Addressing publication bias in meta-analysis: Empirical findings from community-augmented
meta-analyses of infant language development

Sho Tsuji^{1,2}, Alejandrina Cristia², Michael C. Frank³, & Christina Bergmann⁴

¹The University of Tokyo

²Ecole Normale Supérieure, PSL University, EHESS, CNRS

³Stanford University

⁴Max Planck Institute for Psycholinguistics

Author Note

Sho Tsuji, International Research Center for Neurointelligence, Institutes for Advanced Studies, The University of Tokyo. This research was supported by grants from the Berkeley Initiative for Transparency in the Social Sciences, a program of the Center for Effective Global Action (CEGA), with support from the Laura and John Arnold Foundation. The authors were further supported by the H2020 European Research Council [Marie Skłodowska-Curie grant No. 659553], the Agence Nationale de la Recherche [ANR-14-CE30-0003 MechELex, ANR-17-EURE-0017]. Correspondence concerning this article should be addressed to Sho Tsuji, International Research Center for Neurointelligence, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033 Japan, Contact: tsujish@gmail.com

Abstract

Meta-analyses are an indispensable research synthesis tool for characterizing bodies of literature and advancing theories. One important open question concerns the inclusion of unpublished data into meta-analyses. Finding such studies can be effortful, but their exclusion potentially leads to consequential biases like overestimation of a literature's mean effect. We address two key questions using MetaLab, a collection of community-augmented meta-analyses focused on developmental psychology. First, we assess to what extent these datasets include grey literature, and by what search strategies they are unearthed. An average of 11% of datapoints are from unpublished literature, and that standard search strategies like database searches, complemented with individualized approaches like including authors' own data, contribute the majority of this literature. Second, we analyze the effect of including versus excluding unpublished literature on estimates of effect size and publication bias, and find this decision does not affect outcomes. We discuss lessons learned and implications.

Keywords: Meta-analysis, developmental psychology, effect sizes, grey literature

Addressing publication bias in meta-analysis: Empirical findings from community-augmented meta-analyses of infant language development

Meta-analyses are an indispensable research synthesis tool for characterizing bodies of literature and advancing theories. In typical meta-analyses, noisy measurements from multiple independent samples are normalized onto a single scale (typically a measure of effect size) and combined statistically to produce a more accurate measurement. Effects for meta-analysis can come from the published literature, unpublished data, or even the author's own work, but different strategies for identifying datapoints for inclusion can have major consequences for the interpretation of the meta-analytic estimate. In particular, the exclusion of unpublished work can lead to a bias for positive findings and compromise validity. Thus, it is important to assess the utility – and impact – of strategies for including unpublished data. In the present article, we describe our successes and failures with gathering unpublished data for meta-analyses within developmental psychology, and assess how the addition of these datapoints changes the conclusions drawn from meta-analyses.

Community-augmented meta-analyses and MetaLab

Community-augmented meta-analyses (CAMAs, Tsuji, Bergmann, & Cristia, 2014) are a tool for countering some problems faced by traditional meta-analyses. In the original proposal¹, CAMAs are open-access, online meta-analyses which are living documents that can be openly accessed, updated, and augmented (Tsuji et al. 2014). Their dynamic nature avoids a key problem of traditional meta-analyses, which are crystallized at the time of publication and quickly become outdated. Additionally, CAMAs were set up to allow the addition of unpublished datapoints.

MetaLab is a database and browsable web interface that instantiates the CAMA idea (metallab.stanford.edu; Bergmann et al., 2018). The database's focus is Developmental Psychology, and the goal is to eventually cover all subfields on which there are experimental results bearing on infant and child cognition. At present, MetaLab hosts 20 meta-analyses (containing a total of 1,686 effect sizes), covering diverse topics ranging from sensitivity to vowel contrasts (e.g., the sound difference between “ship” and “sheep”; Tsuji & Cristia, 2014) to children's preference for prosocial over anti-social agents (Margoni & Surian, 2018). Most meta-analyses, however, bear on language development, and focus on children aged 5 years or younger.

In the present paper, we analyze 12 meta-analyses in MetaLab for which efforts like search strategy and contact with authors were well documented and accessible to us (containing a total of 1,232 effect sizes; Bergmann & Cristia, 2016; Black & Bergmann, 2017; Carbajal, 2018; Cristia, 2018; Fort et al., 2018; Rabagliati, Ferguson, & Lew-Williams, 2019; Tsuji & Cristia, 2014; Tsui, Byers-Heinlein, & Fennell, 2019; Von Holzen & Bergmann, 2018). We discuss below to what extent our results may generalize to other meta-analyses and fields of psychology.

Unpublished literature in meta-analyses

Since meta-analyses largely build on publicly accessible literature, they face some of the same challenges as primary literature in the context of the replication crisis (Lakens et al., 2017). One key issue concerns the inclusion of unpublished data, that is results that do not appear in the published literature (and hence may not be indexed by all libraries and academic search engines), but are either reported in theses, dissertations, conference abstracts, white papers or internal reports, or not reported publicly at all (i.e. studies that are “file-drawer”).

