items until *post-test 2*, yielding final mean scores of 4.32 (out of 5, $SD = 1.00$, subset 1) and 2.84 (out of 3, $SD = 0.41$, subset 2). For Bayes factor items (subset 3, $n = 206$) and confidence interval items (subset 4, $n = 225$), neither effects of *time* (resp., $\hat{\beta} = 8.74\text{e-}02$ and $\hat{\beta} = -1.24\text{e-}01$) nor interaction effects with *lag* (resp., $\hat{\beta} = 2.68\text{e-}04$ and $\hat{\beta} = -5.14\text{e-}04$) were significant (all ps were larger than the Bonferroni-corrected alpha level of .006). In other words, detected increases and drops in scores from *post-test 1* to *post-test 2* were too small to be considered statistically significant, and thus are interpreted as representing retained learning, with final mean

IMPROVING STATISTICAL INFERENCE

scores of 2.10 (out of 3, $SD = 0.75$, Bayes factors) and 2.25 (out of 3, $SD = 0.80$, confidence intervals). Once again, time lag, i.e. the amount of time that elapsed between the two assessments, did not have a statistically significant impact on changes in scores observed across individuals for any of the subsets of items. Taken together, findings provide evidence that it is possible to improve misconception rates, across all three concepts (p -values, Bayes factors, confidence intervals), as a result of general statistics instruction. Given, however, that there remains room for even further improvement, it would be worth exploring how targeted training which explicitly clarifies these statistical misconceptions enhances learning.

Figure 3 – Rates of improvement. Improvements in learning across all 14 items from week 1 (PQ1) to week 8 (PQ6). Graphs demonstrate mean rate fluctuations in immediate learning (solid lines), i.e. from baseline (*pre-test*) to first post-test (*post-test 1*), and retained learning (dashed lines), from first to second post-test (i.e. *post-test 1* to *post-test 2*). **Left:** 8 *p*-value items (PV1 – PV8). **Right:** 3 Bayes factor items (BF1 – BF3), and 3 confidence interval items (CI1 – CI3).

IMPROVING STATISTICAL INFERENCES

Discussion

The current study provided evidence for the ability to improve misconception rates among online learners with respect to interpreting the meaning of p -values, Bayes factors, and confidence intervals. Mean scores, across the 14 items, improved from a baseline of 8.26 correct responses in week 1 to 11.13 (at *post-test 1*) and finally 11.60 (at *post-test 2*) at week 8. In other words, individuals overall made on average three to four fewer misconceptions as a result of participating in the 8-week MOOC. Specifically, all p -value fallacies (*inverse probability*, *replication*, *clinical or practical significance*, and *effect size*), as well as Bayes factor and confidence interval item subsets, incurred statistically significant improvements in immediate learning, that is when the follow-up test occurred immediately after the concept in question was taught. Concepts that were, on average, either lesser known (i.e. Bayes factors) or more problematic (i.e. confidence intervals), demonstrated unsurprisingly more pronounced improvements in immediate learning, as compared to p -values. Retained learning, as measured by a follow-up assessment in week 8, was observed across all statistical concepts, indicating neither significant increases nor drops in learning from *post-test 1* to *post-test 2*. Taken together, improvements resulting from the MOOC were not only successful in improving learners' ability to correctly interpret statements about statistical concepts, but these improvements were maintained until the final week of the course. What is worth noting of course is that observed increases in performance were certainly far from perfect, with some concepts, such as the group of items measuring then *inverse probability* and *replication* fallacies, yielding relatively minimal improvements. Should such trends reflect a higher level of difficulty in elucidating some of these more tricky statistical concepts, it would be worth devoting greater efforts in explicitly training away such misconceptions alongside traditional instruction.

