Running head: SOUND INFERENCE IN COMPLICATED RESEARCH

Sound Inference in Complicated Research: A Multi-Strategy Approach

Sanjay Srivastava

University of Oregon

Draft date: November 28, 2018

Author Note

This is a working paper. Feedback is welcome.

The author was supported by grants from the National Institutes of Health (MH106879) and the National Science Foundation (1551817). Helpful comments on previous drafts were provided by Jan Vanhove and Simine Vazire.

Correspondence may be addressed to Sanjay Srivastava, Department of Psychology, 1227 University of Oregon, Eugene, OR 97403-1227, USA; Email sanjay@uoregon.edu; Twitter: @hardsci.

Abstract

Researchers are increasingly paying attention to practices that can capitalize on chance. Capitalizing on chance can lead to unreplicable results in the form of false-positive tests, biased estimates, or overfit models. Preregistration of analysis plans, where a researcher anticipates all analysis decisions in advance and makes a plan for each one, has been promoted as one way to make inferences more sound and findings more replicable. Preregistration works by creating *decision independence*: analytic decisions are the same regardless of the potentially spurious features of the specific dataset being analyzed, instead of being overfit to them. But preregistration can be difficult in practice for some complicated research paradigms, including longitudinal studies, statistical modeling, machine learning, and other large multivariate designs and analyses where key decisions may be difficult to anticipate or make in advance. When simple preregistration is not practical, other strategies that create decision independence can be incorporated into it or used in place of it. This manuscript describes six such strategies: standardization, blind analysis, data partitioning, supporting studies, coordinated analysis, and multiverse analysis. All of them can be used to create adaptive preregistrations, and several can also be used when decisions arise that were not anticipated. Used together, they can form an integrative, multi-strategy approach to drawing sound inferences in a wide range of research.

Sound Inference in Complicated Research: A Multi-Strategy Approach

In recent years, psychology has seen greater attention to the problem of capitalizing on chance – identifying and interpreting as meaningful patterns what are actually spurious features of one specific dataset. Capitalizing on chance can lead to replicability problems in the form of false-positive test results, biased estimates, misleading fit statistics, and misspecified models. A variety of terms, some old and some new, refer to a set of interrelated research practices that increase the risk of capitalizing on chance: these include p-hacking, flexible data analysis, researcher degrees of freedom, the garden of forking paths, and questionable research practices.

Researchers are increasingly recognizing that this does not require intention or awareness; rather, capitalizing on chance can result from researchers' well-intentioned and necessary aims of making sense of their data. Human beings are voracious pattern-recognizers who can sometimes overreach, whether it is seeing images in clouds or seeing meaningful effects in noisy data. Rigorous research benefits from safeguards against seeing things that aren't there.

A practice that can help protect against capitalizing on chance is preregistration of analysis plans (Nosek, Ebersole, DeHaven, & Mellor, 2018). Currently popular websites and templates for preregistration cater to common research designs like clinical trials and factorial designs. Templates help researchers anticipate all of the important decisions that will need to be made and make principled plans without capitalizing on chance features of data.

But for researchers who do more complex research, preregistration can be difficult to implement in practice. In longitudinal studies, large surveys, statistical modeling, machine learning, and other kinds of big multivariate designs and analyses, it may be difficult to make a complete plan in advance or even to anticipate all of the decisions that will need to be made. In

addition, sometimes optimal decisions require interacting with the data. Simple, template-based preregistration may seem ill-suited to such applications.

The goal of this manuscript is to discuss a multi-strategy approach for researchers who do complicated research. The approach includes preregistration, but it both expands on it and incorporates other research practices that complement it. The structure of the paper is as follows: I begin by introducing *decision independence* as a key concept in making sound inferences and avoiding overfitting. I review how preregistration can create decision independence. I discuss some of the challenges in creating a preregistration for complicated research. Then I review six other techniques that also create decision independence and can themselves be preregistered or used when preregistration falls short. Finally, I discuss how to bring these techniques together to form an integrative strategy.

Decision Independence Guards Against Overfitting

When researchers report empirical results, the results are typically presented as if they are a function of the data:

$$Result = f(data) \tag{1}$$

For example a researcher might write, “A 2x2 condition-by-gender ANCOVA produced a significant main effect of condition, $F(1, 73) = 6.21, p = .03$.” The F value, 6.21, is a function of the data, calculated by applying standard mathematical formulas.

