Decision-Making Under Uncertainty in Research Synthesis: Designing for the Garden of Forking Paths

Alex Kale
University of Washington
Seattle, Washington
kalea@uw.edu

Matthew Kay
University of Michigan
Ann Arbor, Michigan
mjskay@umich.edu

Jessica Hullman
Northwestern University
Evanston, Illinois
jhullman@northwestern.edu

ABSTRACT

To make evidence-based recommendations to decision-makers, researchers conducting systematic reviews and meta-analyses must navigate a *garden of forking paths*: a series of analytical decision-points, each of which has the potential to influence findings. To identify challenges and opportunities related to designing systems to help researchers manage uncertainty around which of multiple analyses is best, we interviewed 11 professional researchers who conduct research synthesis to inform decision-making within three organizations. We conducted a qualitative analysis identifying 480 analytical decisions made by researchers throughout the scientific process. We present descriptions of current practices in applied research synthesis and corresponding design challenges: making it more feasible for researchers to try and compare analyses, shifting researchers' attention from rationales for decisions to impacts on results, and supporting communication techniques that acknowledge decision-makers' aversions to uncertainty. We identify opportunities to design systems which help researchers explore, reason about, and communicate uncertainty in decision-making about possible analyses in research synthesis.

CCS CONCEPTS

• **Information systems** → **Decision support systems**; • **Human-centered computing** → **Visualization theory, concepts and paradigms**.

KEYWORDS

Representing Uncertainty, Research Synthesis

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 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300432>

ACM Reference Format:

Alex Kale, Matthew Kay, and Jessica Hullman. 2019. Decision-Making Under Uncertainty in Research Synthesis: Designing for the Garden of Forking Paths. In *Proceedings of CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019)*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3290605.3300432>

1 INTRODUCTION

Organizations routinely rely on prepared summaries of empirical evidence to support decision-making. For example, the Navy employs scientists who review scientific literature and collect data internally in order to recommend improvements in training practices. Similarly, Veterans' Affairs employs researchers who meta-analyze the scientific literature on treatments for post-traumatic stress disorder (PTSD) and other conditions in order to recommend the best possible treatment options for veterans struggling with trauma.

When compiling and communicating a summary of scientific evidence, researchers make a series of analytical decisions such as how to combine information from studies conducted with different measures or in different settings. Recent work in reproducible statistics [64, 65, 67, 69, 77], driven by concerns about a “replication crisis”, demonstrates how flexibility in decision-making produces multiple possible sequences of analytical decisions, which an analyst chooses between at their discretion. In the context of research synthesis, alternative possible analyses may lead to alternative understandings of empirical evidence and consequently, opposite or inconclusive recommendations. Faced with a “garden of forking paths” [26], researchers cannot eliminate subjectivity and uncertainty from the scientific process. Instead, scholars suggest that researchers should attempt to understand which analytical decisions impact results [67, 69].

By identifying researcher degrees of freedom [77] as a cause of the “replication crisis”, prior work seems to focus blame on the indiscretions of individual researchers. However, with existing software to tools, it is difficult for researchers to deliberate about and explore the consequences of alternative analyses (e.g., [67, 69]), such that even researchers with honest intentions may struggle to perceive the implications of different choices. In this study, we seek an

in-depth understanding of how researchers make analytical decisions in research synthesis, and where they struggle with uncertainty in the process, in order to identify opportunities to design for decision-making in the garden of forking paths.

We contribute the results of a qualitative analysis of open-ended conversational interviews we conducted with 11 researchers who work to support evidence-based decision-making at three institutions: the Navy, the Medical Center at a large public university, and the Veterans' Affairs Medical Center in a major US city. In our interviews we elicited detailed descriptions of experiences conducting scientific review and analysis, emphasizing the reasoning behind analytical decisions and the strategies that researchers use to manage uncertainty. Based on these interviews, we identify a set of challenges around managing uncertainty in research synthesis: the tension between surveying analysis paths and implementing a specific path, the disconnect between researchers' rationales for analytical decisions and the actual impact of those decisions on findings, and the balance between researchers' skepticism and the need for compelling recommendations. We draw on utility theory [75] and prior work on reliable statistics and uncertainty visualization, in combination with our interviews, to identify opportunities to address these challenges by designing systems which encourage exploration of alternative analyses, elicit and represent researchers' reasoning about analytical decisions, and provide researchers with techniques to communicate uncertainty in the research process.

2 BACKGROUND

Managing Uncertainty in Decision-Making

A large body of research on judgment and decision-making (JDM) has examined how people reason with and make decisions under uncertainty. Canonical work by Tversky and Kahneman [37, 73] established that people often seek to reduce uncertainty, sometimes by substituting heuristic judgments for more complex reasoning. A drive to reduce uncertainty can lead to unwarranted expressions of certainty [52], which has consequences for decision-making individually and at an organizational level (e.g., in public policy).

Decision-making under uncertainty is characterized by feelings of conflict and doubt which block or delay a choice between alternative courses of action [55]. As such uncertainty is broadly defined and associated with a variety of terms [10, 47] such as ambiguity [32, 54], risk [1, 3, 49], unreliability, imprecision, incompleteness, and contradiction [44], as well as error and subjectivity [50]. Based on a literature review of work describing real world decision-making under varying forms of uncertainty, Lipshitz and Strauss [47] developed a framework for understanding how decision-makers in organizations like the military cope with uncertainty. In

their framework uncertainty is either *acknowledged* through preemptive action and planning, *reduced* through rule- or assumption-based reasoning, or *suppressed* through ignoring information or guesswork. We borrow these strategies from Lipshitz and Strauss to characterize analytical decisions described by our participants during interviews.

Systematic Review and Meta-Analysis

Systematic review and meta-analysis are methodologies used to produce a rigorous summary of existing evidence on a topic. Using systematic review, researchers account for scientific literature within a consistent framework and characterize the research to date on a particular topic [58]. As an extension of systematic review, researchers sometimes choose to aggregate quantitative results from studies through a meta-analysis [17, 46, 58]. Meta-analysis produces an estimate of the effect size of an intervention by pooling statistical outcomes from studies conducted under similar conditions.

Ideally, when prescribed procedures are followed, systematic review and meta-analysis offer more robust findings¹ than regular literature review and quantitative analysis, respectively [58]. A typical systematic review starts with a question about the effect of some intervention and database queries to find all the relevant literature. Researchers independently judge which articles to include and exclude from the review and then resolve disagreements. Since working in pairs is standard practice, this is sometimes called dual-review. For each study, researchers document information on effect sizes and contextual factors (e.g., research design, subject populations) in spreadsheets, sometimes called evidence tables. The data collected in evidence tables is then statistically aggregated in meta-analysis. Following these procedures helps researchers answer a targeted research question while mitigating potential biases (e.g., selection bias, confirmation bias) that might otherwise accrue.

In practice, researchers sometimes do not adhere strictly to these standards. Researchers under time pressure may take shortcuts (i.e., rapid review [24, 40, 76]) by surveying the literature through citation trails or by making inclusion and exclusion decisions individually. If the literature on a topic is sparse or studies cover a variety of populations and contexts, it is difficult to follow conventions of systematic review and conventional meta-analysis may be inadvisable [17, 46]. Researchers may choose to conduct a scoping study [2, 45] in which they survey the breadth of literature to identify gaps in knowledge. The impacts of rapid review and scoping study methods on quality of findings have only been studied recently [60, 72], and there is disagreement about best practices. We contribute a characterization of the gap between best practice and actual practice.

¹See Manski's critique [53] about difficulty interpreting meta-analyses.