IMPROVING STATISTICAL INFERENCES

An important contribution of our study was the inclusion of items that involved the interpretation of non-significant p -value items. Up to now, the majority of studies using scales to tap into rates of p -value misconceptions (i.e. Oakes, 1986; Haller & Krauss, 2002; Badenes-Ribera et al., 2015; Lyu, et al., 2018; Rinke & Schneider, 2018) have exclusively evaluated individuals' understanding of significant p -values, with the exception of Lyu et al. (2020) who also looked at interpretations of non-significant p -values for the *inverse probability* and *replication* fallacies. Considering the tendency for articles to report non-significant results as evidence of no effect or no relationship (Schatz et al., 2005; Fidler et al., 2006; Hoekstra et al., 2006; Bernardi et al., 2017), our findings support the need in also clarifying how to correctly make sense of and communicate the implications of null effects. Including these items also allowed us to investigate differences between baseline misconception rates of statistically significant ($p=.001$) versus non-significant ($p=.30$) inferences. Overall, individuals had a tendency to make more mistakes when interpreting p -value items in the context of a significant outcome (vs. non-significant outcome), with significant differences observed, for instance, in the case of the *clinical or practical significance* misconception. Such findings may be simply an artifact of item dependencies; alternatively, we might consider response patterns as a reflection of cognitive biases, such as confirmation bias in light of 'significance chasing' (Ware & Munafò, 2015). Due to "the tendency to emphasize and believe experiences which support one's views and to ignore or discredit those which do not" (Mahoney, 1977, p.161), individuals tasked to interpret a significant result may be more prone to prematurely endorsing a conclusion before questioning its limitations than the converse, i.e. when the outcome (e.g., non-significant p -value) conflicts with one's expectations. This ties in with the idea of inflated interpretations to which Bakan (1966) alluded: When it comes to interpreting research findings, over-stated generalizations, unwarranted extensions to different inferential levels,

IMPROVING STATISTICAL INFERENCES

and downplaying limitations, can lead to unrealistic conclusions and false endorsements of study outcomes, an abuse of “the realm of qualitative interpretation of quantitative effects” (Ioannidis, 2008, p. 643).

It is worth noting that the MOOC pop quizzes were completed voluntarily; neither were they graded nor required to complete in order to advance in or pass the course. As such, it would be of interest to replicate such a study within an educational setting where performance on the quizzes had specific learning goals or consequences, as improvements might be better when students have a stake in their score. In the same vein, as our design was not able to account for additional inter-individual differences between users (e.g., motivation, interest) or in strategies adopted (e.g., working individually vs. in a group, note-taking, etc.), it would be hard to speculate on whether learning effects reflect conservative versus optimistic estimates with respect to the average learner. On the other hand, as this flexibility in use is integral to the MOOC learning platform, work going forward might seek to investigate which characteristics of a MOOC – that would otherwise be absent in a traditional learning environment – may or may not facilitate the learning process. Specifically, with regard to the current line of research, it is reasonable to assume that there exists more than one viable avenue through which to clarify or alleviate statistical misunderstandings, with our study merely pointing to one possibility, namely the use of online training tools. Future work might attempt to capitalize on the growing body of available online resources (e.g, open source instructional materials and syntax, interactive visualizations and demonstrations like shinyapps, etc.) in order to enhance current methods used to teach statistical concepts, explain how to comprehensively make sense of data, and importantly draw and communicate meaningful and accurate statistical inferences.

IMPROVING STATISTICAL INFERENCES

Limitations

One potential limitation of the study was the observation that there may have been too systematic a relationship between the item phrasing and correct answers. Specifically, with the exception of the correct p -value interpretation (item PV8), all False items used positive phrasing (e.g., “You have found the probability of the null hypothesis being true ($p = .001$)” [PV2] or “Obtaining a statistically significant result implies that the effect detected is important” [PV7]), whereas all True statements were reverse items (e.g., “The probability that the results of the given study are replicable is not equal to $1-p$ ” [PV5] or “A Bayes Factor that provides strong evidence for the null model does not mean the null hypothesis is true” [BF2]). It is possible that if participants noticed this pattern, our results could simply be an artifact of this methodological confound. That said, if the majority of individuals successfully used this heuristic to drive their responses, we should realistically expect more consistency in response accuracies across all items (e.g., ceiling effects for both easier as well as more problematic items at post-tests). Because we did not observe such a pattern in our data, it is unlikely that this potential confound can fully explain our findings. Nevertheless, follow-up work should attempt to replicate findings while accounting for these sorts of methodological concerns. Another potential methodological confound is the possibility that improvements were caused by a testing effect, whereby mere exposure to items and their correct answers could have prompted individuals to score higher at subsequent post-tests. Similarly, however, if this occurred we should have arguably observed more pronounced increases in all items, from *post-test 1* to *post-test 2*, as well as greater improvements for items with a shorter time-lag between assessments. Neither of these patterns appeared in our data. This explanation is also less likely given that individuals were exposed to two scale versions containing non-identical items.