---

¹ Although we initially aimed for authors and others to add studies to extant meta-analyses, we now favor a system where a single curator is responsible for updating a given meta-analysis. This preserves the original goal of having up-to-date meta-analyses, and further ensures internal consistency in all meta-analyses. This change in the concept of curation (from crowd-sourcing to centralized), however, does not affect the topics that are broached in this paper, and thus will not be discussed further.

Attempting to access unpublished data is difficult and time-consuming. To begin with, reports on these data, if they exist, tend to not be indexed as carefully as published data and thus are harder to discover. For instance, a search on PubMed would not reveal theses or dissertations, whereas Google Scholar does index some (but not all) thesis archives. Even if a meta-analyst uses Google's Scholar engine, conference abstracts and proceedings in many fields are not indexed, and thus need to be searched manually. In some cases, for instance when conferences in a field favor very short abstracts, one may discover the existence of a study, but be unable to integrate it because there is insufficient information reported. In this case, as well as in the case of studies for which reports do not exist, author contact is the only way to secure the information needed to integrate a study into a quantitative analysis. To this end, one may try to write to all authors who have published on the topic, and ask for data in their file drawers. To access also data collected by others, one can publicly post a call for data, for example via field-specific mailing lists. Thus, meta-analysts who intend to include grey literature can be led to make a significant investment in time to be able to discover and integrate such results. To our knowledge, there is no previous research documenting the effectiveness of these modes of grey literature integration for psychological research. Therefore, in Part one below, we have undertaken to document the efficacy of these diverse methods, i) database searches, ii) citation searches, iii) mailing list calls, iv) cases where author's work was known v) inclusion of own data. Relatedly, we also document the success rate in gathering data based on emailing authors with a request for information.

Although discovering and integrating unpublished data is costly, it is often part of standard meta-analytic practice recommendations (e.g., White, 1994), in the hope that it will reduce publication bias. Indeed, published literature might not constitute an unbiased sample of the data, in turn yielding an overestimation of effect sizes in meta-analyses that only include published literature (e.g., Guyat et al., 2011). Ferguson and Heene (2012) note that at least 25%, and possibly as many as 80%, of meta-analyses in psychology, suffer from significant bias. A vast body of evidence confirms that this is a concern for psychological science in particular. For instance, Bakker, van Dijk, and Wicherts (2012) show convincingly that researchers in psychology typically use small sample sizes (with an inordinate proportion of statistically significant results), rather than larger sample sizes (whose higher precision reduces the likelihood of false positives and negatives). The problem is so widespread that item 15 of the PRISMA checklist specifically asks meta-analysts to "Specify any assessment of risk of bias that may affect the cumulative evidence (e.g., publication bias...)" (Moher Liberati, Tetzlaff, Altman, & The PRISMA Group, 2009) and a systematic review of systematic reviews on the effects of *all sources* of bias identifies the inclusion of grey literature among its key recommendations (Tricco et al., 2008).

Most meta-analyses, and therefore most recommendations for meta-analyses, are based on the medical intervention literature, however. Studies of publication biases from this field may or may not generalize to psychological research. One previous study investigated bias and unpublished data inclusion for 91 meta-analyses published in psychological journals (Ferguson & Brannick, 2012). Surprisingly, they concluded that meta-analyses including unpublished data were more, rather than less, biased than studies based purely on published data. These authors recognize the validity of the grey literature inclusion approach for medical meta-analyses, where registries allow for unbiased discovery of studies, and mandatory pre-registration of studies further precludes analyses that favor specific results (Huic, Marušić, & Marušić, 2011). Since neither of these factors exist for psychology, it may be unwise for psychological meta-analyses to include grey literature because (1) the effort will be too large for the number of effect sizes that can be included ultimately (with a median of fewer than 5% of effect sizes stemming from unpublished data; Ferguson & Brannick, 2012); and (2) unpublished data will be biased because they are discovered mainly via a biased network: the meta-analysis' authors and close colleagues, and prominent authors in the field, all of whom may contribute data that favors a given outcome. In view of these contrasting results between the psychological literature and

the body of meta-analytic best practices research, we revisit the question of what the effects of adding unpublished data are based on our CAMAs.

In Part two, we follow previous literature (e.g., Tricco et al., 2008) and report i) effect size estimates for samples with and without unpublished literature; ii) bias estimates with and without unpublished literature; and iii) potential correlates of study quality. We note that study quality is much harder to measure objectively in basic psychological research than in interventions. In interventions, randomized control trials with a double-blind procedure are undoubtedly better quality evidence for causal links than correlational research. Such hierarchy can be harder to establish for some types of laboratory experiments, where procedures like experimenter blinding or randomization exist, but might be implemented much less systematically and consistently than in intervention studies. However, we can at least inspect some general features that may correlate with data quality, for instance a study's sample size. Some previous work suggests that unpublished data is lower quality by being based on smaller samples (e.g., Tricco et al., 2008).

