Researcher-Centered Design of Statistics: Why Bayesian Statistics Better Fit the Culture and Incentives of HCI

Matthew Kay
Computer Science &
Engineering | dub,
University of Washington
mjskay@uw.edu

Gregory L. Nelson
Computer Science &
Engineering | dub,
University of Washington
glnelson@uw.edu

Eric B. Hekler
School of Nutrition &
Health Promotion
Arizona State University
ehekler@asu.edu

ABSTRACT

A core tradition of HCI lies in the experimental evaluation of the effects of techniques and interfaces to determine if they are useful for achieving their purpose. However, our individual analyses tend to stand alone, and study results rarely accrue in more precise estimates via meta-analysis: in a literature search, we found only 56 meta-analyses in HCI in the ACM Digital Library, 3 of which were published at CHI (often called the top HCI venue). Yet meta-analysis is the gold standard for demonstrating robust quantitative knowledge. We treat this as a user-centered design problem: the failure to accrue quantitative knowledge is not the users' (i.e. researchers') failure, but a failure to consider those users' needs when designing statistical practice. Using simulation, we compare hypothetical publication worlds following existing frequentist against Bayesian practice. We show that Bayesian analysis yields more precise effects with each new study, facilitating knowledge accrual without traditional meta-analyses. Bayesian practices also allow more principled conclusions from small- n studies of novel techniques. These advantages make Bayesian practices a likely better fit for the culture and incentives of the field. Instead of admonishing ourselves to spend resources on larger studies, we propose using tools that more appropriately analyze small studies and encourage knowledge accrual from one study to the next. We also believe Bayesian methods can be adopted from the bottom up without the need for new incentives for replication or meta-analysis. These techniques offer the potential for a more user- (i.e. researcher-) centered approach to statistical analysis in HCI.

Author Keywords

Replication; meta-analysis; Bayesian statistics; small studies; effect size; estimation.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous;

INTRODUCTION

A core focus of the HCI community is on the development of novel ideas and technology artifacts. The focus on novelty is valuable: it establishes a tradition that challenges assumptions about the design of technical systems and often results in insights that translate into more useful, usable, and enjoyable technologies. We, as a community, also value accurately understanding and precisely characterizing the effects of technology, which can be at odds with the first goal. Our community has limited resources to meet this latter goal as we have historically devoted more resources towards novelty and innovation over accurate and precise estimates of the utility of our systems.

Greater emphasis on novelty has led to concerns about the reliability of knowledge accrued in our field. To support quantitative knowledge accrual, we have adopted particular statistical tools, such as frequentist null hypothesis significance testing¹, and quantitative standards, such as $p < .05$, that define what constitutes sufficient evidence for researchers to support their claims. It is well known in the statistics community that results from individual studies—especially with the small sample sizes typical in our community—regularly fail to reliably estimate true effects [17]. To gain more reliable estimates of effects, multiple studies can be aggregated using frequentist meta-analytic techniques, which combine the results from multiple papers to obtain more precise effect size estimates (e.g., the difference between two conditions). However, our community rarely conducts meta-analyses. Paralleling so-called *replication crises* in psychology and medicine, movements such as RepliCHI [26] have called for an increased focus on replication and meta-analysis to effectively accrue quantitative knowledge about the utility and generalizability of our technologies. Others emphasize deeper changes to statistical practice, such as reducing the focus on p -values in favor of effect size estimates and confidence intervals (the *New Statistics* [3]; or in HCI, Kaptein & Robertson [18] or Dragicevic [5]), or the abandonment of frequentist null hypothesis

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHI'16, May 07 - 12, 2016, San Jose, CA, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3362-7/16/05 \$15.00

DOI: <http://dx.doi.org/10.1145/2858036.2858465>

¹ We will use frequentist and NHST interchangeably.

sis significance testing (NHST)² altogether for Bayesian analyses. Similar conversations about the merits of Bayesian analyses are also occurring in psychology [4,16,20].

All of these suggestions, apart from a Bayesian approach, logically extend NHST, the dominant statistical approach used in HCI and related fields such as psychology. NHST is a statistical approach that turns research into a binary question: can we reject or fail to reject a *null hypothesis* (i.e., that there is no effect)? A common use of NHST in HCI is to compare a novel system to a control system and, if a p -value is below the customary $p < .05$ target, then the new system is deemed better, as this p -value suggests that the observed difference is not likely due to chance. In other words, we reject the null hypothesis of no difference.

However, getting a p -value less than .05 can still happen even when there is no true meaningful difference. Replication and meta-analysis allow us to reduce this error in NHST (for example, the probability of falsely rejecting the null hypothesis) by combining the results of many studies of the same phenomenon. However, this requires at least one additional study, the meta-analysis, which necessitates new top-down incentives for conducting and publishing meta-analysis in HCI. By contrast, Bayesian analysis incorporates prior knowledge from other studies of the same and similar phenomena into a paper's quantitative analysis. A series of papers analyzing novel contributions can plausibly accrue knowledge and bypass the need for publishing separate meta-analyses. This allows increased precision of knowledge from the bottom up, within the existing publishing incentives of the field.

We consider the choice of statistical tools to be a user-centered design problem, with researchers as the users. Insisting that we *should* conduct meta-analysis amounts to blaming the users instead of the tools. Instead, we propose changing the tools—from NHST statistics to Bayesian statistics—in order to make quantitative accrual of knowledge easier (and maybe even preferable) within the existing publishing incentives of HCI. It is not researchers, but the statistical tools they have been given, which currently prevents this. Bayesian statistics are more user-centered statistics.