IMPROVING STATISTICAL INFERENCES

Another potential limitation worth discussing is the use of composite scores in the analyses assessing the improvements in accuracy rates over time, namely the LMM analyses. While our 14-item scale presumably tapped into some general construct such as statistics knowledge or savvy, we did not assume that it would be strictly unidimensional. In other words, subfactors representing different statistical concepts or types of misunderstandings would not be unexpected; and in fact, what we found was that a two-factor model (with subdimensions for frequentist vs. Bayesian concepts; see Supplementary Information) fit the data better than the one-factor model. In this way, our reliability coefficient for the 14-item scale (computed on the basis of assumed unidimensionality) yielded only an acceptable estimate ($\omega = .78$). Regarding the LMMs, the poorer the reliability of the items being summed across (to form overall scores), the more likely the beta weights in our regression model should be attenuated, thus leading to underestimation of path strengths. Regarding the 14-item scale reliability, the computed omega coefficient was as mentioned still acceptable; moreover, even though the two-factor model was favoured over the unidimensional model, the fit indices for the one-factor solution displayed reasonably acceptable fit. As such, concerns regarding the use of total score composites for the LMMs assessing overall learning were relatively less problematic. In contrast, the LMMs which used composite subscores obtained from each of the four subsets of items, i.e. the LMMs used to assess improvements in immediate and retained learning, were far more likely to be subject to attenuated estimates. This was unsurprising as each of the subsets consisted of only a handful of items (3 to 5 total) and could combine different types of misconceptions. Notably, items like those of subset 2 (i.e. *effect size fallacy*, *clinical or practical significance fallacy*, and correct *p*-value fallacy) were only incidentally grouped together on account of being relevant to different aspects of the week 4 course module on effect sizes; thus, while inter-related in some capacity, the three items as a subscale yielded a reliability estimate that

IMPROVING STATISTICAL INFERENCES

was fairly poor ($\omega = .58$). Taken together, because the current study was not purely experimental, but rather embedded within a real online course, some experimental design decisions, such as staggering the *post-test 1* items across weeks 1 to 5, served as both a design necessity and limitation, and introduced a trade-off between what would be more logical (from a pedagogical perspective) yet more arbitrary or questionable (from a statistical perspective). In turn, improvement rates may reflect more conservative estimates.

Finally, a key limitation of our work concerns the generalizability of the results. Namely, due to the nature of our scale and experimental design, our findings can only speak to individuals' ability to recognize misconceptions at the level of isolated theoretical (and often abstract) statements. As such, our results cannot be used as the basis to infer whether such learning will transfer into practice (such as interpreting real research results). We feel that the current work is nevertheless an important first step in tackling the improvement of statistical inferences especially in face of claims about the "persistency and deepness of the misconceptions" (Sotos, Vanhoof, Van den Noortgate, & Onghena, 2007) and their perceived resistance to change. Our work demonstrates that not only is it possible to systematically improve statistical misconception rates across individuals, but that the need for improved methods of instruction and clarification is not exclusive to frequentist concepts, like *p*-values and confidence intervals, but also includes alternative measures like Bayes factors. Whether statistical reforms tend toward the improvement of a more widely practiced method (i.e. NHST), versus its abandonment in place of an alternative approach (e.g., Bayesian statistics), challenges surrounding the process of drawing correct statistical inferences will invariably exist no matter the inferential approach used. To increase the ecological validity of studies investigating misconception rates, work going forward should also assess barriers and facilitators to statistical inference in a context that closer simulates the conditions of drawing conclusions from real data.

IMPROVING STATISTICAL INFERENCES

Finally, future work would benefit from experimental studies that directly manipulate the degree of instructional support learners' receive, in order to better understand how supplementary tools and training can reduce the probability of errors individuals make.

Conclusion

The current study provides empirical grounds to contradict the assumption that statistical misconceptions are impervious to correction. Moreover, we believe our findings provide strong grounds to a) advocate for increased efforts in exploring, developing and refining targeted training techniques, specifically geared toward the clarification of statistical misconceptions, and b) recommend that these instructional scaffolds complement the instruction of core statistical concepts.