Finally, we dedicate a third part to in-depth case studies and summaries of lessons learned.

Methods

All data reported in the Results section will be based on a subset of meta-analyses openly available on MetaLab. We include those meta-analyses that are based on a systematic literature search, made efforts to include unpublished data, and documented their data gathering efforts systematically. We define as unpublished anything that is not in a peer-reviewed journal, including work that has appeared only in theses, proceedings, and books or book chapters.

Cohen's d is a standardized effect size based on sample means and their variance. We rely here on d values computed in the MetaLab pipeline, which uses standard formulae to convert the measurements reported in papers to d values (details are reported elsewhere; Bergmann et al., 2018, see also metalab.stanford.edu).

Data and analyses scripts are shared on our Open Science Framework project site (https://osf.io/g6abn/?view_only=a9449ba0bb2b48468038bc7688b3484f - link anonymized for peer review purposes). Analyses have been conducted with the tidyverse (Wickham, 2017) and the metafor (Viechtbauer, 2010) packages in R (R Core Team, 2019).

Results

Of the 20 meta-analyses currently included in MetaLab, 14 meta-analyses (70%) include unpublished data. This proportion is comparable to previous reports, where 63% of recent meta-analyses in Psychology made efforts to include grey literature (Ferguson & Brannick, 2012). Of those meta-analyses that included unpublished data, 12 fit our additional criteria for inclusion in the present analysis, namely being based on a systematic literature search, and systematically documenting data gathering efforts and/or making those efforts accessible to us. For the purpose of the present study, we aggregate two pairs of meta-analyses into one, since the systematic literature review in both cases had originally been conducted on the pair, and the datasets were only later thematically separated in MetaLab². The resulting 10 meta-analyses ranged in size from 11 to 73 papers (median = 23), with 20 to 315 datapoints included (median = 92). Table 1 gives an overview of the meta-analyses, including citations, descriptive statistics on the number of effect sizes by publication status, and how this unpublished literature was found.

Table 1. Overview of meta-analyses and number of published and unpublished effect sizes.

² “Word segmentation” and “Function word segmentation” are aggregated into “Natural word segmentation”; “Native vowel discrimination” and “Non-native vowel discrimination” are aggregated into “Vowel discrimination”.

Citation	Meta-analysis	N effect sizes	N unpublished effect sizes					N data obtained through author email
			Overall	Database search	Citation search	Author known	Own	
Carbajal (2018)	Familiar word recognition	33	12	2	4	4	2	4
Bergmann & Cristia (2016)	Natural word segmentation	298	25	23	0	0	2	12
Von Holzen & Bergmann (2018)	Mispronunciation sensitivity	251	32	0	0	33	0	14
Black & Bergmann (2017)	StatSeg	103	10	10	NA	0	0	2
Tsuji & Cristia (2014)	Vowel discrimination	253	11	4	NA	6	1	6
Fort et al. (2018)	Sound symbolism	44	20	2	NA	10	8	6
Rabagliati, Ferguson, Lew-Williams (2019)	Abstract rule learning	95	4	4	NA	0	NA	3
Cristia (2018)	Phonotactic learning	47	11	0	0	1	10	2
Cristia (2018)	Statistical sound category learning	22	9	0	0	5	2	2
Tsui, Byers-Heinlein, & Fennel (2019)	Switch task	143	11	2	0	0	9	3

Evaluation of data-gathering efforts

Overall, of the total of 1,232 effect sizes contained in our MetaLab subset, 144 effect sizes, or 11.7% of data, were based on unpublished literature. Similar to previous reports (Ferguson & Brannick, 2012), the distribution of percentage of unpublished studies shows a positive skew, with most meta-analyses in our sample having a low percentage of unpublished studies, at around 10% or less (Figure 1).

Figure 1. Number of meta-analyses by percentage of unpublished studies included.

i) Database searches. All database searches followed standard meta-analytic practice, wherein a set of pre-determined keywords were entered into a search engine, and titles of hits and the abstracts of potentially relevant papers were scanned to arrive at the sample of papers eligible for inclusion in the respective meta-analysis. In all meta-analyses included, Google Scholar was the search engine of choice, which performs equivalent to a combined search of multiple databases (Gehanno, Rollin & Darmoni, 2013), although exact replicability of Google Scholar searches might be compromised since it saves users' history. Crucially for us, it includes unpublished work (i.e., pre-prints, conference proceedings, or unpublished manuscripts) in its search results as long as they are available online and indexed. We found that an average of 4.7 datapoints or 37.3% (SD = 42.9) of unpublished data was found by Google Scholar searches.