In the rest of this paper, we briefly compare Bayesian and NHST approaches to statistics and give some background on replication and meta-analysis. We then examine a subset of the HCI literature in the ACM digital library to assess the current state of meta-analyses in the field and establish that current incentives do not encourage meta-analyses, especially at the most prestigious venues. We then run several simulated experiments, representing hypothetical experi-

ments run for separate publications, using a realistic effect size drawn from an existing meta-analysis. We contrast two hypothetical publication worlds: one in which the simulated experiments were each analyzed in a traditional (NHST) manner (as would occur now), and one in which they were analyzed using Bayesian techniques. We demonstrate:

1. **The current state of quantitative knowledge accrual in HCI is poor.** Through an examination of publications in the ACM Digital Library, we demonstrate that little meta-analysis is conducted in the community.
2. **Bayesian analysis provides more precise estimates of previously-studied conditions in each successive study.** The NHST approach only increases the precision of effect sizes if the new study has a larger sample or when a meta-analysis is conducted. In contrast, the Bayesian approach uses prior knowledge to increase the precision of effect sizes for known conditions in each successive study, without requiring a meta-analysis (which is unlikely to be done in HCI).
3. **Bayesian analysis allows more precise comparison of novel conditions against known conditions.** By giving more precise effect size estimates of previously-studied conditions, Bayesian analysis increases the precision of estimated differences between existing and novel conditions, making it suited to novel HCI work.
4. **Bayesian analysis facilitates quantitative knowledge accrual within HCI's existing publishing incentives.** Unlike frequentist analysis, Bayesian analysis can accrue knowledge within individual studies without new top-down incentives for publishing meta-analyses, shifting knowledge accrual into original papers. Benefits like improved precision may help incentivize this adoption.
5. **Bayesian analysis draws more reasonable conclusions from small- n studies.** In small-sample studies of novel techniques, priors can be used to shrink unreasonably large effect sizes. Priors also let us take advantage of known conditions to improve estimates of novel ones. This makes Bayesian analysis particularly attractive to design and engineering researchers running small studies on novel technology.
6. **Bayesian analyses help shift the conversation from “Does it work?” to “How strong is the effect?”, “How confident are we in this estimate?”, and “Should we care?”** While NHST can incorporate effect size estimates and confidence intervals, ultimately, the use of a p -value (a fundamental feature of NHST) translates all questions into binary answers: can we reject or fail to reject the null hypothesis? Bayesian estimates emphasize the probability of an effect given prior knowledge, thus shifting the conversation from “does it work?” to “how strong is the effect?” and “how confident are we?” Frequentist statistics can technically answer these questions but under-emphasizes them. Bayesian statistics better emphasizes the ultimate questions of our work: as practitioners, do

² While there exist Bayesian formulations of NHST based on Bayes factors, we believe they share some problems with frequentist NHST, such as a focus on binary testing rather than precision of estimation and cost/benefit analysis; thus we do not consider them here.

we care enough about these results to adopt new designs, or as researchers, do we care enough to study this more?

BACKGROUND AND MOTIVATION

Interpretation of Bayesian versus frequentist statistics

Interpretations of frequentist statistics are a common source of errors amongst HCI researchers and others. The focus on p -values/significance testing requires users to learn how to interpret the cognitively demanding conceptual double-negative of a p -value, instead of interpreting results as evidence for a hypothesis—a valid interpretation within a Bayesian framework [20]. This interpretation problem has resulted in simplifying the results of NHST tests into the belief that the p -value answers the question, “Does it work?” Unfortunately, this is an inaccurate interpretation and lies at the heart of the argument to shift interpretation towards effect sizes and confidence intervals [3,6] (though 95% confidence intervals are the inverse of $p < .05$, thus still succumbing to the interpretation problem [14]). Effect size estimates and confidence in those estimates is the essential information sought, but NHST logic unintentionally relegates this information as secondary to the p -value.

By contrast, Bayesian analysis provides formal approaches to quantifying our existing beliefs (for example, as a probability distribution over the expected difference in the means of some variable between two conditions), and then updating those beliefs based on new experimental evidence. This gives results expressed as probabilistic evidence for or against hypotheses, emphasizing effect size estimates and confidence in the estimates. This information supports the decisions researchers and practitioners are making with the data: should I incorporate this technology into my practice or, if confidence is low, are the results promising enough to continue to study it? Conceptually, Bayesian statistics are feasibly more user-centered as they emphasize the information needed to support the decisions researchers want to make without the need to interpret double-negative logic.

Replication and meta-analysis in HCI

The statistical tools researchers customarily use in HCI do not help them effectively accrue knowledge from one study to the next, even when the variations in design of novel systems are informed by previous work. The classic strategy for knowledge accrual of a series of NHST studies is literature reviews conducted often in the related work section of a CHI paper that implicitly use the *vote-counting* method of knowledge accumulation [13]. In this method, the number of significant and non-significant findings are counted up to infer if an effect is true or not (e.g., three studies found a significant effect, four did not, therefore this strategy is likely not effective). There are many problems with this approach, particularly when a field utilizes small samples to estimate statistical significance, as many of these significant differences are likely due to chance [17].

A step towards better knowledge accrual is via the use of meta-analyses, where the focus is not on the statistical sig-

nificance of any single study, but instead on combining the results from many studies to estimate the *effect size* (e.g., this system designed for encouraging exercise results in 1,000 more steps per day compared to control) and the confidence in that effect (i.e., that 1,000 step increase could feasibly be as low as 100 steps or as high as 1,900 steps). This strategy relies somewhat on increasing the incentives for replication in the literature, an approach currently advanced by RepliCHI [26]. While encouraging more standalone replication studies and meta-analyses is useful for knowledge accrual, we argue that it has difficulty fitting into HCI culture and the incentives for publishing novel findings. As we will show, few meta-analyses are currently conducted in the community, supporting our intuition.