IMPROVING STATISTICAL INFERENCES

References

- Amrhein, V., Greenland, S., & McShane, B. (2019). Scientists rise up against statistical significance.
- Badenes-Ribera, L., Frías-Navarro, D., Monterde-i-Bort, H., & Pascual-Soler, M. (2015). Interpretation of the p value: A national survey study in academic psychologists from Spain. *Psicothema*, 27(3), 290-295.
- Bakan, D. (1966). The test of significance in psychological research. *Psychological bulletin*, 66(6), 423.
- Belia, S., Fidler, F., Williams, J., & Cumming, G. (2005). Researchers misunderstand confidence intervals and standard error bars. *Psychological methods*, 10(4), 389.
- Berkson, J. (1938). Some difficulties of interpretation encountered in the application of the chi-square test. *Journal of the American Statistical Association*, 33(203), 526-536.
- Bernardi, F., Chakhaia, L., & Leopold, L. (2017). ‘Sing me a song with social significance’: The (mis) use of statistical significance testing in European sociological research. *European Sociological Review*, 33(1), 1-15.
- Carver, R. (1978). The case against statistical significance testing. *Harvard Educational Review*, 48(3), 378-399.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*, 2nd edition. New York: Erlbaum.
- Cohen, J. (1990). Things I have learned (so far). *American psychologist*, 45(12), 1304.

IMPROVING STATISTICAL INFERENCES

- Craig, J. R., Eison, C. L., & Metze, L. P. (1976). Significance tests and their interpretation: An example utilizing published research and ω2. *Bulletin of the Psychonomic Society*, 7(3), 280-282.
- Cumming, G., Fidler, F., Leonard, M., Kalinowski, P., Christiansen, A., Kleinig, A., ... & Wilson, S. (2007). Statistical reform in psychology: Is anything changing?. *Psychological science*, 18(3), 230-232.
- Cumming, G., & Maillardet, R. (2006). Confidence intervals and replication: Where will the next mean fall?. *Psychological methods*, 11(3), 217.
- Cumming, G., Williams, J., & Fidler, F. (2004). Replication and researchers' understanding of confidence intervals and standard error bars. *Understanding statistics*, 3(4), 299-311.
- Edwards, L. J., Muller, K. E., Wolfinger, R. D., Qaqish, B. F., & Schabenberger, O. (2008). An R2 statistic for fixed effects in the linear mixed model. *Statistics in medicine*, 27(29), 6137-6157.
- Falk, R., & Greenbaum, C. W. (1995). Significance tests die hard: The amazing persistence of a probabilistic misconception. *Theory & Psychology*, 5(1), 75-98.
- Feng, W., Tang, J., & Liu, T. X. (2019, July). Understanding dropouts in MOOCs. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 33, pp. 517-524).
- Fidler, F. (2005). *From statistical significance to effect estimation: Statistical reform in psychology, medicine and ecology* (Doctoral dissertation).

IMPROVING STATISTICAL INFERENCES

- Fidler, F. (2006). Should psychology abandon p-values and teach CIs instead? Evidence-based reforms in statistics education. In *Proceedings of the seventh international conference on teaching statistics*. International Association for Statistical Education
- Fidler, F., Thomason, N., Cumming, G., Finch, S., & Leeman, J. (2004). Editors can lead researchers to confidence intervals, but can't make them think: Statistical reform lessons from medicine. *Psychological Science*, 15(2), 119-126.
- George, D., & Mallery, P. (2010). SPSS for Windows step by step. A simple study guide and reference (10. Bask1).
- Gliner, J. A., Vaske, J. J., & Morgan, G. A. (2001). Null hypothesis significance testing: Effect size matters. *Human Dimensions of Wildlife*, 6(4), 291-301.
- Greenland, S., Senn, S. J., Rothman, K. J., Carlin, J. B., Poole, C., Goodman, S. N., & Altman, D. G. (2016). Statistical tests, P values, confidence intervals, and power: a guide to misinterpretations. *European journal of epidemiology*, 31(4), 337-350.
- Haller, H., & Krauss, S. (2002). Misinterpretations of significance: A problem students share with their teachers. *Methods of Psychological Research*, 7(1), 1-20.
- Hoekstra, R., Finch, S., Kiers, H. A., & Johnson, A. (2006). Probability as certainty: Dichotomous thinking and the misuse of p values. *Psychonomic Bulletin & Review*, 13(6), 1033-1037.
- Hoekstra, R., Morey, R. D., Rouder, J. N., & Wagenmakers, E. J. (2014). Robust misinterpretation of confidence intervals. *Psychonomic bulletin & review*, 21(5), 1157-1164.
- Ioannidis, J. P. (2008). Interpretation of tests of heterogeneity and bias in meta-analysis. *Journal of evaluation in clinical practice*, 14(5), 951-957.