But as Gelman and Loken (2013) observed, describing a result as a function of data is incomplete. A result is also a function of a series of decisions made by the researcher during the data analysis workflow:

$$Result = f(data, decisions) \tag{2}$$

The points in an analysis workflow at which decisions are made are known as *researcher degrees of freedom* (Simmons, Nelson, & Simonsohn, 2011). The concept is very general: it includes all decisions about what data to collect and analyze, including inclusion and exclusion rules; how to analyze or model data, including preliminary parts of a multi-stage analysis such as scoring, transformations, and other kinds of preprocessing; and selecting which results to report and interpret. For example, even in a simple example like the one above, the researcher could alternatively have made a decision to collect more or fewer subjects; to use different covariates or none at all; to test the theoretical hypothesis with a simple effect in a subgroup rather than a main effect; to report an analysis of a different outcome variable instead; and much more. Different decisions would have led to different results.

There is nothing inherently problematic about the fact that results are a function of decisions, and in fact it is virtually impossible to imagine real-world empirical research that involves no decisions at all. Typically there are many defensible ways to investigate the same scientific question in data. For example, in a recent demonstration, 29 teams of experienced data analysts were provided with the same research question and the same dataset. No two teams analyzed the data the same way, resulting in 29 different answers, all based on different decisions (Silberzahn et al., 2018).

The problem then is not decisions per se. Rather, the risk of capitalizing on chance grows when the decisions are themselves a function of features that are specific to the data being analyzed:

$$Result = f_1(data, decisions) \quad (3)$$

$$decisions = f_2(data) \quad (4)$$

Thus, if an analyst excludes subjects because a scatterplot showed that they did not fit an expected pattern; adds covariates to a regression because an initial analysis did not produce a significant p-value; iteratively modifies a model to improve a fit statistic; reports a subset of results that make for a compelling story; or any one of multitudes of other possibilities that can occur in a wide range of research settings, the decisions that led to the reported result would be dependent on the data. The analysis or model specification will be overfit to the data at hand, making the results less replicable and potentially biased.

A researcher's decisions can become dependent on data through many mechanisms, including unconscious motivated reasoning, low-level cognitive or perceptual biases, or social influence (such as informal and formal peer review). All can be problematic. More subtly, it is not even necessary that the researcher actually attempted or contemplated more than one way of analyzing the data. Researcher degrees of freedom are defined by the space of analytic possibilities, not by the researcher's behaviors (Gelman & Loken, 2013). For example, if the researcher got an interesting effect with the first analysis and stopped, but in a parallel universe might have examined other specifications if the first result had been different, then there is still a researcher degree of freedom and its resolution is still a function of the data.

An analyst who has already seen or otherwise learned about some features of a dataset cannot "un-know" what is now known, and it is universally difficult for people to even recognize their own biases in operation, let alone stop them (Pronin, 2007). Unfortunately this means that good intentions, introspection, and expertise are no guard against overfitting. The biasing effect may be toward a specific conclusion, such as a researcher's preferred theory. But it may also be toward a counterintuitive result, a well-fitting model, a clean story, a pattern that provides the

researcher with a sense of discovery or of solving an important problem, or a variety of other ends. Any systematic influence on decisions can potentially be biasing.

Within this framework, we can define *decision independence* as a property of an observed result. A result has decision independence if the researcher's decisions that led to it – decisions about what data to collect and analyze, how to analyze it, and what results to report – are not functions of unrepeatable features of the specific dataset being analyzed. A heuristic for thinking about decision independence is to consider a counterfactual: If the data were different, could the decisions have been different? If the answer is “no,” the results have decision independence.

Decision independence is desirable when capitalizing on chance is a concern, but is not a sole criterion for all data analysis. Data-dependent decisions based on manipulation checks and positive controls, inspection of residuals, and other kinds of model checking can often lead to appropriate and justifiable changes in analyses or interpretations that are worth some added risk of capitalizing on chance. Results generated through interaction with the data – sometimes labeled “exploratory”¹ – can lead to serendipitous discoveries, inspire new thinking, and lead to new patterns to be modeled or hypotheses to be tested in future work. Moreover, decision independence is a safeguard but not a guarantee: a researcher could decide, even without knowing the data, to follow an analysis plan that capitalizes on chance or is otherwise not valid. However, decision independence is relevant to empirical claims or inferences that are intended to generalize beyond the specific dataset being analyzed, including inferences from a sample to a