Software for Research Synthesis

Interactive systems for research synthesis offer relatively little support for reasoning about possible analysis paths. Software tends to provide features for common steps in the analysis process such as forms for risk of bias assessment [16, 71], data extraction [16, 71], or creating forest and funnel plots of meta-analytic results [6, 9, 16, 71, 74]. Additionally, most tools focus exclusively on a single stage in the analysis process [5] such as study screening [59] or meta-analysis [6, 9, 74]. Designing research software as a set of isolated optional procedures forces the researcher to conduct alternative analyses sequentially and separately. Further, this design fails to represent the motivations and constraints that drive researchers' decision-making.

Three tools in particular—RevMan [16], Eppi-Reviewer [71], and Rayyan [59]—offer collaboration features, such as role assignment [16], review flow diagrams [71], and coding disagreement overviews [59], which help researchers coordinate and review their work. However, with the exception of highlighting disagreements about study inclusion/exclusion (a decision with implicit alternatives), these features do not help researchers identify and weigh alternative analyses. Eppi-Reviewer [71] allows users to create custom “codesets” and annotations to code the status of studies under review and take notes. However, researchers may overlook the benefits of using these features to document motivations and constraints that guide their analyses in order to support later recall or scrutiny, instead using them inconsistently or not at all. Rayyan [59] uses a pulldown interface to enable researchers to select or create a reason for their decision to include/exclude a study. While this elicitation technique is a promising way of linking reasoning to decisions, Rayyan does not document competing motivations or constraints unless they are provided by different users who disagree about study screening.

We point to opportunities for tools to better support researchers in exploring, reasoning about, and communicating uncertainty about alternative analyses.

3 METHOD

We conducted open-ended conversational interviews with professional researchers to investigate practices in applied research synthesis. The goals of these interviews were to (1) characterize how researchers manage possible analysis paths, (2) gather information about the reasoning behind their choices, and (3) study how interactive systems can support awareness of uncertainty in analytical decision-making.

Sampling Participants

We employed convenience and snowball sampling [19] to find 11 professional scientists to interview for our study.

First, we reached out to four professors at the Medical Center of a large public university who had recently published at least one systematic review or meta-analysis to advocate for evidence-based practices, such as utilitarian healthcare policies. Through these professors, we were able to interview one additional postdoctoral researcher studying teaching strategies in STEM. Second, we recruited two PhDs from the Veterans' Affairs Medical Center in a major US city working on systematic review and meta-analysis to recommend treatments for veterans with post-traumatic stress disorder (PTSD) and substance use disorder (SUD). Last, we used connections in the Navy to recruit four scientists using research synthesis to recommend improvements in training practices for military pilots. Of all the 13 people we contacted, only two declined to interview. Our sample represents people doing research synthesis in formal professional settings.

Interview Guide

We created an open-ended conversational interview guide [48, 56] with the objective of getting participants to discuss their research practices in terms of specific examples. The interview guide was a list of topics of interest regarding decision-points at different stages of the research process: scoping research questions, sampling literature, assessing and organizing evidence, analysis and visualization, and communicating findings (see Supplemental Material ²). For the first five interviews, the guide was formatted as a list of questions, but we abridged these questions to a list of topics [48] covering the same content in order to better accommodate the need to ask questions in terms of the experiences of individual researchers [56], whose work varied in methods and settings. Most often the interviewer broached a topic by asking a question of the form, “How did you...” For example, “When you conducted that literature review, how did you decide which papers to include or exclude from your review?” The interviewer also asked follow-up questions to seek clarification or greater detail about particular analytical decisions. The interview guide structured our conversations around a consistent set of topics while allowing flexibility to probe for a greater depth of description when necessary.

Qualitative Coding Process

All interviews, transcription, qualitative coding, and analysis were conducted by the first author, with iterative feedback on the coding and analysis framework from the other two authors. The analyses presented in this paper represent the perspectives of our participants systematically curated in an interpretive framework which was developed through discussions among the authors.

²https://github.com/kalealex/analysis_paths_research_synthesis/tree/master/interview

Figure 1: A pictorial representation of strategies used to navigate the garden of forking paths. The diagram at the top depicts observed temporal orderings of these strategies, e.g., multiplexing across choices in a loop of acknowledgement and reduction.

Transcription. The first author listened to audio recordings of interviews and transcribed all episodes of interest, omitting from transcription only small talk that was obviously not relevant to research synthesis.

Creating a Coding Framework. To build familiarity with the content in the early interviews, the first author used open coding [19] to describe what participants said about their research practices (e.g., literature review, meta-analysis, communication). Open coding helped us determine what we could reasonably infer from our interviews, similar to the use of grounded theory in related prior work [10]. Three provisional themes emerged. Participants described *decisions* they made throughout the scientific process, the *reasons and constraints* which guided those decisions, and specific ideas about *how software features could support their work*.

We used the themes identified in open coding to develop a more targeted framework (see Analytical Framework) characterizing analytical decisions and the reasoning behind them. For each decision, we made a set of categorical judgments about the nature of the decision and the rationale provided and recorded contextual information about the practices described. We continued to code instances where interviewees stated needs for software support in order to help us identify design opportunities in software for research synthesis. We kept track of these codes in a spreadsheet with one row per decision (see Supplemental Material ³).

Fine-tuning this framework was an iterative process. The first author coded analytical decisions and noted issues that came up during coding. Then, the first author presented these issues to collaborators for feedback and discussion. The final coding scheme (described below) represents our consensus about how to best characterize the practices described by our

participants as they relate to the challenges of navigating the garden of forking paths.

Analytical Framework

At the heart of the framework are **analytical decisions** described by participants. We were primarily concerned with how each decision acted on the space of possible analysis paths, for example, by surveying multiple paths or selecting one path through some procedure (Fig. 1). Following Lipshitz and Strauss [47], we coded each decision as an instantiation of one of the following three strategies:

- (1) *Acknowledge [Ack.]* uncertainty by accounting for different possibilities and planning to confront or avoid potential risks. This includes decisions to explore possible analysis paths, orient the research toward a broad set of issues, check for potential problems, and provide details or caveats in order to preempt misinterpretations of scientific evidence.
- (2) *Reduce [Red.]* uncertainty by gathering information, applying rules or conventions, making assumptions, or adopting a specific procedure to exert control over uncertainty. This includes decisions to seek information on a particular topic, to select and encode specific evidence, and to follow a rule (e.g., inclusion criteria) at a decision-point.
- (3) *Suppress [Sup.]* uncertainty by eliminating possibilities through intuition, guessing, or other procedures which accumulate unnecessary error. We only coded decisions as suppression when participants expressed that an alternative analysis path would contribute less error to their analysis.

These strategies entail different levels of justification for analytical decisions and tend to have a *temporal order* (Fig. 1, top). Imagine a researcher conducting a meta-analysis on the effect of mindfulness on depression. The researcher first

³https://github.com/kalealex/analysis_paths_research_synthesis/tree/master/analysis

decides to search the “gray literature” (i.e., unpublished or unconventional sources), orienting the scope of their search before engaging in targeted information retrieval, an example of the acknowledgement strategy. Decisions which acknowledge uncertainty often steer research broadly, surveying possible analysis paths and associated trade-offs prior to definitive choices about how to implement the analysis. Next, while reviewing the gray literature for their meta-analysis, the researcher decides to search unpublished dissertations for relevant data, an example of the reduction strategy. Decisions which reduce uncertainty use specific procedures to navigate analysis paths, sometimes pruning away possible paths by omission. Reduction often builds on acknowledgement by implementing the broad goals identified in acknowledgement through a specific approach to analysis. While still engaged in data collection for their meta-analysis on mindfulness, the researcher chooses not to examine the risk of bias in individual studies in their sample. The researcher knows this is not ideal but rationalizes this decision because they have limited time to conduct their analysis. Decisions which suppress uncertainty are often necessary but unjustifiable compromises in response to situational factors which are sometimes beyond the researcher’s control. Suppression substitutes for acknowledgment or reduction strategies.

