Paper #3 of 3

Abstract (1,000 word max, not including figures or tables)

This presentation describes an example of an analysis of the impact of an education intervention on a student outcome measure that utilizes a Bayesian approach that incorporates results from prior studies. Our goal for the description of the methodology is that an analyst with a moderate amount of experience conducting statistical analyses and programming equations with statistical software such as SAS, R, Stata, or SPSS could reproduce the Bayesian analysis results using the data from the presentation and accompanying paper, and therefore fully understand the approach. To date there have been few examples of application of Bayesian approaches that incorporate prior information in impact analyses of education interventions. For example, in a recent review of the use of Bayesian statistics in educational research (König & van de Schoot, 2018), the authors identified only one study that utilized results of prior studies in their Bayesian analysis, and that study was not an impact evaluation that attempted to estimate an intervention impact on student outcome measures. Our presentation will make a contribution to the SREE community by providing guidance and an easily replicable example with real data for how to actually conduct a Bayesian impact analysis that incorporates information from a current study and relevant prior studies.

The Bayesian impact analysis approach will be contrasted to the classical impact analysis approach which uses only data from the current study and typically produces an estimate of the size of the intervention impact and a p-value or confidence interval for the impact estimate. The p-value or confidence interval conveys the likelihood of obtaining an impact estimate as or more extreme than the one obtained if the true impact of the intervention were zero. A meta-analysis may use the results from the current study and relevant prior studies of the same or similar interventions to produce an estimate of the average impact across the collection of studies. The meta-analysis would also typically produce a p-value or confidence interval that conveys the likelihood of obtaining an average impact estimate as or more extreme than the one obtained if the true impact of the intervention(s) were zero. The Bayesian analysis utilizes the results from the current study and relevant prior studies to produce results that can be stated in the form of following example: "*There is a very high probability that text-message-based advising has a positive impact on immediate college enrollment",* and *"There is a very low probability that text-message-based advising increases the proportion of students who enroll in college immediately after high school by more than two or three percentage points."*

Because there are multiple decision-points in the Bayesian analysis and different choices can be made at each point, it is important to determine if the results found through this study's chosen approach are sensitive to particular decisions. The example analysis in this presentation a primary analysis and several sensitivity analyses. The sensitivity analyses involved an alternative method for identifying relevant prior studies, alternative assumptions about the effects of publication bias, and alternative analysis statistical methods to estimate the "posterior distribution". Differences across these dimensions between the primary approach and the sensitivity analyses are summarized in Exhibit 1.

A central aspect of the Bayesian analysis is that it incorporates prior knowledge into the current analysis. Deke & Finucane (2019) recommend using only prior knowledge that comes from prior studies and not using expert judgement or non-informative prior information such as assuming an arbitrary mean an infinite variation, as the basis of prior knowledge. But they also recommend assessing the sensitivity of results to the choice of which prior studies to use. We used two approaches to identifying relevant prior studies. The first utilized a broad-based literature search strategy, while the second utilized only studies in the What Works Clearinghouse (WWC) database of reviewed studies. The first was chosen as primary because it identified studies of interventions that were the most closely aligned with the intervention implemented in the current study.

Because publications serve as a primary source for prior knowledge, it is important to have an understanding of what studies may or may not be published and to be explicit about any assumptions made. Publication bias occurs when the results of a study (e.g., a finding of a large effect or no effect) influence whether the study is published or the results are reported. Because there is no way to know if publication bias has occurred, Deke and Finucane (2019) recommend conducting sensitivity analyses with several different assumptions regarding the presence of publication bias. Per recommendations we present the primary and sensitivity analysis approaches summarized in Exhibit 1. The details of how to apply the assumptions used in each approach will be presented so that attendees can replicate results.

The Bayesian analysis is described as being comprised of three broad steps:

- Step 1: Calculate a meta-analytic mean and standard error of impact estimates from relevant prior studies. The results are called "the priors."
- Step 2: Use the priors from Step 1 in combination with the results from the current study to estimate the mean and standard deviation of impacts in the "posterior distribution." A primary and a sensitivity analysis approach are provided for this step.
- Step 3: Using the mean and standard deviation of the posterior distribution, calculate the probability that the intervention impact is greater than a particular number (e.g., greater than zero or greater than three percentage points).

The steps are described in sufficient detail such each can be replicated using the data provided in the paper.

Conclusions from the sensitivity analyses include the following: The results were not sensitive to how the relevant prior studies were identified; the results were quite sensitive to one of the assumptions about publication bias; and the results were not at all sensitive to the method for estimating the posterior mean.

Exhibit 1. Details about Key Dimensions of the the Primary Approach and Each Sensitivity Analysis

Analysis	Prior Studies	Publication Bias	Analysis Method Used to Estimate Posterior Distribution
Primary Approach	PS-1. Relevant prior studies identified by literature review	PB-1. No adjustment for potential publication bias	AM-1. Closed-form equations used to estimate posterior distribution
Sensitivity Analysis 1	PS-2. Relevant prior studies identified in What Works Clearinghouse database of reviewed studies	PB-1. No adjustment for potential publication bias	AM-1. Closed-form equations used to estimate posterior distribution
Sensitivity Analysis 2	PS-1. Relevant prior studies identified by literature review	PB-2. Adjustment made for potential publication bias – assume published effects are inflated by a factor of two	AM-1. Closed-form equations used to estimate posterior distribution
Sensitivity Analysis 3	PS-1. Relevant prior studies identified by literature review	PB-3. Adjustment made for potential publication bias – assume that all published effects are zero	AM-1. Closed-form equations used to estimate posterior distribution
Sensitivity Analysis 4	PS-1. Relevant prior studies identified by literature review	PB-1. No adjustment for potential publication bias	AM-2. Iterative procedure used to estimate posterior distribution

References

Deke J. & Finucane M., (2019) Moving Beyond Statistical Significance: The BASIE (BAyeSian Interpretation of Estimates) Framework for Interpreting Findings from Impact Evaluations. OPRE Report #2019-35. January.

Christoph König & Rens van de Schoot (2018) Bayesian statistics in educational research: a look at the current state of affairs, Educational Review, 70:4, 486-509, DOI: 10.1080/00131911.2017.1350636