¹ In my experience, the word “exploratory” can cause confusion. Some scientists use “exploratory” to mean speculative or question-driven research that is not conceived as testing a well-elaborated prior theory. In discussions of statistical inference, “exploratory” often refers to analyses that are done interactively with the data (i.e., without decision independence). These two senses – a scientist's mindset vs. a procedural mode of data analysis – do not always overlap, and research can be exploratory in one sense but not the other.

population, predictive modeling, or corroboration of a model or theory from a hypothesis-testing analysis. Thus it will frequently be a desirable feature of results.

Preregistration of Analysis Plans Can Create Decision Independence

Preregistration, broadly defined, means creating a verifiable record of knowledge or plans prior to some part of the research process. Preregistration can serve a number of goals. It creates transparency, by allowing other researchers to verify claims about when key decisions were made. It can be a planning tool, since the process of writing down plans can prompt more systematic thinking and facilitate conversations with collaborators and other peers.

When preregistration includes an analysis plan, it can be used as a procedural safeguard against capitalizing on chance (Nosek et al., 2018). A researcher can create an analysis plan before the data is collected, or for existing or secondary datasets, before looking at or otherwise learning about the data. Thus, the preregistration can be influenced by background knowledge but not by any features that are specific to the particular dataset to be analyzed.

An effective preregistered analysis plan has two important elements. First, the researcher must create a *decision inventory* that identifies all of the researcher degrees of freedom throughout the analysis workflow. This process can be aided by templates if relevant ones are available (e.g., van 't Veer & Giner-Sorolla, 2016; Open Science Framework, 2016), reviews of common researcher degrees of freedom (Wicherts et al., 2016), or peer review (Chambers, 2013).

Second, for each item in the decision inventory, the researcher must make a plan. Plans can be made on the basis of prior theory or best practice guidelines. For example, a theory of how attentive and inattentive respondents' reaction times differ can guide a plan for which trials to exclude; or a plan for how to score a questionnaire might be based on a validated algorithm.

When explicit guidance is not available, researchers can draw on experience to make an educated choice. A plan can sometimes be a contingency that follows a predetermined rule, such as a running a test of distributional assumptions to determine which of several different analyses to use.

If the decision inventory is complete, and the plans in it are clear and specific, then following a preregistered analysis plan will create decision independence. Researchers and readers can evaluate a preregistration by how well it does that, as well as whether it is sound on other statistical grounds (including whether contingency plans could unintentionally capitalize on chance). Contrary to some misconceptions, preregistration prevents neither justifiable departures from a pre-data plan nor reporting of additional analyses, which often have great value for discovery (Gernsbacher, 2018). Preregistration makes them transparently distinguishable from the planned analyses that have decision independence, helping authors and readers to calibrate their interpretations.

Practical Barriers to Preregistration

Preregistration has been used successfully by researchers in a variety of areas and fields (Nosek et al., 2018). However, for some complex research designs and paradigms, practical barriers remain (Tackett et al., 2017). We can identify two kinds of barriers associated with the two elements of preregistration.

First, it may be difficult to create a complete decision inventory in advance. For example, in longitudinal research it is not unusual for the science to change during the course of a study. Analyses or methods that are considered state-of-the-art at the start of a study may fall out of favor or be eclipsed by new developments. Hypotheses or questions that a study was designed to answer might become settled or reframed, while new questions that a study is capable of

answering might emerge. Or in any study, researchers may simply miss important decision points during the planning phase.

Second, even when a researcher can make a complete decision inventory, it may be difficult to make a specific plan for every item without seeing some features of the data. For example, in complex modeling studies, it may be necessary or desirable to select variables, select a model, or tune hyperparameters based on performance in real data, and no such dataset may exist before a study is to be run. And in some stages or kinds of data analysis, best practices may involve visual inspection or other forms of interaction with the data.

Faced with these practical barriers, some researchers might conclude that preregistration is not for them. Given that the underlying risk that preregistration safeguards against – capitalizing on chance – would not go away, would that mean that large swaths of research might have to settle for less sound inferences?