Figure 2. Number of unpublished datapoints obtained through different sources of recovering grey literature, by meta-analysis.

ii) Citation searches. Some unpublished studies might not be available online and thus not detectable by search engines. They might, however, be discoverable through searching the reference lists of available studies. The percentage of unpublished datapoints gathered based on citation searches was an average of 0.4 datapoints or 3.3% (SD = 10.5).

iii) Mailing list requests. In order to reach a relevant audience to recover potential grey literature, authors of six of the included meta-analysts requested contributions via professional email lists. Strikingly, these attempts did not lead to a single reply with information that could be added to the meta-analyses.

iv) Author's work known. In addition to the more formal routes described above, a meta-analyst can get to know an author's eligible work at a conference, or via informal communication with experts in the field. Our estimate of datapoints added via this route is an average of 5.9 datapoints or 31.8% (SD = 35.9). Since a meta-analyst is often an expert in the topic of their meta-analysis, this can be a very fruitful route - one of our special cases below will illustrate how helpful this strategy can be for the data gathering process.

v) Own data. A meta-analyst might also contribute their own unpublished data to their meta-analysis. In MetaLab, unpublished data from meta-analysts' own research accounted for an average of 3.4 datapoints or 27.5% (SD = 33.7). Previous reports have documented a difference in own published and unpublished contributions, with an average of 5.89% of published datapoints, and 12.94% of unpublished datapoints, being based on meta-analysts' own data (Ferguson & Brannick, 2012). If we assess our data with this metric, we find that an average of 9.6% of published datapoints, and 27.5% of unpublished datapoints, are based on own data in MetaLab. While this is in accordance with previous reports, note, however, that there are much more total published than unpublished datapoints in the datasets. If we look at the absolute number of datapoints added, meta-analysts added more published (an average of 6.1) than unpublished (an average of 3.4) own datapoints to their datasets.

vi) Emails to authors. Meta-analysts can chose to contact authors of papers eligible for inclusion in their meta-analysis with request for additional information from eligible literature (whether published or unpublished), and whether they are aware of any grey literature. Since it was impossible to recover the number of effect sizes added based on these requests, we instead counted the number and ratio of authors contacted, and how many of these contacts were responsive and lead to data that could be added to the meta-analysis. An average of 12.7 authors were contacted. Out of all authors contacted, an average of 9.8 or 85.1% (SD = 19.5) were responsive and 5.4 or 49.6% (SD = 28.1) provided data that could be added to the meta-analysis.

vi) Community contributions. Although the original CAMA idea entailed that, ultimately, the research community would take over the curation of datasets and infrastructure in a bottom-up fashion, this model has not proven feasible, both for lack of contributions and difficulty in curating erroneous entries. Instead, MetaLab now consists of a governing board that, aided by funding sources³, maintains and expands the infrastructure. Dedicated curators for each dataset are in charge of updating individual meta-analyses.

Meta-analyses with published and unpublished results

The second part of our analyses evaluated the effect of including grey literature on publication bias. We evaluate the effect of including or excluding unpublished literature on effect size and bias estimates in our samples, as well as attempt to assess potential correlates of study quality.

³ <https://www.bitss.org/projects/metabolab-paving-the-way-for-easy-to-use-dynamic-crowdsourced-meta-analyses/>

Figure 3. Mean meta-analytic effect size estimates and associated confidence intervals. Values are based on meta-analytic regression models for each dataset. Shapes represent different subsets per meta-analysis based on their publication status, lines indicate the 95% CI.

i) Effect size estimates by publication status. The mean Cohen's d effect size across meta-analyses is $d = 0.22$ (range -0.34 - 0.77). The mean is $d = 0.24$ (range -0.17 - 0.66) for published, and $d = 0.15$ (range -0.34 - 0.77 for unpublished studies, consistent with publication bias (greater effects for journal-published studies) as well as a difference in data quality (lower for not journal-published studies). However, such an analysis ignores factors that are known to vary across studies.

Therefore, in order to more specifically assess the effect of inclusion of unpublished datapoints into a meta-analysis, we constructed meta-analytic regression models for the full dataset as well as datasets including only the published or unpublished datapoints (see e.g., Tricco et al., 2008). The model for each meta-analysis included infant age as a predictor, since this factor has been consistently found to explain variance in effect sizes (Bergmann et al., 2018). While testing method is another such factor, we refrained from including it in our analyses, since data subsets differ in the amount of testing methods included, with some subsets being comprised of data stemming from only one testing method. Infant age was centered before analyses. Our random effects structure allowed shared variance for datapoints stemming from the same paper, and accounted for the dependence between datapoints stemming from the same infant participants contributing multiple effect sizes (model = `rma.mv(d, d_var, mods = ~ age, random = ~ same_infant | paper/data_point)`). Since estimates become imprecise with too small datasets, we excluded from this analysis all cases where the subsets of published or unpublished data only contained fewer than 10 datapoints. This excluded four datasets from the analysis of unpublished datasets. Since the meta-analytic regression of one dataset in the subset of published data that contained 11 published datapoints did also not converge, this dataset was also excluded from the analysis.