We also coded aspects of decision context. We recorded the **stage in the research process** (e.g., *question formation, literature review, meta-analysis*) for each decision. This allowed us to compare the relative frequencies of different strategies at different stages in the research process, roughly following the way that Wicherts et al. [77] break up researcher degrees of freedom into different phases of analysis.

We coded the **reason given for each decision** (e.g., striving for *reproducibility, limited availability of information, standard practices*). Sometimes these reasons were not explicitly stated, and we had to infer them from the broader context of our discussions with participants. Coding the reasons for each decision allowed us to examine what factors motivated participants to engage in each strategy.

We extended the notion of the garden of forking paths to **communicative and organizational decisions**. We coded a distinction between decisions which have a *direct impact* on the final written report of results and decisions which have an *indirect impact* on results through interpretive and communicative aspects of the scientific process. This distinction allowed us to highlight analytical decisions with impacts that may be difficult or impossible to quantify.

We also coded **metadata** such as the *participant*, their *field of study*, and the *goal* of their project. Lastly, we noted when a decision was associated with a particular *need for software features* or a particular *threat to validity*. These passages informed our discussion of challenges and opportunities in designing for the garden of forking paths.

4 RESULTS

Figure 2: Frequencies of analytical decisions and percentages of strategies coded at each stage of the research process.

Decision-Making Strategies

We coded 480 analytical decisions in our interviews. The majority were decisions to reduce uncertainty (297 of 480 decisions; 61.9%), followed by decisions to acknowledge uncertainty (139 of 480 decisions; 29.0%) and suppress uncertainty (44 of 480 decisions; 9.2%).

Acknowledgement. *Acknowledging uncertainty* was most prevalent early in the research process when researchers define questions and objectives (24 of 48 decisions; 50.0%) (Fig. 2). Participants identified this early scoping of the review as very important to a straightforward and systematic analysis. “If I spend time thinking about and writing down what my population, intervention, etc. is for my question, then that makes my literature review that much more efficient. Because then, as I run across a new study, it should be that much easier to say, ‘Is it in or is it out? Does it give me one of my outcomes that I’ve specified?’ If it doesn’t, then it’s not included. It prevents you from getting mired down in indecision.” [Ack.] (P3). Multiple researchers described a similar practice of considering possible analysis paths early on and writing a scope description to guide later analytical decisions. Although researchers often rely on frameworks (e.g., PICOTS, described by Nelson [58])

to guide scope development, links between scope and subsequent decisions are maintained in digital or paper notes if they are documented at all.

We also see relatively more *acknowledgement* during literature review (37 of 117 decisions; 31.6%), quantitative analysis (10 of 30 decisions; 33.3%), and communication (39 of 119 decisions; 32.8%) than in data collection (17 of 85 decisions; 20.0%) and meta-analysis (12 of 72 decisions; 16.7%) (Fig. 2). Researchers acknowledge uncertainty in later stages of research mostly to confront irreducible sources of uncertainty. For example, one researcher in the Navy described using caveats to qualify information gathered through interviews with Navy personnel. *“Usually it ends up being a time issue or access. [We] did not have time to go out and validate that this is how they actually perform the work on the job. We just take people’s word for it.”* [Ack.] (P2). In a similar case of acknowledgement to avoid potential misunderstandings of evidence, another researcher described checking the quality of available data. *“I’ll use data visualization as I’m going through the process to see if there are things that look really weird. Like, I’ll make a little ggplot and look to see if things are falling in the range that I expect them to.”* [Ack.] (P9). When available information or resources constrain analysis paths, researchers tend to rely on acknowledgement to communicate limitations or check impacts on data quality.

Reduction vs Suppression. The high frequency of strategies to *reduce* uncertainty at every stage in analysis (Fig. 2) suggests that researchers often employ rule-based reasoning when implementing their scientific review and analysis. These rules are a mix of standard practices and lab-specific procedures. For example, one researcher studying behavioral treatments for PTSD augmented the common practice of meta-analyzing only between-subjects experiments by looking separately at within-subjects evaluations to validate the treatment effect within an individual. *“Most meta-analyses only look at the between-group [studies], and they just assume that there is sufficient within-subject change to warrant doing any of it. I don’t think that is as useful. That’s the other big thing we are adding with this meta-analysis are these within-subjects tests to contextualize the between-group [effect]. I’ve never seen a meta-analysis in my area that does both, but systematic reviews do.”* [Red.] (P10). All researchers expressed awareness of best practices and deviated from them at times, but researchers had different attitudes about when and how much it was appropriate to bend the rules.

Sometimes researchers have little choice but to follow a path which they know is not ideal. For example, a Navy scientist measuring in-flight blood-oxygen levels to fill a gap in existing evidence described a decision to use finger-mounted monitors rather than more precise head-mounted monitors. *“Constraints of the experiment make it so that you can’t get*

big head-mounted monitors and have a pilot wear their helmet at the same time. You have to make concessions where you can and realize that there’s going to be some variability and error in your data just based on where your monitors are at.” [Sup.] (P5). The participant went on to describe how they would acknowledge this measurement error in their written report. When the best option available to researchers adds a source of error to analysis, the line between strategies of reduction and suppression is blurred, and researchers need to document how their analysis is constrained in order to alert stakeholders to potential *suppression* of uncertainty.

Although we code *suppression of uncertainty* infrequently (44 of 480 decisions; 9.2%), we suspect that suppression is underrepresented in our sample because researchers may not recognize or admit when decisions introduce greater error than other viable alternatives. For example, *“We are for sure not including any qualitative research. The only two outcome measures we are interested in are exam score and failure rate. So maybe active learning changed how students feel about [the classroom climate], and maybe they have qualitative data from a survey. We are not including that either, and that’s driven completely by our research question.”* [Red.] (P4). Although the researcher attributed the decision to ignore qualitative work to their research question, earlier in our interview they described potential biases in the framing of their research question. *“Before we started this meta-analysis, we wondered if maybe class size would be an explanatory variable in how well active learning works, or subject area, or whether it’s an intro class or an upper-division class. Things like that, our relatively small research team came up with those things, so there’s totally a bias present in the things that we omitted and the things that we included.”* [Sup.] (P4). Researchers need ways to represent and keep track of the reasoning behind analytical decisions in part because, as this researcher put it, *“Defining where the personal bias is coming in and where specifically it’s problematic would be important.”* (P4). We argue that helping researchers document sources of bias, and how they influence analytical decisions throughout a research project, would serve to identify situational sources of uncertainty which are opaque in current research practice.

Communicative and Organizational Practices

We made a distinction between decisions which have *direct impact* on the content presented in the written report of findings and decisions which have *indirect impact* on findings. Decisions with indirect impact do not change the quantitative or qualitative evidence presented in the written report of findings, but they change the way evidence is recorded in work documents and framed in meetings with collaborators or presentations to stakeholders. These social and communicative aspects of the research process have the potential to impact how findings and recommendations are interpreted.

Decisions with *indirect impact* rely more on *acknowledging* uncertainty (79 of 201 decisions; 39.3%) than decisions with direct impact (60 of 279 decisions; 21.5%). Often these are decisions about how to organize or manage a review. “We create a timeline and meet once a week to share what we’ve learned... and see where we have some challenges that we need to work on. If we’ve got to jump in and help someone else, then we’ll do that too.” [Ack.] (P6). Decisions to discuss issues with collaborators often serve to acknowledge possible ways of handling subsequent decisions with direct impact, such as how to sample the literature. “When the [search] term isn’t clear, it’s a nightmare... That’s sort of this ongoing iterative process, and honestly I don’t think there’s a way around it... The conversations with collaborators have been the best way to [narrow the scope] in my experience.” [Ack.] (P9).