Fortunately, the answer is “no.” Even “marvelously complex” work, such as is common in areas like genetics and neuroscience, can benefit from preregistration (Bishop, 2017). If the larger goal is to create decision independence, then researchers can take an expanded view of preregistration as one part of a multi-strategy approach that still does that. First, preregistration can be more adaptive and flexible than many researchers may realize. Rather than writing a single *deterministic preregistration* with a simple rule for every decision, researchers can write *adaptive preregistrations* with plans to deploy flexible strategies (so long as the flexible strategies themselves preserve decision independence), sometimes iteratively in interim registrations as different parts of the data are observed. Second, when preregistration falls short – like when decisions must be made that were not anticipated in a preregistration – researchers can use other strategies that protect decision independence. Used well, a multi-strategy approach can

make complicated research designs – and the inferences drawn from them – every bit as rigorous and replicable as simpler designs with deterministic preregistrations.

Adaptive Preregistration And Beyond: Other Strategies for Sound Inference

In this section, I review six procedural and statistical safeguards to create decision independence, with references for further reading; and I discuss how they can either be incorporated into an adaptive preregistration or used for unanticipated decisions. These strategies are summarized in Table 1.

Standardization

Standardization means making decisions for how to handle situations that recur across multiple studies. For example, an individual researcher or a community of researchers might standardize the way neuroimaging data is preprocessed for a commonly used paradigm, the way that outliers in response time data are excluded for a particular task, or the way that scale scores from a particular questionnaire are to be calculated. Standardization has an if-then structure: if some paradigm, task, or questionnaire is used, then the standardized approach will be used. Thus standardization is like a general preregistration for a class of situations, rather than a specific study.

Standardization can be done for a single lab or workgroup by creating a standard operating procedure, or SOP (Lin & Green, 2016). An SOP is a document that describes default practices to be used when decisions arise that were not anticipated in a study-specific preregistration. Labs can develop SOPs for situations they commonly encounter in their work. For example, an infant lab might define exclusion rules for babies who “fuss out” during experiments, or a personality lab might standardize how they score commonly used self-report questionnaires. Frank (2018) reported developing a priority list for pruning random effects from

multilevel models that do not converge or produce errors such as non-positive definite matrices. This allows analysts in the lab to begin data analysis with “maximal” models that are considered best practice but are vulnerable to estimation problems (Barr, Levy, Scheepers, & Tily, 2013). The SOP allows the models to be respecified to yield proper solutions without an analyst who has seen the data making discretionary decisions.

At a fieldwide level, standardization can take the form of consensus “best practice” statements and common paradigms that delineate how to apply certain methods. To create decision independence, such standardization must be as specific as a preregistration would be. It must also go beyond a mere suggestion and become a socially enforced norm within a research community. For example, Weidman, Steckler, and Tracy (2017) contrasted how the Big Five structural model has evolved in the personality psychology community with how affective states are typically measured in the emotion literature. Widely used Big Five instruments have standard scoring systems that delineate what constructs are measured by the instrument, which items are indicators of which constructs, and how scales are to be scored. By contrast, Weidman and colleagues’ coding of articles in a top emotion journal showed that decisions about which items measure which constructs were often ad hoc and not repeated across studies. Weidman et al. (2017) argued that greater standardization of measurement would increase the replicability of published emotion research and facilitate efforts to synthesize results across studies.

Standardization only creates decision independence if it constrains researcher degrees of freedom. Standards that are vague, conflicting, or flexible, or ones that can be taken up or dropped at will, do not provide much constraint. Lin and Green (2016) recommended using public, timestamped SOPs linked from study-specific preregistrations. The preregistration takes priority when the documents conflict. Field-wide standardization is facilitated by published

validation research on common methods and measures, but it requires the members of research communities to reach consensus on a single standard. It also requires a healthy culture of peer review and critical discourse wherein researchers communicate norms and hold one another accountable for following them in the absence of a verifiable plan to do otherwise.

Blind Analysis

Blind analysis has been developed in physics to prevent investigators' biases from affecting decisions during data analysis (Klein & Roodman, 2005; Lyons, 2008). MacCoun and Perlmutter (2015; 2017) proposed extending it to psychology. Broadly speaking, blind analysis is “a method that hides some aspect of the data or result to prevent experimenter's bias” (Klein & Roodman, 2005, p. 147). It is not a single technique; rather, it is an approach to data analysis. The idea is to create a reduced or perturbed dataset in which an analyst can make decisions without having access to information that would bias the result.