Figure 3 shows the resulting effect size estimates and associated confidence intervals. There is no clear pattern in terms of higher meta-analytic effect size estimates for published datasets, consistent with previous reports (e.g., Chow & Ekholm, 2018; Guyat et al., 2011), and confidence intervals for the respective sets mostly overlap. Thus, when known factors structuring variance (age, meta-analysis) are accounted for, it does not appear to be the case that the inclusion of grey literature affects meta-analytic conclusions in terms of overall effect size.

ii) Bias estimates with and without unpublished literature. In order to evaluate the impact of inclusion or exclusion of unpublished literature on bias estimates, we assessed each individual meta-analysis by means of funnel plot asymmetry, a classical diagnostic for identifying potential publication bias (Egger, Smith, Schneider, & Minder, 1997). We again included as moderators infant age b . The distribution of test statistics for Egger's test for funnel plot asymmetry did not differ whether we assessed datasets under exclusion of unpublished studies ($z_{\min} = -1.90$, $z_{\text{mean}} = 3.53$, $z_{\max} = 12.44$), or when grey literature was included ($z_{\min} = -1.89$, $z_{\text{mean}} = 3.33$, $z_{\max} = 13.41$) (Fig. 4). In both subsets, the same 4 out of 10 datasets showed significant funnel plot asymmetry, suggesting that adding grey literature did not improve this indicator of publication bias⁴.

In order to evaluate the effect of studies' publication status more directly, we ran a random-effects meta-analytic regression for each of the six meta-analyses that had 10 or more unpublished datapoints, using publication status (unpublished, published) as a moderator. Publication status was

⁴ A more quantitative evaluation of this difference proved difficult. For instance, a non-parametric bootstrapping approach is unreliable for meta-analyses with fewer than about 50-100 studies, which is the case for the majority of the meta-analyses included in the present analysis.

deviation-coded. We again included infant age, centered, as an additional moderators. We assessed whether publication status had a significant effect in each of these meta-analyses, and found that this was the case in two out of six cases (with a third case being marginally significant). In all of these, grey literature had a negative impact on effect sizes (Table 2). Thus, even though including grey literature did not cause a significant change in publication bias, it did explain a significant proportion of variance in several datasets. Further, grey literature inclusion may also be relevant in other cases but we failed to detect it due to meta-analyses being small, and thus low-powered.

Table 2. Meta-analytic regression coefficients for the effect of publication status by dataset. Note that we only include those six meta-analyses that had 10 or more unpublished datapoints. A negative beta indicates a lower effect size for unpublished than published studies. Se = standard error, ci.lb = lower boundary of confidence interval, ci.ub = upper boundary of confidence interval

	beta	se	z	p	ci.lb	ci.ub
Familiar word recognition	0.12	0.063	1.916	0.055(*)	-0.003	0.242
Mispronunciation sensitivity	-0.599	0.103	-5.788	< .001*	-0.802	-0.396
Natural word segmentation	-0.071	0.038	-1.846	0.065(*)	-0.146	0.004
Phonotactic learning	0.091	0.079	1.155	0.248	-0.064	0.246
Sound symbolism	-0.127	0.062	-2.032	0.042*	-0.249	-0.005
Switch task	0.029	0.076	0.379	0.705	-0.119	0.177

iii) Correlates of study quality. There are no uniformly agreed-upon markers of study quality in experimental psychology research that can easily be assessed. Since unpublished effect sizes are often based on smaller samples (Tricco et al., 2008), we assessed whether sample size could serve as a proxy of quality. However, a descriptive assessment of sample size revealed no difference in sample size for published ($M = 21.7$, $SD = 9.9$) and unpublished studies ($M = 22.5$, $SD = 10.3$) in the meta-analyses under study.

In addition to sample sizes, we attempted to assess one other potential indicator of study quality, namely internal correlations in within-participant designs. Since weighted meta-analytic regression requires an estimate of these correlations, some meta-analysts have gathered this measure. A higher internal correlation might suggest less noise in the measure, thus potentially indicating higher study quality. We first checked whether the degree of internal correlation correlated positively with child age (since measure precision tends to improve as children age). However, our assessment showed that internal correlation and child age anti-correlated significantly [$r = -0.28$, $t(417) = -5.88$, $p < .001$]. Since this result contradicted our initial assumption, we did not pursue this possibility further.

Overall, therefore, we were not able to show any relationship between potential measures of study quality and publication status, likely due in part to the lack of such objective criteria in experimental studies in psychology.

Figure 4. Funnel plots by dataset. Effect size estimates for published data are represented by black points, and estimates for unpublished data by yellow points. The mean effect size for the full dataset is shown as a red dashed line (invisible when the means for full and published data sets overlap), and the grey shaded funnel corresponds to a 95% CI around this mean. The mean effect size for the subset of published studies is shown as an orange dashed line, and the transparent dashed funnel corresponds to a 95% CI around this mean. The grey dashed line shows an effect size of zero. In the absence of bias, we should expect all points to fall symmetrically inside the funnel.