Within a Bayesian approach, prior beliefs can be derived from previous work, allowing knowledge accrual from study to study without requiring a separate meta-analysis. To derive priors in HCI, we can capitalize on the fact that partial replication is common to the field in the form of the comparison of a new technique against the state-of-the-art. As we will show, incorporating prior quantitative results into new analyses using a Bayesian framework is straightforward in these cases, allowing us to accrue quantitative knowledge without the need for top-down incentives for meta-analysis. We will also discuss how to use prior work to set prior expectations on the size of an effect even when not conducting a partial replication. In contrast to traditional meta-analysis, Bayesian analysis allows the effect sizes in successive studies to be estimated more precisely in each study. This could feasibly fit into the publishing incentives for HCI: knowledge accrues with each individual, novel study (easily published at CHI), making it unnecessary to publish standalone meta-analyses (less publishable at CHI).

HCI RESEARCHERS CONDUCT FEW META-ANALYSES

To assess the current state of quantitative knowledge aggregation in HCI, we conducted a review of meta-analyses accessible through the ACM Digital Library, as many of the most prominent HCI publication venues are archived there (e.g. CHI, CSCW, UIST, UbiComp, TOCHI). We searched the extended library for the terms *meta-analysis*, *meta-analyses*, *metaanalysis*, or *metaanalyses* in the abstract or title fields on August 17, 2015, yielding 509 unique results. We examined abstracts and eliminated 151 domain-specific statistical methods and techniques, mostly in biology and machine learning. We examined the full-text of the remaining papers. We found 40 dissertations, which we discarded since their results may have been published in other venues. We found 56 HCI (or HCI-related) papers with quantitative meta-analyses, defined as modeling effect sizes or using traditional meta-analysis based on the results of multiple studies found from a literature search with inclusion criteria. Only 3 were published at the venues above [15,25,28]. This low number prompted us to search the DL full text for “meta-analysis” for the top venues, yielding 159 results. The top 3 results were the meta-analyses we had already found, and we did not find any others after reviewing the

abstracts (and full text as needed) from this additional search. Most meta-analyses were in other journals and communities, from management information systems and HICSS to specialized venues (e.g. ICMI '06: Proceedings of the 8th international conference on multimodal interfaces).

Our search suggests that meta-analyses are not being rewarded within the community. The CHI conference is regarded by many as the top publication venue for work in human-computer interaction. Given the paucity of meta-analysis at CHI, it seems clear that the incentives do not currently exist for such work to be published there (or prioritized in the field overall). However, within the frequentist paradigm, a meta-analysis is the gold standard for quantitative knowledge accrual, representing the best estimates we can make. We must either build new incentives for meta-analysis, or find another way to accrue knowledge within the existing incentives. We believe that Bayesian analysis may be that other way.

CONTRASTING FREQUENTIST AND BAYESIAN KNOWLEDGE ACCRUAL USING SIMULATED EXPERIMENTS

By way of explaining the differences between frequentist meta-analysis and a Bayesian incremental approach to knowledge accrual, in this section we provide an example of these approaches applied in two different hypothetical worlds. Specifically, we examine a series of 4 simulated, hypothetical experiments on the effects of progress indicators on completion rates of online surveys.

Domain: Varying progress indicators in online surveys

We chose this domain because it will be familiar to the HCI audience (as many researchers in our field make use of online surveys), and because a meta-analysis has previously been conducted in this domain by Villar *et al.* [24], providing us with realistic effect sizes to use in our simulations.

That meta-analysis looked at experiments comparing the effects of different types of *progress indicators* on survey completion rates. A progress indicator is any type of textual or graphical display communicating how much of the survey has been completed so far (“10%”, a graphical progress bar, etc.). Progress indicators can be distinguished by the relationship between the true progress and the displayed progress. A *constant* indicator communicates the true progress. Progress in a *fast-to-slow* indicator starts fast, telling the participant they have made more progress than they actually have near the beginning of the survey, then slows down later. By contrast, a *slow-to-fast* indicator starts slow, then speeds up near the end of the survey.

In their meta-analysis, Villar *et al.* [24], found that using a *slow-to-fast* progress indicator decreased the probability that a person would complete the survey. They found an effect size (as a log odds ratio³) of ~ -0.45 .⁴ A log odds ratio

³ A log odds ratio of 0 indicates no difference between conditions. The log odds ratio is the log of the ratio of the odds of someone completing the survey in one condition com-

pared to another condition. It is regularly used in comparing probabilities between two conditions because (unlike, say, differences of proportions), it is unbounded, which simplifies analysis. It is related to logistic regression in that coefficients of a logistic regression can be interpreted as log odds ratios.

Simulation Method

To compare Bayesian and frequentist approaches, we will simulate 100 hypothetical “worlds” in which we know the true effect of different progress bar types on completion rates, and then run the same series of 4 experiments in each world. Each experiment could represent an experiment run by different authors, thus representing the error that would be present from a single study, the current primary mechanism for estimating effect sizes. We will conduct analyses on each world as if 1) all authors take a frequentist approach or 2) all authors take a Bayesian approach.

For the purposes of our simulations, we will consider the true effect of a *slow-to-fast* progress bar on the log-odds of the completion of a survey to be -0.45 , as suggested by the meta-analysis of Villar *et al.* [24]. We will also surmise a similarly-sized effect of *fast-to-slow* progress indicators, in the opposite direction, of 0.45 .⁵ This approach supports examination of knowledge accrual in both worlds and allows comparison to the known fact (because we define it in the simulation) that there is an effect size of -0.45 (*slow-to-fast*) and 0.45 (*fast-to-slow*) compared to no indicator.

In each world, we simulate the results of 4 experiments:⁶

- **Experiments 1-3** have between-subjects designs comparing a *fast-to-slow* progress indicator against a control condition of no indicator.
- **Experiment 4** also compares a *fast-to-slow* progress indicator against a control condition, but adds an additional *slow-to-fast* indicator. We can think of this experiment as representing one of the common ways that partial replication happens in the HCI community: through comparison to previous state-of-the-art results. Perhaps some authors,

pared to another condition. It is regularly used in comparing probabilities between two conditions because (unlike, say, differences of proportions), it is unbounded, which simplifies analysis. It is related to logistic regression in that coefficients of a logistic regression can be interpreted as log odds ratios.