Developing a good blind analysis strategy requires careful thought about which features of a dataset could potentially produce biases and which features must be preserved to accomplish the goals of the analysis. Blind analysis can be applied to data preparation steps like screening for outliers or selecting a transformation. For example, a dataset consisting of just dependent variables (with condition codes removed) can be visually inspected for outliers and distributional assumptions. Clinical trials sometimes use a form of blind analysis in triple-blind designs (Dawes, 2010), where the data analyst does not know which condition is which (e.g., conditions may be randomly labeled “A” and “B” instead of “Treatment” and “Control”). An extension of that described by MacCoun and Perlmutter (2015) is *cell scrambling* in factorial designs, which preserves F statistics but randomizes which effects each one maps onto. Dutilh et al. (2017) provide an example of how complex modeling decisions can be made while protecting from bias

by shuffling the link between the to-be-modeled data from a perceptual task and another variable it was hypothesized to relate to. Additional examples are in MacCoun and Perlmutter (2017).

Blind analysis creates partial decision independence – the goal is for decisions to be independent of the potentially biasing features of the data (such as those related to the focal hypotheses) but not from other features that the analyst needs to see in order to make decisions (such as the presence of outliers, the shapes of distributions, etc.). A good strategy for blind analysis will require both substantive domain knowledge and careful reflection to decide which features an analyst should and shouldn't be able to see. Blind analysis can also be combined with interim registrations: after decisions are made from a blind analysis, a researcher can register the decisions before working with the full dataset in order to create a transparent workflow.

Data Partitioning and Holdout Samples

The practice of partitioning data to avoid overfitting is common in engineering and computer science, though it apparently originated in psychology (Larson, 1931; credited by Arlot & Celisse, 2010). With the growing popularity of modeling and machine learning in behavioral science, more psychologists are rediscovering its benefits.

Data partitioning is useful when a researcher does not have enough prior knowledge to comfortably create a deterministic preregistration, perhaps because of novel features of a study design. For example, a researcher may wish to look at data before deciding how to create ad hoc scale scores, model a novel experimental task, select predictors or covariates for an analysis, specify a model or algorithm, etc. Prior to working with data, the analyst can partition it into two or more sets. One partition is designated as a holdout set (also called a test set) and set aside unanalyzed. The remainder of the data can be used interactively by the analyst to make decisions about scoring, preprocessing, variable selection, model specification, etc. Depending on the

complexity of the analysis, a single training dataset may be sufficient, or more partitions may be desirable (for example, in neural network modeling it may be valuable to do “learning” in one set and hyperparameter tuning in another). At the end, after all analysis decisions are made and the researcher has settled on a final model, the analysis is then run in the holdout set. Because of the risk of overfitting in training data, inferences are only made from the holdout data.

More complex cross-validation procedures, such as leave-one-out and k-fold, have also been developed (see Arlot & Celisse, 2010, for a review of both “simple” holdout samples and algorithmic cross-validation). Algorithmic cross-validation can be effective, but it requires knowing all analytic possibilities in advance so that all modeling steps, including preprocessing, can be accounted for in the cross-validation process. By contrast, holdout samples have the advantage that they place almost no restrictions on how much the analyst can explore and interact with data during training. Holdout samples are also relatively simpler to use and understand.

Using holdout samples in this way is a procedural analog of running an exploratory study and then using the results to guide a preregistered analysis plan for a replication. The main difference is that in classical replication the second dataset is collected later, whereas in data partitioning all data is collected up front but some is set aside unobserved. But in both cases, decision independence is established by basing decisions on interactive analyses in one dataset and drawing inferences from another.

Preregistration can work well in conjunction with data partitioning. Prior to first contact with data, a researcher can preregister a plan to create and use a holdout sample, including the size of the partition, the way it will be selected (randomly or otherwise), and how the final model will be evaluated in the holdout data (e.g., whether parameters will be freely estimated in the

holdout data or fixed based on training results, what statistics will be calculated, and what inference criteria will be used). This ensures that decisions about whether and how to use a holdout sample have been made without conducting any analyses that included the held-out data, which can create data leakage (Kaufman, Rosset, Perlich, & Stitelman, 2012) and undermine decision independence. After training, the researcher can write an interim registration describing the decisions that were made during training and specifying the analysis that will be run in the holdout data (Santos & Grossman, 2018). Such a combination can provide the best of all worlds: the flexibility of interactive exploration, the rigor of decision independence, and the transparency of preregistration.