Other times, decisions with *indirect impact* are about checking the evidence in the review to preempt potential problems later on. “Before the literature search is a tenth over, you should revisit your template against the literature you’ve examined. You can double-check your template [for the evidence table] to make sure it’s right, and you adjust it.” [Ack.] (P7). By checking that the columns of the evidence table adequately reflect important themes in the literature, the researcher ensures that their review is comprehensive. Decisions with indirect impact are scattered across notes, spreadsheets, and personal correspondences, yet they play an important role in shaping the scientific process behind the scenes.

In contrast, decisions with *direct impact* tend to rely more on the *reduction* strategy (190 of 279 decisions; 68.1%) than decisions with indirect impact (107 of 201 decisions; 53.2%). For example, one participant described a relaxed approach to systematic review which did not rely on evidence tables. “We were kind of looking for whether a study was really sound or not. Some were more easily brushed aside because it seemed like their research design was not something we would have followed, so we trusted each other to weed some of those out early on.” [Red.] (P6). This researcher’s team excluded and reviewed studies outside the context of a consistent template. When asked if they would use an evidence table, this participant said, “If there was a spreadsheet already created, and I could tailor it to what might be most important for our particular project, then I think I would use that. But as far as using it every time and pulling out every piece of [information], I would literally have to have an intern do that for me because I don’t have time.” (P6). Shortcuts, such as forgoing dual-review and evidence tables, are essential to researchers in the Navy because they operate on tight timelines.

Motivations

During our qualitative analysis we inductively coded 21 different reasons for researchers’ decisions. We then grouped

Figure 3: Categorized reasons for analytical decisions.

these 21 reasons into six themes which capture the motivations behind researchers’ decision-making (Fig. 3).

- **Principles** (93 of 480 decisions; 19.4%): ideals which researchers adhere to such as *comprehensiveness* of the review, *reproducibility* in judgments, *consistency* in how evidence is evaluated.
- **Social Factors** (55 of 480 decisions; 11.5%): strictly communicative or interpersonal influences on decision-making such as a desire for *ease of understanding* or building *confidence through consensus*.
- **Domain-Specific Factors** (67 of 480 decisions; 14.0%): concerns related to the domain under study such as cases where the researcher relies on *domain knowledge* or seeks *conceptual clarity* regarding key constructs.
- **Conventions** (73 of 480 decisions; 15.2%): *standard practices* which often reflect established knowledge about how to deal with issues like *statistical power*.
- **Self-Imposed Constraints** (107 of 480 decisions; 22.3%): factors which are under the control of the individual researcher such as *research questions*, *preferences*, or acting on a sense of *caution* or a *lack of knowledge* about research procedures.
- **External Constraints** (85 of 480 decisions; 17.7%): factors beyond the researchers’ control such as *limited availability of information*, *limited time and attention* for the review, *limited control of research objectives*, or *limited alternatives* to the chosen analysis path.

In our interviews, decisions to *acknowledge* uncertainty were most often motivated by principles (46 of 139 decisions; 33.1%), domain-specific factors (32 of 139 decisions; 23.0%),

and self-imposed constraints (21 of 139 decisions; 15.1%). This suggests that acknowledging uncertainty is often voluntary and carefully thought out on the part of researchers. Researchers are particularly motivated by the principle of comprehensiveness (75 of 93 decisions; 80.6%), often reporting that they provide detailed background information in their written report, check data quality, and code possible covariates in their evidence tables. These practices serve the purpose of a complete and thorough review.

The motivations for decisions to *reduce* uncertainty are fairly balanced, with self-imposed constraints (81 of 297 decisions; 27.3%) and conventions (60 of 297 decisions; 20.2%) being most prevalent. By analyzing participants' motivations, we find that not all decisions to reduce uncertainty are equally justified, and this is further evidence of a blurred line between reduction and suppression strategies. For example, one participant described flexibility in defining the search terms used to sample the literature. *"It's helpful when you have a sense of how many studies you're talking about. In the psychiatric conditions, there was a study looking at anger. Anger is not exactly—there's not a diagnosis associated with that. Initially, I wanted to include it, but there's only a couple of studies for which that is relevant, so it seemed like it was cleaner to just exclude anger because it's not a diagnosis. It's just two studies; it's not a big deal... We don't want to do a systematic review with two studies; we don't want to do a systematic review with 100 studies because it's probably not targeted enough. It seemed like some of that process was finding a happy medium."* [Red.] (P9). By sampling in order to hit a target sample size, the researcher may ignore constructs which are theoretically important but hard to measure.

In contrast, another researcher described a decision to *include a group with two studies in a hierarchical meta-analysis* (P3) (i.e., meta-regression) in order to address their research question. Later, when speaking in general terms, they said that using groups of studies which are too small in meta-regression leads to unrepresentative results and poor statistical power. These decisions about how to group studies highlight a trade-off between conventions—adhering to statistical best practices by having large enough samples in each group—and self-imposed constraints—evaluating a specific research question by using a certain grouping factor in a statistical model. The frequency of external constraints as a motivator for suppression (33 of 44 decisions; 75.0%) and the thin line between suppression and reduction strategies suggest that researchers would benefit from documenting the reasoning behind their analytical decisions.

5 DESIGN FOR THE GARDEN OF FORKING PATHS

Our findings point to challenges in designing interactive systems to support research synthesis which faithfully represent

the relevant empirical evidence while accounting for organizational and individual constraints and goals in the synthesis process. Specific challenges that emerge from our results include: 1) a tension in current practice between researchers' strong desire to consider multiple possible analysis paths, as evidenced by the prevalence of acknowledgement, and the frequent influence of self-imposed constraints in motivating the reduction of uncertainty; 2) a lack of consistent representations and processes for capturing their rationales for decisions, limiting researchers' ability to reflect on their rationales; and 3) a trade-off between researchers' desire to acknowledge uncertainty (e.g., by stating caveats) and the need to tell a clear story with their analysis (e.g., filtering evidence based on their research question).

Research Synthesis as Expected Utility Maximization

The inherent subjectivity in scoping a systematic review or meta-analysis, combined with the influence of external constraints, suggests that the single "best" analysis may not always be the analysis that is chosen in an organization. To guide reflection on how interactive software might better support researchers' endeavors, we frame research synthesis as a process aimed at *maximizing the expected utility of analysis* as perceived by the researcher within an organization.

An expected utility framework assumes that a person will decide between multiple actions by reasoning about the expected consequences of each in combination with the value they place on those consequences, which is defined by a utility function [75]. While the consequences of actions (e.g., betting strategies) in a classic example of utility are monetary payoffs, we suggest that an adaption of utility theory for research synthesis might frame the relevant consequence of an action, or a choice of analysis, as the accuracy or faithfulness with which the analysis captures the true state of the world. Since the true state of the world is unknown, there is uncertainty about the faithfulness of any analysis. In research synthesis perceived uncertainty takes the form of the researcher's subjective sense of confidence in the accuracy of possible analyses. The criterion for choosing between actions is the average utility of each action given the uncertainty.

Under utility theory a decision-maker should be able to choose the higher utility alternative given any pair of actions (or alternatively, should be able to declare them equal in utility) [36, 75]. This assumption implies that interactive systems for research synthesis could help researchers choose the paths that maximize their expected utility by helping them to 1) accurately perceive the consequences of different analysis paths (i.e., actions). Because the true state of the world, and thus the exact consequences of an analysis, may be impossible to judge, researchers may use heuristics such as looking at the results of different paths or considering competing values to make decisions. Similarly, interactive

systems could also support researchers directly in 2) identifying how much they value the accuracy of different paths.