⁴ While their results use probability of drop-out, we use probability of completion.

⁵ It is worth noting that this effect may be larger than the true effect in the real world, since Villar *et al.* [24] found the *fast-to-slow* effect is likely closer to .3; but for our purposes it remains a realistic effect size based on the *slow-to-fast* results, and simplifies the example by mirroring *slow-to-fast*.

⁶ See <https://github.com/mjskay/bayes-for-chi> for simulation data and code (DOI: 10.6084/m9.figshare.2061669)

Figure 1. Forest plots of effects from the frequentist (A) and Bayesian (B) analyses applied to one of our simulated worlds with 100 participants per condition.

having seen the success of *fast-to-slow* indicators, wished to know how the opposite type of indicator might perform (or perhaps conducted this experiment as part of work to establish a more complete theory explaining *why* we see these particular results).

For simplicity of exposition, we assume the same experimental design in each case: a between-subjects design with 100 participants per condition (thus, 200 participants in experiments 1-3 and 300 in experiment 4). This is similar to the number of participants in the studies in Villar *et al.*'s meta-analysis, and is a reasonable number to expect to respond to an online survey. The between-subjects design is necessary primarily because it is difficult to ask someone to take the same survey twice and observe their drop-out rates.

Frequentist analysis

In the frequentist analysis, we conduct a logistic regression in each experiment to model the probability of *completion* based on the *progress indicator condition*: *control* (no indicator), *fast-to-slow*, and *slow-to-fast* (experiment 4 only). In addition, after all four experiments are analyzed, we conduct a meta-analysis on the log-odds ratios for the effect of the *fast-to-slow* progress indicator, as in Villar *et al.* [24]. This yields a final, more precise estimate of the effect of that indicator based on the preceding four experiments.

Bayesian analysis

In the Bayesian analysis of each world, we also conduct a logistic regression in each experiment to model the probability of *completion* based on *progress indicator*. However, we do not conduct a final meta-analysis. Instead, starting with experiment 2, we build upon the previous results by using the posterior distribution (i.e., estimates from the previous study) of the estimated effect of the *fast-to-slow* progress indicator in experiment i as the prior for that effect in

experiment $i + 1$. This could happen, for example, if the author of experiment $i + 1$ had read the previous paper, and therefore was able to use the posterior estimate from that paper in their analysis (previous study's posterior becomes the next study's prior). This results in incremental knowledge accumulation without a meta-analysis.

In experiment 4, we must also place a prior on the new *slow-to-fast* indicator. We use a Cauchy⁷ distribution centered at 0 (no effect) with a scale equal to the furthest point in the 95% credibility interval⁸ of the estimated effect of *fast-to-slow* in experiment 3. This prior is weakly-informed: it expresses a belief that *fast-to-slow* might reasonably have about twice the effect (positively or negatively) that *slow-to-fast* does compared to control condition. While it is beyond the scope of this paper to discuss the various strategies for setting priors, this is a core topic in any book on Bayesian analysis (see e.g. [21]).

Results

In a single world

Before contrasting the Bayesian and frequentist results across all simulated worlds, we will first walk through the

⁷ The Cauchy distribution is similar to the Normal distribution, but with fatter tails. Gelman recommends it for use as a weakly-informed prior because the fatter tails express less certainty in the location of the effect [9].

⁸ Roughly, a credibility interval is the Bayesian analog to a confidence interval. Unlike a confidence interval, however, a proper credibility interval *is* an expression of the probability of the location of a parameter (the confidence interval is not, despite its common misinterpretation [14]; again a reason frequentist statistics are not user-friendly).

results of one simulation. The results of the frequentist analysis are shown above in Figure 1A and the results of the Bayesian analysis in Figure 1B. Each figure shows a forest plot with 95% CIs (*confidence intervals* in the frequentist case; *credibility intervals* in the Bayesian case). In the frequentist case, a 95% confidence interval that does not overlap 0 is equivalent to a p value of less than .05, suggesting that the null hypothesis (i.e., no effect) can be rejected with 95% confidence. The dashed vertical lines indicate the true effect sizes from which the data was simulated.

In the frequentist analysis, we have a promising first result in experiment 1. This is followed by two borderline results in experiments 2 and 3. Looking strictly at p -values, experiment 4 fails to replicate the result of experiment 1, though it does find some evidence of a difference between *fast-to-slow* and *slow-to-fast* progress indicators. Finally, the meta-analysis is able to combine the previous estimates into a more precise and accurate estimate of the true effect—assuming it is conducted and published.

Note that, because all of these experiments are run using the same number of participants, the confidence intervals are all approximately the same width; the only ways to increase our precision (i.e., decrease CI width) in the frequentist world are to increase the power of our experiment/analysis (for example by increasing our sample size, using a within-subjects design, or including covariates that explain some of the variation in the response) or by conducting meta-analysis. This limitation is not particularly helpful to the authors of experiments 1-4, since they may not have the resources to recruit more participants.

In addition, the small variations in intervals from experiments 1-3 represent vastly different conclusions if we reduce the results to vote counting of null hypothesis tests: experiments 1 and 3 reject the null ($p < .05$); experiment 2 does not. This highlights the problem with reducing estimation to a binary choice (“effect” or “no effect”): these estimates are all similar, but the decision to reject (or not) the null hypothesis hinges on whether the 95% confidence interval overlaps 0 in that particular study, thus resulting in false conclusions as per the vote-counting method discussed earlier (current HCI standard practice).