IMPROVING STATISTICAL INFERENCES

Kalinowski, P., Fidler, F., & Cumming, G. (2008). Overcoming the inverse probability fallacy: A comparison of two teaching interventions. *Methodology*, 4(4), 152-158.

Kirk, R. E. (1996). Practical significance: A concept whose time has come. *Educational and psychological measurement*, 56(5), 746-759.

Kuznetsova A, Brockhoff PB, Christensen RHB (2017). “lmerTest Package: Tests in Linear Mixed Effects Models.” *_Journal of Statistical Software_*, *82*(13), 1-26. doi: 10.18637/jss.v082.i13 (URL: <https://doi.org/10.18637/jss.v082.i13>).

Lakens, D. (2019, April 9). The practical alternative to the p-value is the correctly used p-value. <https://doi.org/10.31234/osf.io/shm8v>

Lecoutre, M. P., Poitevineau, J., & Lecoutre, B. (2003). Even statisticians are not immune to misinterpretations of Null Hypothesis Significance Tests. *International Journal of Psychology*, 38(1), 37-45.

Leys, C., Ley, C., Klein, O., Bernard, P., & Licata, L. (2013). Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology*, 49(4), 764-766.

Lyu, Z., Peng, K., & Hu, C. P. (2018). P-value, Confidence Intervals and Statistical Inference: A New Dataset of Misinterpretation. *Frontiers in Psychology*, 9, 868.

Lyu, X. K., Xu, Y., Zhao, X. F., Zuo, X. N., & Hu, C. P. (2020). Beyond psychology: prevalence of p value and confidence interval misinterpretation across different fields. *Journal of Pacific Rim Psychology*, 14.

IMPROVING STATISTICAL INFERENCES

- Mahoney, M. J. (1977). Publication prejudices: An experimental study of confirmatory bias in the peer review system. *Cognitive therapy and research*, 1(2), 161-175.
- Morey, R. D., Hoekstra, R., Rouder, J. N., Lee, M. D., & Wagenmakers, E. J. (2016). The fallacy of placing confidence in confidence intervals. *Psychonomic bulletin & review*, 23(1), 103-123.
- Nickerson, R. S. (2000). Null hypothesis significance testing: a review of an old and continuing controversy. *Psychological methods*, 5(2), 241.
- Oakes, M. W. (1986). *Statistical inference*. Epidemiology Resources.
- Rinke, E. M., & Schneider, F. M. (2018). Probabilistic misconceptions are pervasive among communication researchers. <https://doi.org/10.31235/osf.io/h8zbe>
- Rozeboom, W. W. (1997). Good science is abductive, not hypothetico-deductive. *What if there were no significance tests*, 335-391.
- Schatz, P., Jay, K. A., McComb, J., & McLaughlin, J. R. (2005). Misuse of statistical tests in Archives of Clinical Neuropsychology publications. *Archives of Clinical Neuropsychology*, 20(8), 1053-1059.
- Shafer, G. (2019). On the nineteenth-century origins of significance testing and p-hacking. *Available at SSRN 3461417*.
- Shaver, J. P. (1993). What statistical significance testing is, and what it is not. *The Journal of Experimental Education*, 61(4), 293-316.

IMPROVING STATISTICAL INFERENCE

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological science*, 22(11), 1359-1366.

Sotos, A. E. C., Vanhoof, S., Van den Noortgate, W., & Onghena, P. (2007). Students' misconceptions of statistical inference: A review of the empirical evidence from research on statistics education. *Educational Research Review*, 2(2), 98-113.

Wagenmakers, E. J., Wetzels, R., Borsboom, D., & Van Der Maas, H. L. (2011). Why psychologists must change the way they analyze their data: the case of psi: comment on Bem (2011).

Ware, J. J., & Munafò, M. R. (2015). Significance chasing in research practice: causes, consequences and possible solutions. *Addiction*, 110(1), 4-8.

Wasserstein, R. L., & Lazar, N. A. (2016). The ASA's statement on p-values: context, process, and purpose. *The American Statistician*, 70(2), 129-133.