Acknowledging Alternative Analyses

Challenge: Trying a set of analyses instead of just one. Our interviews suggest that one of the core challenges in designing for research synthesis is supporting judgments about when a given analysis is justified. The pitfall for researchers is the thin line between exploratory analysis and using researcher degrees of freedom to select an analysis path which confirms a particular hypothesis or point of view. An episode from one of our interviews illustrates this ambiguity: “Through those discussions that we have, we’ll have suggestions for different types of analysis we can conduct that might split the hair a little bit differently and give us additional information that we can then decide, ‘Okay, what is the best way to report this now that we’ve looked at it both of these ways?’” (P6). In cases like this, we argue that any reduction strategy used to select a particular analysis may actually be suppressing uncertainty because presenting one analysis path among multiple paths observed may reflect confirmation bias as much as the true state of the world.

Prior work on reliable statistics [67, 69] suggests that researchers should survey possible analysis paths and multiplex across them by exploring *and reporting* a reasonable subset of analyses. Because comprehensiveness was the most prevalent motivation for analytical decisions in our analysis, we suspect that researchers are naturally incentivized to explore possible analyses and that they will do so to the extent that tools make it feasible. Being able to directly evaluate how different paths impact results would provide researchers with relevant information for determining what paths have higher utility given their various motivations and constraints. For example, when the differences in results between paths are minimal, the researcher may decide to choose a path that “trades-off” accuracy with other desirable properties like minimizing time or resources spent, or maximizing consensus or interpretability⁴. When results are highly variable, the researcher should exercise extreme caution and consider reporting multiple analyses if possible.

Of course, researchers may not always be motivated report on multiple analyses paths if they prefer a certain result, raising the ethical dilemma that developing tools which make it easier to explore alternative analyses may facilitate the cherry-picking of results. The need to explore multiple analyses and the risk that this multiplexing will facilitate cherry-picking remains a tension in the broader set of recommendations in the literature on reducing bias in analysis.

⁴An alternative framing of decisions about paths in our study results could define utility as a function of path accuracy alone, and choices of non-optimal paths as “satisficing” under bounded rationality [66]

Opportunity: Multiverse analysis. Each possible analysis in the garden of forking paths produces a distribution of estimates which represents error in the analysis process. This is distinct from uncertainty about which path most accurately represents the evidence (i.e., the true underlying effect). Prior work [67, 69] suggests that, ideally, researchers should quantify and report uncertainty about analytical decisions by running a subset of possible analyses to see which decisions impact results, a procedure called multiverse analysis. Similar work on model comparison techniques (e.g., [51, 53, 61]) suggests that researchers should build and compare multiple models expressing assumptions of varying strength in order to separate uncertainty of evidence from uncertainty of assumptions in their analysis. These convergent lines of research suggest an opportunity for interactive systems to elicit a set of analysis paths under consideration and support researchers in interactively comparing outcomes from multiple quantitative analyses.

Opportunity: Visualizing the garden of forking paths. In order to support researchers in exploring, comparing, and reporting on multiple possible analyses, interactive systems need an explicit way to represent analytical decisions and elicit information about possible analyses from researchers. Based on prior work on interactive visualization of scientific workflows [12], we point out the opportunity to represent the garden of forking paths using interactive diagrams which help researchers map out decision-points and their influence on one another as nodes and edges. We propose other opportunities related to reasoning with and communicating uncertainty which build on this representation.

Representing Reasoning About Analysis

Challenge: Shifting attention from rationales to impacts. When prompted to describe the reasoning behind analytical decisions, researchers tend to focus on factors which rationalize their choices, often appealing to standard practices and research questions rather than examining the impacts of their choices on the results of analysis. For instance, we compare two cases where researchers made opposite decisions about whether or not to include a group of two studies in a meta-regression (see Results: Motivations). Both researchers were aware that two studies would not give them enough statistical power to make a precise effect size estimate, but instead they rationalized their decisions by appealing to conventions (P9) and research questions (P3). We observe that researchers document their rationales inconsistently across personal correspondences, notes, and work documents. This may contribute to difficulty weighing these motivations alongside the outcomes of alternative analysis paths. Considering rationales in the absence of information about outcomes leads to a sense of utility that is driven primarily by the perceived

value of a choice independent of its consequences. We argue that this problem would be mitigated by interactive systems that formally represent researchers' reasoning about analysis paths and attempt to shift researchers' attention to the impacts of decisions on the results of analysis.

Opportunity: Attributing rationale. While researchers often have a rationale for their decisions, they seem to lack the tools to externalize and review the motivations and constraints which shape their analysis. For example, one researcher expressed the need for a way to track conceptual uncertainty when coding evidence tables in Excel. *"If [a flag or annotation is] associated with the value in a cell, I think that would help coders feel more confident even if you're not ultimately going to do anything about it. And it might help your quality control checks at the end."* (P4). In agreement with this researcher, prior work on how uncertainty is represented in visual analytics systems [62] suggests that helping users maintain awareness of sources of uncertainty improves the calibration of confidence in the accuracy of the analysis. Given an interactive visualization of the garden of forking paths, a system could enable researchers to create custom flags associated with factors which influence analytical decision-making such as rules, assumptions, and conventions, as well as guidelines (e.g., scope descriptions and research questions) and constraints (e.g., limited time and attention). Researchers could use these flags to represent their reasoning at a given decision-point by attributing their choice to a set of rationales. We speculate that mapping out rationales for alternative analyses might highlight trade-offs between different motivations and thus prompt researchers to reflect on and potentially update their decisions.

Opportunity: Aligning subjective and statistical uncertainty. Since the true state of the world is unknown, researchers must rely on their subjective sense of the accuracy of analyses (i.e., confidence) to decide between alternative paths. However, some of the rationales given by researchers (e.g., conventions) are essentially appeals to authority, suggesting that researchers sometimes feel uncomfortable with their ability to judge the consequences of alternative analyses. Prior work [33] shows that comparing predicted effect size distributions to observed outcomes helps some form expectations about effects that better align with statistical uncertainty. Expressing subjective expectations promotes more active reasoning [41, 42], and using visual representations offloads information from working memory [18, 57], freeing up attention for metacognitive reflection [63] about how effects are understood. This suggests an opportunity to help researchers calibrate their expectations about the impacts of analytical decisions through comparisons of subjective uncertainty, elicited by a system using graphical or other formats (e.g., [28, 33, 41–43]), and statistical uncertainty, achieved

through computation. Based on prior work [15], we speculate that the experiential learning that occurs through such prediction can ultimately help foster confidence as well as accuracy in one's predictions.

Communicating Uncertainty

Challenge: Varying tolerance for uncertainty. Researchers we interviewed tend to remain skeptical about their findings, but they often need to present findings in a way that offers convincing support for recommendations. Contributing to this tension, it has been argued that decision-makers [4, 8, 11, 52, 70, 78, 79] (and people in general [21, 22, 25]) have limited tolerance for uncertainty. *"They want a visualization that shows gross effect, it gets down to the point, so they are no longer wanting a string of visualizations to illustrate every point. If you write about the methodology, you write about your data crunching, and then you show your visualization for your final result, that is what decision-makers are hungry for and are expecting now... They do care about the fidelity of the data, they just don't want a chart on it."* (P11). In practice, this Navy researcher told a simplified and compelling story advocating for a promising training program. They emphasized the trade-off between downplaying uncertainty to get decisions made and doing just the opposite when there was concern about the safety of Navy personnel. *"If it's a really important decision that you need to impress upon them, that there is a significant thing to consider, then I will sometimes do the findings a disservice by presenting the problem and burying the potential compelling use case or storyline."* (P11).