Case studies

Case study 1: Expanding the pool of meta-analyses. In addition to gathering unpublished data and missing datapoints for extant CAMAs, it was also sought to expand the pool of CAMAs available on MetaLab. For this purpose, a call for contributions to MetaLab was issued, with a \$1000 cash prize for the top three most extensive meta-analyses submitted. To advertise this challenge, announcements were sent to professional mailing lists, the literature was searched for extant meta-analyses fitting the scope and their authors were contacted. Information on ongoing meta-analyses efforts were informally gathered and distributed. These efforts resulted in six eligible submissions for the challenge, four of which are already integrated in MetaLab. Considering that these will finally constitute 27.2% (6/22) of meta-analyses on MetaLab, this strategy substantially expanded the database at relatively low cost compared to the cost of performing new meta-analyses.

Case study 2: Gathering grey literature through in-person author contact. One of the MetaLab datasets on sound symbolism in infancy (Fort et al., 2017) assesses the development of the

bouba-kiki effect, whereby humans associate pseudowords like “bouba” with round objects, and pseudowords like “kiki” with spiky objects. Experiments examining this effect have had mixed results in infant populations. Anecdotally, several researchers have failed to find a consistent bouba-kiki effect, but have faced problems publishing these null results. Encountering others researchers with similar results through conference presentations, they encouraged each other to share their unpublished data and decided to conduct a meta-analysis on the phenomenon with both published and unpublished results. This meta-analysis reveals that there is overall evidence for a bouba-kiki effect in infants, however, it is smaller than suggested by the published literature alone. This small effect size, combined with habitually small sample sizes in infants studies, likely explains the divergence in findings between attempts to elicit the bouba-kiki effect. In the case of this meta-analysis, presenting null results at a relevant conference, and thus making others aware of their existence, proved a highly effective way to assemble grey literature.

Discussion

Publication bias is considered a key problem of the meta-analytic literature, and the exclusion of grey literature from meta-analyses is a potential cause. Since the difficulty of accessing such grey literature is a reason for the lack of unpublished studies in meta-analyses, we assessed the amount of grey literature gained based on various strategies in datasets assembled in the open-access database MetaLab (Bergmann et al., 2018). In the following, we will discuss these strategies from the viewpoint of lessons learned and recommendations for future meta-analysts. We further assessed the impact of including such grey literature on publication bias. These analyses show that our efforts had only a moderate impact on publication bias in our datasets, a result we will discuss in light of previous literature and the nature of grey literature gathered in our datasets.

Lessons learned from efforts gathering missing data

If an article does not report all data necessary for estimating effect sizes, contacting the original authors of the article is the only way to possibly obtain these missing data. Although this is an effortful endeavor, our analysis of data gathering efforts shows that it is worthwhile, because authors contacted individually by email were highly responsive, and sent data useful for computing effect sizes in almost half of the cases. While we do not have comparative data for other approaches, the successful author contacts by meta-analysts in MetaLab are based on highly individualized emails to the respective first and/or last authors of an article. Along with outlining the general aim of the meta-analysis, we would mention both the authors’ and their article’s name, and explain in detail the nature of data needed from them. Habitually, we would send one follow-up email in case we didn’t get a reply. In addition to these efforts, it has been shown that adding data-sharing agreements to requests for primary data can improve responsiveness (Polanin & Terzian, 2019).

Lessons learned from efforts gathering grey literature

On average, 11.4% of effect sizes were based on unpublished studies in the present data set. Standard tools like database and citation searches, which arguably take less time and effort than more personalized efforts like knowing of an author’s work, can already contribute an important amount of grey literature. Although it is difficult to estimate the exact time effort required for each, in case of database searches some grey literature would be included in the search automatically depending on the engine chosen, and thus the search itself would not require more time investment. Although citation searches are an added effort, a meta-analyst would be recommended to carefully read the included literature, which is not too many steps away from the citation search itself.

Knowing the author of a study (whether in person or not), proved to be a fruitful strategy to gather data. Although this strategy might be more dependent on an individual meta-analysts’ network and the availability of unpublished studies in conferences and other places enabling personal contact,

our analyses and case study on sound symbolism illustrate that this is a promising way to include unpublished studies, and should be on the meta-analyst's mind as a possible way to gather data. Ideally, though, we think it would be desirable that there were better indexing of unpublished literature (by authors and conference organizers uploading their unpublished work to searchable archives). A more reliable index of grey literature would reduce the individual's dependence on high effort strategies and enable discovery via the standard database search strategies, which at the same time increases transparency. Finally, in line with previous literature (Ferguson & Brannick, 2012), the meta-analyses included here contained a relatively high proportion of effect sizes stemming from the meta-analysts' own data. A meta-analysts' own unpublished data, or such data in their network, might induce bias, since owning such data might serve as a motivator to conduct an MA, or own data might be more likely included in an MA before peer review. Thus, such potential bias might counteract the otherwise bias-reducing effect of adding unpublished literature to an MA.