IMPROVING STATISTICAL INFERENCES

APPENDIX Q1
Scale Version 1

Task: Please read through the following statements, and mark each as True or False. Note that several or none of the statements may be correct. *[Options: True / False / I don't know]*

Given: Let's suppose that a research article indicates a value of $p = .001$ in the results section ($\alpha = .05$).

Subset	Item		Fallacy	Correct A
1	PV1	You have absolutely proven your alternative hypothesis (that is, you have proven that there is a difference between the population means).	Inverse probability	F
	PV2	You have found the probability of the null hypothesis being true ($p = .001$).	Inverse probability	F
	PV3	The null hypothesis has been shown to be false.	Inverse probability	F
	PV4	The p -value gives the probability of obtaining a significant result whenever a given experiment is replicated.	Replication	F
	PV5	The probability that the results of the given study are replicable is not equal to $1-p$.	Replication	T
2	PV6	The value $p = .001$ does not directly confirm that the effect size was large.	Effect size	T
	PV7	Obtaining a statistically significant result implies that the effect detected is important.	Clinical or practical significance	F
	PV8	The p -value of a statistical test is the probability of the observed result or a more extreme result, assuming the null hypothesis is true.	Correct interpretation	T
3	BF1	When a Bayesian t-test yields a $BF = 0.1$, it is ten times more likely that there is no effect than that there is an effect.	N/A	F
	BF2	A Bayes Factor that provides strong evidence for the null model does not mean the null hypothesis is true.	N/A	T
	BF3	A Bayes Factor close to 1 (inconclusive evidence) means that the effect size is small.	N/A	F
4	CI1	The specific 95% confidence interval observed in a study has a 95% chance of containing the true effect size.	N/A	F
	CI2	If two 95% confidence intervals around the means overlap, then the difference between the two estimates is necessarily non-significant ($\alpha = .05$).	N/A	F
	CI3	An observed 95% confidence interval does not predict that 95% of the estimates from future studies will fall inside the observed interval.	N/A	T

IMPROVING STATISTICAL INFERENCES

APPENDIX Q2
Scale Version 2

Task: Please read through the following statements, and mark each as True or False. Note that several or none of the statements may be correct. *[Options: True / False / I don't know]*

Given: Let's suppose that a research article indicates a value of $p = .30$ in the results section ($\alpha = .05$).

Subset	Item		Fallacy	Correct A
1	PV1	You have absolutely proven the null hypothesis (that is, you have proven that there is no difference between the population means).	Inverse probability	F
	PV2	You have found the probability of the null hypothesis being true ($p = .30$).	Inverse probability	F
	PV3	The alternative hypothesis has been shown to be false.	Inverse probability	F
	PV4	The p -value gives the probability of obtaining a significant result whenever a given experiment is replicated.	Replication	F
	PV5	The probability that the results of the given study are replicable is not equal to $1-p$.	Replication	T
2	PV6	The value $p = .30$ does not directly confirm that the effect size was small.	Effect size	T
	PV7	Obtaining a statistically non-significant result implies that the effect detected is unimportant.	Clinical or practical significance	F
	PV8	The p -value of a statistical test is the probability of the observed result or a more extreme result, assuming the null hypothesis is true.	Correct interpretation	T
3	BF1	When a Bayesian t-test yields a $BF = 10$, it is ten times more likely that there is an effect than that there is no effect.	N/A	F
	BF2	A Bayes Factor that provides strong evidence for the alternative model does not mean the alternative hypothesis is true.	N/A	T
	BF3	A Bayes Factor close to 1 (inconclusive evidence) means that the effect size is small.	N/A	F
4	CI1	The specific 95% confidence interval observed in a study has a 5% chance of not containing the true effect size.	N/A	F
	CI2	To draw the conclusion that the difference between the two estimates is non-significant ($\alpha = .05$), it is necessary that the two 95% confidence intervals around the means do not overlap. ¹	N/A	F
	CI3	An observed 95% confidence interval does not predict that 95% of the estimates from future studies will fall inside the observed interval.	N/A	T

¹Note: Typo observed after implementation of the scale. While item correctly signals a misinterpretation about CIs, it is not identical to the misunderstanding measured in Version 1. To correspond, should be corrected to "To draw the conclusion that the difference between the two estimates is non-significant ($\alpha = .05$), it is necessary that the two 95% confidence intervals around the means overlap."