Opportunity: Visualizing possible outcomes. We argue that interactive systems for research synthesis should provide a set of visualizations for distributions of possible outcomes, which attempt to alleviate specific aversions that decision-makers have toward uncertainty. One major aversion to uncertainty is that many people, even experts [7, 68], find it hard to understand [37]. Prior work in psychology [29, 30], statistical reasoning [27, 28, 31], and data visualization [23, 33, 34, 38, 39] suggests a remedy to this problem: people reason about uncertainty most accurately when it is framed as frequencies of events, rather than probabilities or summary statistics. As such, visualizations of *quantitative uncertainty* should convey the possibility of multiple outcomes using frequency framing to circumvent misunderstandings of uncertainty. When it is important to convey uncertainty in possible outcomes with high fidelity, hypothetical outcome plots [34, 38, 43] and quantile dotplots [23, 33, 39] are valuable visualization formats. By presenting discrete outcomes, these formats align with results from prior work in statistics pedagogy suggesting that people develop better statistical intuitions via simulations [13, 14, 20, 35, 68].

In contrast, when decision-makers have low tolerance for uncertainty because they want a clear yes-or-no answer, visualizations could be designed to emphasize modal results while still displaying uncertainty, as supported by most static representations of distributions, including intervals. Additionally, propagating uncertainty in the effect distribution to derived measures that may be more closely aligned with the decision-maker's utility function, such as money or time saved, might help decision-makers appreciate uncertainty information for the purpose of making more informed decisions. *"There are these situations where you face the moral quandary of having to tell a company that you wasted their time [on inconclusive results] or giving them information that isn't very useful to them [because the analysis misrepresents available evidence to give an exaggerated sense of certainty]."* (P5). We argue that providing researchers with communication techniques which counteract decision-makers' specific aversions to uncertainty would mitigate the misconception that an uncertain result is not presentable and does not contain useful information. Of course, the best way to identify the impact of the techniques we propose is to evaluate decision-making given different communication strategies.

Opportunity: Conveying qualitative uncertainties. In our interviews we find that researchers express *qualitative* forms of uncertainty such as the assumptions and constraints behind their analysis by writing caveats in limitations sections or preparing supplemental presentation slides. Descriptive accounts of uncertainty help researchers align the expectations of decision-makers with the quality of available scientific evidence. In the words of one researcher, *"There is no such thing as a perfect study, but maybe if there was some way to make that information more readily available to researchers so that I clearly see a trend in limitations that are going on in this area of research. How am I not going to let that affect mine, or how can I at least convey that skepticism to my potential employer so that they know what might impact their results?"* (P5). Prior work [52, 53] suggests that uncertainty about assumptions is often overlooked by decision-makers, leading to an exaggerated sense of certainty about analytical results. Drawing on the opportunities we identify for representing reasoning about analytical decisions, a system could use rationales and their associations with decisions to create summary representations of assumptions, constraints, and limitations which should inform the decision-maker's sense of utility. Future work should explore the design space for representing and communicating these qualitative uncertainties.

Limitations

Limitations of our procedure for sampling analytical decisions introduce imprecision into the frequencies of decision-making strategies presented in our results. The sample of

researchers we interviewed may not be representative of the broader population of researchers working in applied settings. Additionally, we rely on each participant's account of whether decisions are justified to inform our coding of suppression. While this means that our data reflect the perspectives of our participants rather than our opinions, it also means that suppression may be relatively underrepresented in our sample because researchers may feel disinclined to admit to choosing less-than-ideal analysis paths.

We draw on expected utility theory to provide a basis for identifying "better" ways of supporting research synthesis despite the inevitability of organizational constraints the process. However, our treatment is not intended to be a formal model. How to characterize research synthesis as it occurs in organizations using theories like expected utility maximization or bounded rationality is left to future work.

6 CONCLUSION

We present a qualitative analysis of how researchers conducting applied research synthesis navigate the garden of forking paths: a series of analytical decision-points, each of which has the potential to influence findings. Based on our interviews and analysis, we identify a set of design challenges around making it more feasible for researchers to try and report multiple analyses, shifting researchers' attention from rationales for decisions to impacts of decisions, and supporting uncertainty communication techniques which address specific aversions to uncertainty among decision-makers. Considering evidence from our interviews in light of prior work, we point out opportunities for interactive systems to support research synthesis by helping researchers map the garden of forking paths, document their reasoning about analysis paths, and effectively communicate uncertainties impacting their analytical decision-making.

7 ACKNOWLEDGEMENTS

This work was funded by US Navy STTR Contract N68335-17-C-0410 in partnership with Stottler Henke Assoc, and NSF award #1749266.

REFERENCES

- [1] B. F. Anderson, D. H. Deane, K. R. Hammond, and G. H. McClelland. 1981. *Concepts in judgment and decision research*. Praeger, New York.
- [2] H Arksey and L O'Malley. 2005. Scoping Studies: Towards a Methodological Framework. *International Journal of Social Research Methodology* 8, 1 (2005), 19–25.
- [3] Kenneth J. (Kenneth Joseph) Arrow. 1965. Aspects of the theory of risk-bearing. (1965).
- [4] Maya Bar-hillel and Efrat Neter. 1993. How Alike Is It Versus How Likely Is It: A Disjunction Fallacy in Probability Judgments. *Journal of Personality and Social Psychology* 65, 6 (1993), 1119–1131.
- [5] Leon Bax, Ly Mee Yu, Noriaki Ikeda, and Karel G.M. Moons. 2007. A systematic comparison of software dedicated to meta-analysis of