In the Bayesian analysis, the result of the first experiment is virtually identical to the frequentist world (we used a weakly-informed Cauchy(0, 2.5) prior for logistic regression parameters recommended by Gelman *et al.* [9]). However, in contrast to the frequentist world, in each subsequent experiment our estimate of the effect size becomes more precise. The authors of experiments 2 and 3 make a stronger contribution to the field by building on the results of prior work, rather than borderline failed replications. In experiment 4, the estimated effect of the fast-to-slow indicator is similar to that of the frequentist meta-analysis, *one paper early*. **Bayesian analysis helps us learn faster and with fewer publications.**

Figure 2. Results of the frequentist (A) and Bayesian (B) analyses of all simulated worlds, 100 participants per condition.

Besides the benefit of getting quantitative knowledge accrual into the literature without requiring publication of meta-analysis, this also has additional benefits for the authors of experiment 4: note that, even though they are testing a new technique that they don’t have explicit priors for

(the *slow-to-fast* indicator), the prior knowledge of the effect of the *fast-to-slow* indicator helps them estimate the effect of the novel technique more precisely. This is because the more precise estimate of *fast-to-slow* also helps makes the estimate of the difference between *fast-to-slow* and *slow-to-fast* a little more precise. In other words, more precise estimates of techniques we've seen lead to more precise estimates of comparisons to new techniques, which even makes estimates of those new techniques a little more precise. **Bayesian analysis helps us apply old knowledge to novel questions.**

In many worlds

We now step up to consider the effects of the two analysis approaches in all 100 simulated worlds. Each point in Figure 2A represents the mean estimated effect from the frequentist analysis in one simulated world. Figure 2B shows the mean estimated effects from the Bayesian analyses.

We can see that the pattern observed in our single example world holds true across simulations: the estimated effect becomes narrower (more precise) with each experiment in the Bayesian analysis, and the final estimate for *fast-to-slow* resembles the frequentist meta-analysis, one publication early. In addition, the estimates for *slow-to-fast* are more precise in the Bayesian analysis of experiment 4 due to the use of prior knowledge, even though we have never seen that condition before. This increased precision (and accuracy) is reflected in the root-mean-squared error of those estimates compared to their true effects in experiment 4:

	Frequentist	Bayesian
fast-to-slow – control	0.27	0.17
slow-to-fast – control	0.27	0.20
slow-to-fast – fast-to-slow	0.26	0.22

Table 1. Root mean-squared error of estimates in experiment 4 with 100 participants per condition.

Note that the estimated effects in the frequentist analysis all have approximately the same error, reflective of the power of the experiment. The Bayesian approach gives us estimates with less error by building knowledge as we go.

BAYESIAN ANALYSIS OF SMALL SAMPLES

Due to limited resources and an emphasis on novelty, HCI studies are often conducted with fewer participants than a traditional power analysis would suggest is prudent [10,18]. With a frequentist analysis, this increases the probability of what Gelman calls a *magnitude error* [8]: because the confidence intervals are so wide, the only effects that reach significance are those that overestimate the effect size.

That said, we believe that there are many reasons why small-*n* studies are conducted in HCI, including limited resources and the emphasis on novelty discussed earlier as fundamental for contributions to the field. Thus, we ask: can Bayesian analysis help make better use of our limited

A. Fast-to-slow – control in experiment 1, one simulation

B. Fast-to-slow – control in experiment 1, all simulations

Figure 3. The effects of shrinkage towards 0 in experiment 1 with 20 participants, from a single simulation (A) and from all simulations (B).

resources while still acknowledging the importance of novelty? To assess this, we repeated our simulations with 20 participants per condition instead of 100 and examined how the use of priors could be used to reduce false conclusions both from single studies and across studies.

Results

At only 20 participants per condition, in the frequentist world a large effect size (a log odds ratio of nearly 1.5) is required to reject the null hypothesis—3 times the known (because we defined it) effect size. To put this in context, imagine a survey whose completion rate with no progress indicator would be 50%. The correct log odds ratio of 0.45 would imply that switching to a fast-to-slow indicator would raise the completion rate to 61%, a modest effect. However, a log odds ratio of 1.5—the size needed to reject the null hypothesis—would imply that a fast-to-slow indicator would yield an 82% completion rate on this survey. This gross overestimate exemplifies a magnitude error [8]. Even though we haven't (in our hypothetical world) studied this particular effect, we do have prior knowledge about anticipated effect sizes based on standard small, medium, and large effects in humans. In the Bayesian analysis, we encode this knowledge in the prior we set on the effect; we adopt a weakly-informed prior from Gelman [9] (we address prior-setting in more detail in the discussion).

Bayesian analysis effectively weighs how strong our prior knowledge is against how much evidence we have: with larger samples (e.g., our 200-participant experiments) the evidence more strongly influences the estimate. With only 20 participants, the prior has more influence. This enables

us to quantify our skepticism and mitigate the risk of making a magnitude error in small samples. In the frequentist world we might intuitively dismiss large effect sizes in small studies as unreasonable, but not have a formal way to specify this skepticism. In the Bayesian world we can consistently and quantitatively apply this intuition by encoding it in a prior and using the prior to shift unreasonably large effects towards zero. This is called **shrinkage**.

We can see the effects of shrinkage by comparing an inflated estimate (by chance, a common occurrence with small sample studies) from experiment 1 with its corresponding Bayesian estimate (Figure 3). Compared to the frequentist estimate, the Bayesian estimate is shrunk towards 0. The resulting posterior is still quite diffuse: we haven't learned all that much from the small study. That said, we did learn *something* from our 20 participants: The Bayesian analysis weights what we knew before (specifically, that small effect sizes are generally anticipated if nothing more is known) against how much evidence we have (in this case, observations of a small number of participants).