The effect of adding unpublished literature on effect sizes and publication bias

Previous literature has reported higher effect sizes among published effect sizes, potentially leading to an overestimation of effect sizes if unpublished literature is not included. In contrast, other authors have warned against the inclusion of grey literature because it would increase bias (Ferguson & Brannick, 2012). Neither of these problems were apparent in our data set: meta-analytic regressions per dataset showed largely overlapping confidence intervals and bias estimations were virtually identical, when unpublished effect sizes were included as opposed to excluded.

Adding grey literature might not be equally illuminating, necessary, or damaging in every field of psychology. With regard to the advantages and disadvantages of including grey literature in the case of infant literature, we suggest that such literature will improve the overall quality of a database, for at least two reasons. First, the infant literature is tremendously underpowered and benefits from a larger body of studies to better estimate true effect sizes. Second, in order to conduct infant experiments, a researcher habitually needs to undergo training and make use of a specialized laboratory, making it unlikely that unpublished data are especially prone to being badly designed or executed. Third, if grey literature is added the way we suggest in the context of CAMAs such that each study is coded based on publication status, meta-analysts and database users can decide for themselves whether or not to include unpublished datapoints in their assessments.

A relatively large amount of unpublished data in the present dataset was based on potentially biased data, namely the meta-analyst's own data or data based on direct contact with study authors. Potentially, reducing this bias in the future might impact the difference between effect sizes and measures of publication bias between published and unpublished studies in the present dataset.

Limitations and opportunities

Meta-analyses in MetaLab are not a random or representative selection of meta-analyses in psychology. In fact, two of the included meta-analyses themselves have not yet been published (Carbajal, 2018; Tsui et al., 2019), and two more have been peer-reviewed for proceedings papers (Black & Bergmann, 2017; Von Holzen & Bermann, 2018), and thus they have not, or to a moderate degree, been affected by the review process. The meta-analysis authors were often students who were doing a systematic review of a literature to which they were planning to contribute, and thus may not have had as much of a vested interest in supporting one or another theory as more established researchers might do. Finally, most of them come from a cluster of researchers (including the authors of this paper) who strived to follow best practices guidelines such as the PRISMA statement. (Moher et al., 2009).

While these characteristics might limit comparability with other attempts, the transparent data gathering process by relatively unbiased meta-analysts might enable us to assume that the biases found in the datasets can be attributed to the literature itself more than to the meta-analytic process. Then,

considering that all meta-analysts of datasets included in the present sets attempted to gather grey literature, but publication bias was still prevalent, the MetaLab subset re-emphasizes a broader problem of the field, namely the lack of publicly available indexing of grey literature.

Another characteristic of MetaLab is its basis on the CAMA approach, which from the outset was meant to function as a natural home for file-drawer studies in addition to published studies. Opening the file-drawer in this way can also help us to estimate how many studies are filtered out by the peer review process in the future. Metalab's growth to now 20 datasets indicates its success. The website's visibility has led to numerous conference presentations on the included meta-analyses as well as invitations to provide tutorials, which in turn have inspired others to start their own meta-analyses and become curators, or else to render more visible extant meta-analyses.