- causal studies. *BMC Medical Research Methodology* 7, February 2007 (2007). <https://doi.org/10.1186/1471-2288-7-40>
- [6] Leon Bax, Ly Mee Yu, Noriaki Ikeda, Harukazu Tsuruta, and Karel G.M. Moons. 2006. Development and validation of MIX: Comprehensive free software for meta-analysis of causal research data. *BMC Medical Research Methodology* 6, 50 (2006), 1–11. <https://doi.org/10.1186/1471-2288-6-50>
  - [7] Sarah Belia, Fiona Fidler, Jennifer Williams, and Geoff Cumming. 2005. Researchers misunderstand confidence intervals and standard error bars. *Psychological methods* 10, 4 (2005), 389.
  - [8] Silvia Bonaccio and Reeshad S Dalal. 2006. Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes* 101 (2006), 127–151. <https://doi.org/10.1016/j.obhdp.2006.07.001>
  - [9] M Borenstein, L Hedges, J Higgins, and H Rothstein. 2005. *Comprehensive Meta-Analysis 2*. Englewood, NJ, Biostat. (2005).
  - [10] Nadia Boukhelifa, Marc-Emmanuel Perrin, Samuel Huron, and James Eagan. 2017. How Data Workers Cope with Uncertainty : A Task Characterisation Study. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2017). <https://doi.org/10.1145/3025453.3025738>
  - [11] David V Budescu and Adrian K Rantilla. 2000. Confidence in aggregation of expert opinions. *Acta Psychologica* 104 (2000), 371–398.
  - [12] Steven P Callahan, Juliana Freire, Emanuele Santos, Carlos E Scheidegger, Cláudio T Silva, and Huy T Vo. 2006. VisTrails: visualization meets data management. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*. ACM, 745–747.
  - [13] Beth Chance, Robert del Mas, and Joan Garfield. 2004. Reasoning about sampling distributions. In *The challenge of developing statistical literacy, reasoning and thinking*. Springer, 295–323.
  - [14] B Chance, J Garfield, and B delMas. 1999. A model of classroom research in action: Developing simulation activities to improve students’ statistical reasoning. *52nd Session of the International Statistical Institute, Helsinki, Finland* (1999).
  - [15] Beth Chance, Joan Garfield, and Robert delMas. 2000. Developing Simulation Activities To Improve Students’ Statistical Reasoning. (2000).
  - [16] The Cochrane Collaboration. 2014. Review Manager (RevMan) 5.1.0. Copenhagen: The Nordic Cochrane Centre. (2014).
  - [17] Harris Cooper, Larry V. Hedges, and Jeffrey C. Valentine. 2009. *The Handbook of Research Synthesis and Meta-Analysis*. Russell Sage Foundation.
  - [18] Richard Cox. 1999. Representation construction, externalised, cognition and individual differences. *Learning and Instruction* 9, 4 (1999), 343–363. [https://doi.org/10.1016/S0959-4752\(98\)00051-6](https://doi.org/10.1016/S0959-4752(98)00051-6)
  - [19] J. W. Creswell and C. N. Poth. 2018. *Qualitative inquiry & research design: Choosing among five approaches*. SAGE Publications, Inc.
  - [20] Geoff Cumming and Neil Thomason. 1998. Statplay: Multimedia for statistical understanding, in pereira-mendoza (ed. In *Proceedings of the Fifth International Conference on Teaching Statistics, ISI*. Citeseer.
  - [21] Shawn P. Curley and J. Frank Yates. 1985. The Center and Range of the Probability Interval as Factors Affecting Ambiguity Preferences. *Organizational Behavior and Human Decision Processes* 36 (1985), 273–287.
  - [22] Hillel J Einhorn and Robin M Hogarth. 1985. Ambiguity and Uncertainty in Probabilistic Inference. *Psychological Review* 92, 4 (1985).
  - [23] Michael Fernandes, Logan Walls, Sean Munson, Jessica Hullman, and Matthew Kay. 2018. Uncertainty Displays Using Quantile Dotplots or CDFs Improve Transit Decision-Making. In *Conference on Human Factors in Computing Systems - CHI '18*. <https://doi.org/10.1145/3173574.3173718>
  - [24] R Ganann, D Ciliska, and H Thomas. 2010. Expediting systematic reviews: methods and implications of rapid reviews. *Implementation Science* (2010), 5–56. <https://doi.org/10.1186/1748-5908-5-56>
  - [25] P. Gardenfors and N. Sahlin. 1983. Decision making with unreliable probabilities. *Brit. J. Math. Statist. Psych.* 36 (1983), 240–251.
  - [26] Andrew Gelman and Eric Loken. 2014. The statistical crisis in science. *American Scientist* 102, 6 (2014). <https://doi.org/10.1511/2014.111.460>
  - [27] Gerd Gigerenzer and Ulrich Hoffrage. 1995. How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review* 102 (1995), 684–704.
  - [28] Daniel G Goldstein and David Rothschild. 2014. Lay understanding of probability distributions. *Judgment and Decision Making* 9, 1 (2014), 1.
  - [29] L Hasher and R T Zacks. 1984. Automatic processing of fundamental information: the case of frequency of occurrence. *The American psychologist* 39, 12 (1984), 1372–1388. <https://doi.org/10.1037/0003-066X.39.12.1372>
  - [30] Ralph Hertwig, Greg Barron, Elke U Weber, and Ido Erev. 2004. Decisions from experience and the effect of rare events in risky choice. *Psychological science* 15, 8 (2004), 534–539.
  - [31] U. Hoffrage and G. Gigerenzer. 1998. Using natural frequencies to improve diagnostic inferences. *Academic Medicine: Journal of the Association of American Medical Colleges* 73, 5 (May 1998), 538–540.
  - [32] Robin M Hogarth. 1987. Judgement and choice: The psychology of decision. (1987).
  - [33] Jessica Hullman, Matthew Kay, Yea-Seul Kim, and Samana Shrestha. 2018. Imagining Replications: Graphical Prediction & Discrete Visualizations Improve Recall & Estimation of Effect Uncertainty. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)* (2018). <http://idl.cs.washington.edu/papers/imagining-replications>
  - [34] Jessica Hullman, Paul Resnick, and Eytan Adar. 2015. Hypothetical Outcome Plots Outperform Error Bars and Violin Plots for Inferences about Reliability of Variable Ordering. *PloS one* 10, 11 (2015).
  - [35] D Mills Jamie. 2002. Using computer simulation methods to teach statistics: A review of the literature. *Journal of Statistics Education* 10, 1 (2002).
  - [36] Niels Erik Jensen. 1967. An introduction to Bernoullian utility theory: I. Utility functions. *The Swedish journal of economics* (1967), 163–183.
  - [37] D. Kahneman. 2011. *Thinking, fast and slow*. Farrar, Straus and Giroux, New York.
  - [38] Alex Kale, Francis Nguyen, Matthew Kay, and Jessica Hullman. 2019. Hypothetical Outcome Plots Help Untrained Observers Judge Trends in Ambiguous Data. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)* (2019). <http://idl.cs.washington.edu/papers/hops-trends>
  - [39] Matthew Kay, Tara Kola, Jessica Hullman, and Sean Munson. 2016. When (ish) is my bus? User-centered visualizations of uncertainty in everyday, mobile predictive systems. In *Proceedings of the 34th Annual ACM Conference on Human Factors in Computing Systems (CHI '16)*.
  - [40] S Khangura, K Konnyu, R Cushman, J Grimshaw, and D Moher. 2012. Evidence summaries: the evolution of a rapid review approach. *Syst Rev* (2012), 1–10. <https://doi.org/10.1186/2046-4053-1-10>
  - [41] Yea-Seul Kim, Katharina Reinecke, and Jessica Hullman. 2017. Explaining the Gap: Visualizing One’s Predictions Improves Recall and Comprehension of Data. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17* (2017), 1375–1386. <https://doi.org/10.1145/3025453.3025592>
  - [42] Yea-Seul Kim, Katharina Reinecke, and Jessica Hullman. 2018. Data Through Others’ Eyes: The Impact of Visualizing Others’ Expectations on Visualization Interpretation. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)* (2018). <http://idl.cs.washington.edu/papers/others-expectations>