Overall, Bayesian analyses allow for information to be gleaned in proportion to the strength of available evidence. In contrast, estimates from individual small sample studies in the frequentist world are only corrected via meta-analysis. Prior to a meta-analysis, however, NHST can create more confusion in effect size estimates if only significant results are published (*publication bias*). In the above example, this would mean that the effect size estimate would need to be nearly 3 times the actual effect within a single small sample study. If, by chance, this large effect were found (as was the case with our hypothetical scenario), it would create an inflated effect size estimate in the literature that is unlikely to be corrected due to the lack meta-analyses in the HCI community.

In many worlds

We can see the effects of shrinkage in the first experiment when we look at all of the 20-participant simulations (Figure 4): the most extreme estimates in experiment 1 are moved slightly towards 0. This has the effect of reducing the overall error in experiment 1 by discounting unreasonably large estimates that occur due to chance:

	Frequentist	Bayesian
fast-to-slow – control	0.61	0.56

Table 2. Root mean-squared error of estimate in experiment 1 with 20 participants per condition, reflecting the effect of Bayesian shrinkage on discounting unreasonably large effects in small-*n* studies.

Compared to the frequentist world, the Bayesian world also offers more precise and accurate estimates across studies (as represented in the root mean square error). By the time we reach experiment 4, the difference in error is dramatic: the estimate for *fast-to-slow* has nearly half the error in the Bayesian world (.36 versus .66), and we again get better

Figure 4. Results of the frequentist (A) and Bayesian (B) analyses of all simulated worlds, 20 participants per condition.

estimates of the novel condition, *slow-to-fast* (see Table 3 below). In contrast, effect size estimates are only corrected with meta-analysis in the frequentist world, thus the correction is unlikely in HCI because of a lack of meta-analyses.

	Frequentist	Bayesian
fast-to-slow – control	0.66	0.36
slow-to-fast – control	0.68	0.51
slow-to-fast – fast-to-slow	0.83	0.60

Table 3. Root mean-squared error of estimates in experiment 4 with 20 participants per condition

DISCUSSION

Bayesian analysis increases the value of small-*n* studies of novel work

The traditional solution to the problems associated with low-power studies (and one HCI researchers are often admonished to adopt) is to spend resources recruiting more participants. In other words, the frequentist solution to low power is *not to run low-powered studies*.

However, researchers developing complex new systems or interaction techniques have the expertise, time, and resources for that type of work; spending their limited resources on running larger studies may be a poor allocation of work across the research community as an accurate effect size might not be their first priority. These researchers already (in our view, rightly) protest that they are asked to run *pro forma* evaluations when their primary contributions are in engineering or design (see e.g., Greenberg and Buxton [12]); telling them not only to run evaluations but to recruit more participants amounts to blaming the users.

Statistics can be a tool for communication and collaboration. Bayesian analysis can feasibly support collaborations that take advantage of specialization. We see researchers that produce novel systems and interaction work as having a symbiotic relationship with others who have the resources and expertise for larger quantitative work, but perhaps not the expertise for novel engineering. The latter researchers might find a novel technique in the literature, adapt it to some domain based on users' needs (or collaborate directly with its inventors, as often happens in novel health sensing work), and evaluate it more extensively. In this context, the goal of small, early studies then becomes to demonstrate face validity of a technique and provide a rough first estimate of its effectiveness, not to find a (likely over-estimated in terms of magnitude) significant difference. For this, Bayesian analysis helps draw reasonable conclusions from small-*n* studies. It provides a more nuanced and accurate tool for evaluating contributions and combining them given the varying skills and resources of researchers.

Part of the goal of this paper is to release novel work in HCI from the chains of meaningless *p*-values from small-*n* studies. We believe that small, early evaluations of novel work are still valuable, but that their output should be a probability distribution of expected effect size, whether or not it overlaps 0. “No effect” as determined via *p*-values should not be a barrier to publication of novel design work when we know that any effect that is found in a small study

is suspect. Instead, the novel work should be (and already often is) judged on the merits of design and engineering, not a *pro forma* small-*n* evaluation. Bayesian analysis provides a richer conceptual understanding and role for these initial evaluations and helps to quantify information (i.e., effect sizes and confidence in those effect sizes) to support the questions implied by the community: should I incorporate this novel tool into my practice or, if not confident enough, is further research warranted?

Bayesian analysis fits into how statistical practice is shaped in HCI

The HCI community is large and multi-disciplinary; therefore, we believe that statistical practice in HCI is best shifted in a bottom-up fashion. For example, Wobbrock *et al.* [27] at CHI 2011 introduced a nonparametric analysis technique to the community—the aligned rank transform (ART)—applicable to various forms of data, including Likert scales. Since then, this approach has been cited 170 times.⁹ This adoption did not require new top-down incentives for improved analysis, but spread study-to-study and researcher-to-researcher.

We believe that Bayesian analysis can be adopted gradually in individual studies, sidestepping the difficulty of shifting an entire multi-faceted field from the top down. Statistical practice in scientific fields tends towards a model of mentorship and of drawing upon approaches found in prior work—e.g., as other papers begin adopting techniques like ART, readers of those papers will use similar techniques when conducting their own analyses in follow-up work. This is how Bayesian analysis can spread: when readers see it used in a paper they wish to build upon, the analysis offers a direct way to do that, teaching them by example [1]. Such a paper also provides priors for the next researcher. In this manner such analyses can spread in the community, slowly building and accumulating knowledge.