References

- Bakker, M., van Dijk, A., & Wicherts, J. M. (2012). The rules of the game called psychological science. *Perspectives on Psychological Science*, 7(6), 543-554.
<https://doi.org/10.1177/1745691612459060>
- Bergmann, C. & Cristia, A. (2016). Development of infants' segmentation of words from native speech: A meta-analytic approach. *Developmental Science*, 19(6), 901-917.
<https://doi.org/10.1111/desc.12341>
- Bergmann, C., Tsuji, S., Piccinini, P.E., Lewis, M.L., Braginsky, M., Frank, M.C., & Cristia, A. (2018). Promoting replicability in developmental research through meta-analyses: Insights from language acquisition research. *Child Development*, 89, 1996-2009.
<https://doi.org/10.1111/cdev.13079>
- Black, A., & Bergmann, C. (2017). Quantifying infants' statistical word segmentation: A meta-analysis. In: *Proceedings of the 39th Annual Conference of the Cognitive Science Society*, (pp. 124-129). Austin, TX: Cognitive Science Society.
- Carbajal, M.J. (2018). *Doctoral dissertation*.
- Chow, J. C., & Ekholm, E. (2018). Do published studies yield larger effect sizes than unpublished studies in education and special education? A meta-review. *Educational Psychology Review* 30(3), 727-744. <https://doi.org/10.1007/s10648-018-9437-7>
- Cristia, A., (2018). Can infants learn phonology in the lab? A meta-analytic answer. *Cognition*, 180, 312-327. <https://doi.org/10.1016/j.cognition.2017.09.016>
- Egger, M., Smith, G. D., Schneider, M., & Minder, C. (1997). Bias in meta-analysis detected by a simple, graphical test. *British Medical Journal*, 315(7109), 629-634.
<https://doi.org/10.1136/bmj.315.7109.629>
- Ferguson, C. J., & Brannick, M. T. (2012). Publication bias in psychological science: prevalence, methods for identifying and controlling, and implications for the use of meta-analyses. *Psychological Methods*, 17(1), 120-128. <https://doi.org/10.1037/a0024445>
- Ferguson, C. J., & Heene, M. (2012). A vast graveyard of undead theories: Publication bias and psychological science's aversion to the null. *Perspectives on Psychological Science*, 7(6), 555-561.
<https://doi.org/10.1177/1745691612459059>
- Fort, M., Lammertink, I., Peperkamp, S., Guevara-Rukoz, A., Fikkert, P., & Tsuji, S. (2018). SymBouki: a meta-analysis on the emergence of sound symbolism in early language acquisition. *Developmental Science*, 21(5), e12659. <https://doi.org/10.1111/desc.12659>
- Gehanno, J. F., Rollin, L., & Darmoni, S.(2013). Is the coverage of Google Scholar enough to be used alone for systematic reviews. *BMC Medical Informatics and Decision Making*, 13(1), 7.
<https://doi.org/10.1186/1472-6947-13-7>
- Guyatt, G. H., Oxman, A. D., Montori, V., Vist, G., Kunz, R., Brozek, J., ... & Williams Jr, J. W. (2011). GRADE guidelines: 5. Rating the quality of evidence-publication bias. *Journal of Clinical Epidemiology*, 64(12), 1277-1282. <https://doi.org/10.1016/j.jclinepi.2011.01.011>
- Huić, M., Marušić, M., & Marušić, A. (2011). Completeness and changes in registered data and reporting bias of randomized controlled trials in ICMJE journals after trial registration policy. *PloS one*, 6(9), e25258. <https://doi.org/10.1371/journal.pone.0025258>
- Lakens, D., van Assen, M. A. L. M., Anvari, F., Grange, J. A., Gerger, H., Hasselman, F., ... Zhou, S. (2017). Examining the Reproducibility of Meta-Analyses in Psychology: A Preliminary Report. *BITS Preprint*. <https://doi.org/10.31222/osf.io/xfbjf>
- Margoni, F., & Surian, L. (2018). Infants' evaluation of prosocial and antisocial agents: A meta-analysis. *Developmental Psychology*, 54(8), 1445. <https://doi.org/10.1037/dev0000538>
- Moher, D., Liberati, A., Tetzlaff, J., Altman, D. G., The PRISMA Group (2009). Preferred Reporting Items for Systematic Reviews and Meta-Analyses: The PRISMA Statement. *BMJ*, 339:b2535.
<https://doi.org/10.1136/bmj.b2535>

- Polanin, J. R., & Terzian, M. (2019). A data-sharing agreement helps to increase researchers' willingness to share primary data: results from a randomized controlled trial. *Journal of Clinical Epidemiology*, 106, 60-69. <https://doi.org/10.1016/j.jclinepi.2018.10.006>
- Rabagliati, H., Ferguson, B., & Lew-Williams, C. (2019). The profile of abstract rule learning in infancy: Meta-analytic and experimental evidence. *Developmental Science* 22: e12704. <https://doi.org/10.1111/desc.12704>
- R Core Team (2019). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL: <https://www.R-project.org/>.
- Tricco, A. C., Tetzlaff, J., Sampson, M., Fergusson, D., Cogo, E., Horsley, T., & Moher, D. (2008). Few systematic reviews exist documenting the extent of bias: a systematic review. *Journal of Clinical Epidemiology*, 61(5), 422-434. <https://doi.org/10.1016/j.jclinepi.2007.10.017>
- Tsui, A. S. M., Byers-Heinlein, K., & Fennell, C. T. (2019). Associative word learning in infancy: A meta-analysis of the switch task. *Manuscript under review*. <https://doi.org/10.1037/dev0000699>
- Tsuji, S., Bergmann, C., & Cristia, A. (2014). Community-augmented meta-analyses: Toward cumulative data assessment. *Perspectives on Psychological Science*, 9(6), 661-665. <https://doi.org/10.1177/1745691614552498>
- Tsuji, S. & Cristia, A. (2014). Perceptual attunement in vowels: A meta-analysis. *Developmental Psychobiology*, 56(2), 179-191. <https://doi.org/10.1002/dev.21179>
- Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, 36(3), 1-48. URL: <http://www.jstatsoft.org/v36/i03/>
- Von Holzen, K., & Bergmann, C. (2018). A meta-analysis of infants' mispronunciation sensitivity development. *Proceedings of the 40th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society, Inc.
- White, H. D. (1994). Scientific communication and literature retrieval. In H. Cooper & Hedges, L. V. (Eds). *The handbook of research synthesis* (pp. 41-56). New York, NY: Russell Sage Foundation.
- Wickham, H. (2017). tidyverse: Easily Install and Load the 'Tidyverse'. R package version 1.2.1. <https://CRAN.R-project.org/package=tidyverse>