- [43] Yea-Seul Kim, Logan Walls, Peter Krafft, and Jessica Hullman. 2019. A Bayesian Cognition Approach to Improve Data Visualization. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019).
- [44] George J. Klir and Bo Yuan. 1995. *Fuzzy Sets and Fuzzy Logic: Theory and Applications*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- [45] D Levac, H Colquhoun, and K K O'Brien. 2010. Scoping studies: advancing the methodology. *Implementation Science* 5, 69 (2010), 1–9. <https://doi.org/10.1186/1748-5908-5-69> arXiv:arXiv:1011.1669v3
- [46] M.W. Lipsey and D.B. Wilson. 2001. *Practical Meta-Analysis*. SAGE Publications.
- [47] Raanan Lipshitz and Orna Strauss. 1997. Coping with Uncertainty: A Naturalistic Decision-Making Analysis. *Organizational Behavior and Human Decision Processes* 69, 2 (1997), 149–163. <https://doi.org/10.1006/obhd.1997.2679>
- [48] John Lofland, David Snow, Leon Anderson, and Lyn H. Lofland. 2006. *Analyzing social settings: A guide to qualitative observation and analysis*. Wadsworth, Cengage Learning, Belmont=, CA.
- [49] K. R. MacCrimmon and D. A. Wehrung. 1986. *Taking risks*. Free Press, New York.
- [50] Alan M MacEachren, Anthony Robinson, Susan Hopper, Steven Gardner, Robert Murray, Mark Gahegan, and Elisabeth Hetzler. 2005. Visualizing Geospatial Information Uncertainty: What We Know and What We Need to Know. *Cartography and Geographic Information Science* 32, 3 (2005), 139–160. <https://doi.org/10.1559/1523040054738936>
- [51] Charles F. Manski. 2003. *Partial Identification of Probability Distributions: Springer Series in Statistics*. Springer.
- [52] Charles F Manski. 2018. Communicating uncertainty in policy analysis. *Proceedings of the National Academy of Sciences of the United States of America* (2018). <https://doi.org/10.1073/pnas.1722389115>
- [53] Charles F Manski. 2018. *The Lure of Incredible Certitude*. Working Paper 24905. National Bureau of Economic Research. <https://doi.org/10.3386/w24905>
- [54] James G March. 1976. Ambiguity and choice in organizations. (1976).
- [55] J. G. March and H. A. Simon. 1958. *Organizations*. Wiley, New York.
- [56] M Miles, M Huberman, and J Saldana. 2014. *Qualitative Data Analysis: A Methods Sourcebook* (3 ed.). SAGE Publications Inc., Thousand Oaks, CA, Chapter 2.
- [57] Hedwig M. Natter and Dianne C. Berry. 2005. Effects of active information processing on the understanding of risk information. *Applied Cognitive Psychology* 19, 1 (2005), 123–135. <https://doi.org/10.1002/acp.1068>
- [58] Heidi D. Nelson. 2014. *Systematic Reviews to Answer Health Care Questions*. Wolters Kluwer Health/Lippincott Williams & Wilkins, Philadelphia, PA.
- [59] Mourad Ouzzani, Hossam Hammady, Zbys Fedorowicz, and Ahmed Elmagarmid. 2016. Rayyan - a web and mobile app for systematic reviews. (2016). <https://doi.org/10.1186/s13643-016-0384-4>
- [60] Mai T. Pham, Andrijana Rajić, Judy D. Greig, Jan M. Sargeant, Andrew Papadopoulos, and Scott A. McEwen. 2014. A scoping review of scoping reviews: Advancing the approach and enhancing the consistency. *Research Synthesis Methods* 5, 4 (2014), 371–385. <https://doi.org/10.1002/jrsm.1123> arXiv:arXiv:1011.1669v3
- [61] Juho Piironen and Aki Vehtari. 2017. Comparison of Bayesian predictive methods for model selection. *Statistics and Computing* 27, 3 (01 May 2017), 711–735. <https://doi.org/10.1007/s11222-016-9649-y>
- [62] Dominik Sacha, Hansi Senaratne, Bum Chul Kwon, Geoffrey Ellis, and Daniel A. Keim. 2016. The Role of Uncertainty, Awareness, and Trust in Visual Analytics. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016). <https://doi.org/10.1109/TVCG.2015.2467591>
- [63] Gregory Schraw, Kent J. Crippen, and Kendall Hartley. 2006. Promoting self-regulation in science education: Metacognition as part of a broader perspective on learning. *Research in Science Education* 36, 1-2 (2006), 111–139. <https://doi.org/10.1007/s11165-005-3917-8>
- [64] R. Silberzahn, E. L. Uhlmann, D. P. Martin, P. Anselmi, F. Aust, E. Awtry, Š. Bahník, F. Bai, C. Bannard, E. Bonnier, R. Carlsson, F. Cheung, G. Christensen, R. Clay, M. A. Craig, A. Dalla Rosa, L. Dam, M. H. Evans, I. Flores Cervantes, N. Fong, M. Gamez-Djokic, A. Glenz, S. Gordon-McKeon, T. J. Heaton, K. Hederos, M. Heene, A. J. Hofelich Mohr, F. Högden, K. Hui, M. Johannesson, J. Kalodimos, E. Kaszubowski, D. M. Kennedy, R. Lei, T. A. Lindsay, S. Liverani, C. R. Madan, D. Molden, E. Molleman, R. D. Morey, L. B. Mulder, B. R. Nijstad, N. G. Pope, B. Pope, J. M. Prenoveau, F. Rink, E. Robusto, H. Roderique, A. Sandberg, E. Schlüter, F. D. Schönbrodt, M. F. Sherman, S. A. Sommer, K. Sotak, S. Spain, C. Spörlein, T. Stafford, L. Stefanutti, S. Tauber, J. Ullrich, M. Vianello, E.-J. Wagenmakers, M. Witkowiak, S. Yoon, and B. A. Nosek. 2018. Many Analysts, One Data Set: Making Transparent How Variations in Analytic Choices Affect Results. *Advances in Methods and Practices in Psychological Science* (2018). <https://doi.org/10.1177/2515245917747646>
- [65] Joseph P Simmons, Leif D Nelson, and Uri Simonsohn. 2011. False-Positive Psychology : Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science* 22, 11 (2011), 1359–1366. <https://doi.org/10.1177/0956797611417632>
- [66] Herbert A Simon. 1972. Theories of bounded rationality. *Decision and organization* 1, 1 (1972), 161–176.
- [67] Uri Simonsohn, Joseph P. Simmons, and Leif D. Nelson. 2015. Specification Curve: Descriptive and Inferential Statistics on All Reasonable Specifications. *SSRN* (Nov 2015). <https://doi.org/10.2139/ssrn.2694998>
- [68] Emre Soyer and Robin M Hogarth. 2012. The illusion of predictability: How regression statistics mislead experts. *International Journal of Forecasting* 28, 3 (2012), 695–711.
- [69] Sara Steegen, Francis Tuerlinckx, Andrew Gelman, and Wolf Vanpaemel. 2016. Increasing Transparency Through a Multiverse Analysis. *Perspectives on Psychological Science* 11, 5 (2016), 702–712. <https://doi.org/10.1177/1745691616658637>
- [70] Lyn M Van Swol and Janet A Snizek. 2005. Factors affecting the acceptance of expert advice. *British Journal of Social Psychology* 44 (2005), 443–461. <https://doi.org/10.1348/014466604X17092>
- [71] J Thomas, J Brunton, and S Graziosi. 2010. EPPI-Reviewer 4: software for metafor synthesis. EPPI-Centre Software. London: Social Science Research Unit, UCL Institute of Education. (2010).
- [72] Andrea C. Tricco, Jesmin Antony, Wasifa Zarin, Lisa Striffler, Marco Ghassemi, John Ivory, Laure Perrier, Brian Hutton, David Moher, and Sharon E. Straus. 2015. A scoping review of rapid review methods. *BMC Medicine* 13, 1 (2015). <https://doi.org/10.1186/s12916-015-0465-6> arXiv:arXiv:1011.1669v3
- [73] Amos Tversky and Daniel Kahneman. 1975. Judgment under uncertainty: Heuristics and biases. In *Utility, probability, and human decision making*. Springer, 141–162.
- [74] Wolfgang Viechtbauer. 2010. Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software* 36, 3 (2010), 1–48. <http://www.jstatsoft.org/v36/i03/>
- [75] John von Neumann, Oskar Morgenstern, and Ariel Rubinstein. 1944. *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton University Press. <http://www.jstor.org/stable/j.ctt1r2gkx>
- [76] A Watt, A Cameron, L Sturm, T Lathlean, W Babidge, S Blamey, K Facey, D Hailey, I Norderhaug, and G Maddern. 2008. Rapid reviews versus full systematic reviews: an inventory of current methods and practice in health technology assessment. *Int J Technol Assess Health Care* 24 (2008), 133–139. <https://doi.org/10.1017/S0266462308080185>

- [77] Jelte M. Wicherts, Coosje L.S. Veldkamp, Hilde E.M. Augusteijn, Marjan Bakker, Robbie C.M. van Aert, and Marcel A.L.M. van Assen. 2016. Degrees of freedom in planning, running, analyzing, and reporting psychological studies: A checklist to avoid P-hacking. *Frontiers in Psychology* 7, Nov (2016). <https://doi.org/10.3389/fpsyg.2016.01832>
- [78] Ilan Yaniv and Dean P Foster. 1995. Graininess of Judgment Under Uncertainty: An Accuracy-Informativeness Trade-Off. *Journal of Experimental Psychology: General* 124, 4 (1995), 424–432.
- [79] Ilan Yaniv and Dean P. Foster. 1997. Precision and Accuracy of Judgmental Estimation. *Journal of Behavioral Decision Making* 10 (1997), 21–32.