Bayesian analysis is increasingly accessible

Even 15 years ago, Bayesian analysis was arguably impractical for most researchers due to a lack of tools and computational power. However, tools for building and running Bayesian models are now widespread, and have mature support in languages already used for data analysis, such as R and Python. These tools include modeling languages like JAGS [23] and Stan [2] (both with R packages, and Stan includes a Python interface), and Python-specific libraries like emcee [7] and PyMC [22]. In addition, literature aimed at practicing researchers has made Bayesian modelling accessible; we particularly recommend Kruschke's *Doing Bayesian Data Analysis* [21] (which includes a table of common frequentist analyses and their Bayesian equiva-

⁹ According to Google Scholar, accessed 2016-01-07: https://scholar.google.com/scholar?cites=16254127723353600671&as_sdt=5.48&sciodt=0.48&hl=en

lents), as well as his proposed BEST test,¹⁰ a robust Bayesian alternative to the *t*-test [19]. Other accessible articles have also been written about practical concerns in Bayesian analysis, including discussions of how to choose priors [9,11]. Still, statistical tools, whether frequentist or Bayesian, can exhibit high barriers to entry or silently fail, providing a poor interface. We believe there is a fruitful area of work in designing better tools and interfaces for statistical methods.

Challenges and opportunities in setting priors

A particular difficulty in Bayesian analysis, in contrast to frequentist analysis, lies in setting priors for novel conditions. Careful suggestions on prior-setting is an essential topic that requires further work but preliminary recommendations can be made. Within HCI, we propose the use of informed skeptical priors for novel conditions. Such priors shrink unreasonably large estimates, particularly in small-*n* studies, making them well-suited for the HCI community.

For example, for generalized linear models (which encompass many models in HCI, including linear and logistic regression), reasonable priors on coefficients can often be specified as a Student-*t* distribution. This symmetric, bell-shaped family is parameterized by degrees of freedom (*df*), location, and scale. It includes the Normal (as $df \rightarrow \infty$) and Cauchy ($df = 1$) distributions. A lower *df* means fatter tails, encoding a higher probability that we may see a more extreme effect size. To express skepticism, we set the prior's location to 0, the equivalent of "no effect" here. The scale parameter expresses the range of reasonable effect sizes around 0. For example, one can set a small scale around 0 to express skepticism in large effects. This will require more robust evidence to shift these beliefs, thus making them less susceptible to inflated effect size estimates common to frequentist analysis of small samples [8]. We saw the benefit of this strategy in the reduction in error in our small-*n* simulations.

Skeptical priors acknowledge the known problem of inflated effect sizes in small samples, which makes them a good fit for HCI. However, there are many strategies for setting priors. For example, another approach, which would mimic a frequentist approach, would be to use uninformed priors (e.g., a flat prior on regression coefficients), yielding results similar to frequentist confidence intervals.

Overall, Bayesian analysis puts a greater onus on the analyst to describe their priors, but this cost brings several benefits, particularly mathematical specification of skepticism. Even for novel conditions, effect sizes of previously studied phenomena (e.g., *fast-to-slow*) can be used to establish informed skeptical priors of an unknown, but related phenomena (e.g., *slow-to-fast*). That said, prior-setting can be daunting for those new to Bayesian analysis. There is fruit-

ful work to be done in guiding researchers in this process, perhaps via interactive systems for prior specification.

Practical impact of research through cost/benefit analysis

Finally, we wish to address another common thread of discussion in the HCI community, a perhaps more existential one: how can we have practical effects on real-world deployed systems? How can practitioners derive value from results in the HCI literature? We believe that the language of statistical significance is not the language of practitioners or business; cost/benefit analysis is. The results of a Bayesian analysis can easily be incorporated into cost/benefit analysis: given the probability distribution of an estimated effect, we can simply apply a cost function to it.

For example, imagine a market research company that wishes to evaluate the cost/benefit of switching from an existing survey tool that does not have a *fast-to-slow* progress indicator to one that does. This would incur some costs for converting the survey into a new format. It would also have an estimated benefit in that the company could recruit fewer participants to reach a desired sample size, in proportion to the expected increase in completion rate. This company could take the probability distribution of estimated completion rate in both cases (*whether or not* the difference has passed the statistical significance filter) and use it to derive a probability distribution of expected cost relative to the probable benefit, and then decide a course of action to minimize cost. This simplifies the translation of research into real-world use, and gives a way to put practical effect sizes in context.

CONCLUSION

Bayesian analysis allows us to learn more quickly by building on previous results. It also fits more effectively into the publication incentives of HCI than approaches to improving knowledge accrual within the NHST framework, such as meta-analysis. At the same time, it is compatible with calls for more replication (RepliCHI), and allows us to make stronger claims about novel work through comparison to well-studied conditions. This, combined with a shift to an emphasis on probable effect sizes instead of simply "significant" differences between conditions, will help free design and engineering researchers from the shackles of meaningless *p*-values in small-*n* studies, while also allowing the field to make better use of the results of such studies. In short, Bayesian statistics are more user-centered statistics.

ACKNOWLEDGEMENTS

Matthew Kay was supported by the Intel Science and Technology Center for Pervasive Computing (ISTC-PC). Gregory Nelson was supported by the National Science Foundation (NSF) Graduate Research Fellowship Program under Grant No. DGE-1256082 and under NSF grant SCH-1344613. Eric Hekler's work was supported, in part, by a grant from the Robert Wood Johnson Foundation (PI: Hekler, 71995).

¹⁰ Somewhat glibly, BEST stands for *Bayesian estimation supersedes the t-test*

REFERENCES

1. Joel Brandt, Mira Dontcheva, Marcos Weskamp, and Scott R. Klemmer. 2010. Example-centric programming. *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10*, ACM Press, 513. Retrieved January 7, 2016 from <http://dl.acm.org/citation.cfm?id=1753326.1753402>
2. Bob Carpenter, Andrew Gelman, Matt Hoffman, et al. 2015. Stan: A Probabilistic Programming Language. *Journal of Statistical Software* In press.
3. Geoff Cumming. 2014. The new statistics: why and how. *Psychological science* 25, 1: 7–29. Retrieved July 9, 2014 from <http://pss.sagepub.com/content/early/2013/11/07/0956797613504966.abstract>
4. Zoltan Dienes. 2011. Bayesian Versus Orthodox Statistics: Which Side Are You On? *Perspectives on psychological science : a journal of the Association for Psychological Science* 6, 3: 274–90. Retrieved January 7, 2016 from <http://pps.sagepub.com/content/6/3/274.abstract>
5. Pierre Dragicevic. 2016. Fair Statistical Communication in HCI. In *Modern Statistical Methods for HCI*, Judy Robertson and Maurits Kaptein (eds.). Springer.
6. Pierre Dragicevic. 2016. From Statistical Significance to Fair Statistical Communication in HCI. In *Modern Statistical Methods for HCI*, Judy Robertson and Maurits Kaptein (eds.). Springer.
7. Daniel Foreman-Mackey, David W. Hogg, Dustin Lang, and Jonathan Goodman. 2013. emcee : The MCMC Hammer. *Publications of the Astronomical Society of the Pacific* 125, 925: 306–312. Retrieved April 21, 2015 from <http://arxiv.org/abs/1202.3665>
8. A. Gelman and J. Carlin. 2014. Beyond Power Calculations: Assessing Type S (Sign) and Type M (Magnitude) Errors. *Perspectives on Psychological Science* 9, 6: 641–651. Retrieved November 17, 2014 from <http://pps.sagepub.com/content/9/6/641.short>
9. Andrew Gelman, Aleks Jakulin, Maria Grazia Pittau, and Yu-Sung Su. 2008. A Weakly Informative Default Prior Distribution for Logistic and Other Regression Models. *The Annals of Applied Statistics* 2, 4: 1360–1383.
10. Andrew Gelman and David Weakliem. 2009. Of Beauty, Sex and Power. *American Scientist* 97, 4: 310–316.
11. Andrew Gelman. 2006. Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis* 1, 3: 515–533. <http://doi.org/10.1214/06-BA117A>
12. Saul Greenberg and Bill Buxton. 2008. Usability evaluation considered harmful (some of the time). *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08*, ACM Press, 111. Retrieved September 25, 2015 from <http://dl.acm.org/citation.cfm?id=1357054.1357074>
13. Larry V. Hedges and Ingram Olkin. 1980. Vote-counting methods in research synthesis. *Psychological Bulletin* 88, 2: 359–369.
14. Rink Hoekstra, Richard D. Morey, Jeffrey N. Rouder, and Eric-Jan Wagenmakers. 2014. Robust misinterpretation of confidence intervals. *Psychonomic Bulletin & Review* 21, 5: 1157–1164. Retrieved July 1, 2015 from <http://www.ncbi.nlm.nih.gov/pubmed/24420726>
15. Kasper Hornbæk and Effie Lai-Chong Law. 2007. Meta-analysis of correlations among usability measures. *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '07*: 617. <http://doi.org/10.1145/1240624.1240722>
16. George S. Howard, Scott E. Maxwell, and Kevin J. Fleming. The proof of the pudding: An illustration of the relative strengths of null hypothesis, meta-analysis, and Bayesian analysis.
17. John P A Ioannidis. 2005. Why most published research findings are false. *PLoS medicine* 2, 8: e124. Retrieved July 9, 2014 from <http://journals.plos.org/plosmedicine/article?id=10.1371/journal.pmed.0020124>
18. Maurits Kaptein and Judy Robertson. 2012. Rethinking statistical analysis methods for CHI. *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12*, ACM Press, 1105. Retrieved September 25, 2015 from <http://dl.acm.org/citation.cfm?id=2207676.2208557>
19. John K Kruschke. 2013. Bayesian estimation supersedes the t test. *Journal of experimental psychology. General* 142, 2: 573–603. Retrieved September 25, 2015 from <http://www.ncbi.nlm.nih.gov/pubmed/22774788>
20. John K. Kruschke. 2010. Bayesian data analysis. *Wiley Interdisciplinary Reviews: Cognitive Science* 1, 5: 658–676. <http://doi.org/10.1002/wcs.72>
21. John K. Kruschke. 2011. *Doing Bayesian Data Analysis*. Elsevier Inc.
22. Anand Patil, David Huard, and Christopher J Fonnesbeck. 2010. PyMC: Bayesian Stochastic Modelling in Python. *Journal of statistical software* 35, 4: 1–81. Retrieved December 2, 2015 from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3097064&tool=pmcentrez&rendertype=abstract>
23. Martyn Plummer. 2003. JAGS: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling. *Proceedings of the 3rd International*

- Workshop on Distributed Statistical Computing (DSC 2003)*. <http://doi.org/10.1.1.13.3406>
24. Ana Villar, Mario Callegaro, and Yongwei Yang. 2013. Where Am I? A Meta-Analysis of Experiments on the Effects of Progress Indicators for Web Surveys. *Social Science Computer Review* 00, 0: 1–19. <http://doi.org/10.1177/0894439313497468>
25. Suzanne Weisband and S Kiesler. 1996. Self Disclosure on Computer Forms : Meta-Analysis and Implications. *ACM Digital Library CHI*, 96: 3–10. <http://doi.org/10.1145/238386.238387>
26. Max Wilson, Wendy E. Mackay, Ed Chi, Michael Bernstein, and Dan Russell. RepliCHI - CHI should be replicating and validating results more: discuss. *CHI EA '11: Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems*. <http://doi.org/10.1145/1979742.1979491>
27. Jacob O Wobbrock, Leah Findlater, Darren Gergle, and James J Higgins. 2011. The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only ANOVA Procedures. *CHI '11*: 143–146.
28. Nick Yee, Nick Yee, Jeremy N Bailenson, Jeremy N Bailenson, Kathryn Rickertsen, and Kathryn Rickertsen. 2007. A meta-analysis of the impact of the inclusion and realism of human-like faces on user experiences in interfaces. *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '07*: 1. <http://doi.org/10.1145/1240624.1240626>