The considerable benefits of data partitioning do not come free. The total sample must be large enough to support useful partitions. There is also no simple rule for deciding how large to make the partitions. It is often better to use a larger training set and a smaller holdout set. But specific details will depend on how much exploration will be done in the training set; how much power or precision is desirable in the holdout set; and particulars of both the phenomena being studied and the analyses or modeling approaches being used.

Supporting Studies

Data partitioning is very effective when researchers can collect a large enough dataset to support it (or for existing data, have enough data available). But sometimes it is not cost-effective or practical to obtain enough data to use partitioning effectively. Additionally, unanticipated decisions may arise after a researcher has started working with a full dataset, making it impossible to create a cleanly held-out set. In those cases, it is sometimes possible to run an inexpensive supporting study to guide and constrain decisions in an expensive main study.

In a supporting study, a researcher partially replicates some of the easier or cheaper elements of a main study design, possibly with extensions such as additional measures. The researcher can then treat the supporting study like a training set, conducting analyses in it to guide principled decisions that will be applied in the expensive main study.

An example of a supporting study is in Tackman, Srivastava, Pfeifer, and Dapretto (2017). In a longitudinal dataset originally collected for other purposes, we noticed some items in a questionnaire that resembled typical Big Five conscientiousness items, raising the possibility of using the existing data for a study of adolescent personality development. Rather than guess which items would form a reliable and valid composite – and risk overfitting if we needed to iteratively adjust the scale, as is often required in scale development – we ran a supporting study. We administered the questionnaire in a large enough sample to do sound psychometric analyses, selected items to form a conscientiousness scale based on analyses in that independent data, cross-validated within the supporting study, and then applied our final scoring decisions in the main data. This gave us confidence that our scale construction decisions were not unintentionally capitalizing on chance features of either the main data or the supporting study.

When used in this way, supporting studies are an inversion of the more usual, recommended approach of conducting validity research on new methods and measures prior to deploying them to test novel hypotheses (Flake, Pek, & Hehman, 2017; Strauss & Smith, 2009). Supporting studies are not a substitute for programmatic construct validation. Rather, they are a way to make up for the absence of such work when unexpected opportunities or challenges arise during research.

Supporting studies will only create decision independence if decisions about whether and how to conduct them are not biased by information from the main data. Because researchers will

often be inspired to do them after already working with the main data, they may not be as “pure” as other approaches to creating decision independence. It is helpful if researchers are transparent about what exposure to data they have already had when they decide to run a supporting study. Once the decision to run a supporting study has been made, researchers should also take care not to further explore parts of the main data that might bias their decisions in the supporting study. Researchers can create interim registrations of the decisions made in the supporting studies before returning to the main dataset for further transparency.

Coordinated Data Analysis

Researchers who work with existing datasets face specific challenges in doing replicable research, and they have generated a variety of practices and solutions for that kind of work (Weston, Ritchie, Rohrer, & Przybylski, 2018). One promising approach is coordinated analysis (Hofer & Piccinin, 2009). In a coordinated analysis, researchers who work with existing datasets that have similar features, such as measures of the same constructs, similar samples, etc. form collaborative networks to work on a common research problem. Analytic decisions are made collectively, with a goal of creating a harmonized protocol that can be carried out at all sites.

Coordinated analysis can create decision independence, and thus avoid capitalizing on chance, if analysis decisions are made without biasing knowledge in all datasets. For example, a researcher who works extensively with one longitudinal dataset might discover an interesting pattern in it but worry that it may be a spurious result. The researcher might identify other research groups that have longitudinal studies with the same or parallel measures and propose a coordinated data analysis. If analysis decisions are based on the first dataset, then results in the other datasets will have decision independence.

Coordinated analysis is facilitated by field-wide standardization, which increases the likelihood that different studies will contain parallel procedures and measures. Conversely, because existing datasets almost always differ in some ways, coordinated analysis can produce discrepant results across datasets for nonspurious reasons such as differences in measurement, sampling, longitudinal intervals, etc. The similarity of results across datasets will therefore often depend on the generalizability of the phenomenon being studied. This may complicate interpretations of results, though it may also be a desirable feature of a study.

Multiverse Analysis

The emphasis so far has been on finding the one best plan for each decision point. An alternative approach, however, is to exhaustively examine (or randomly sample from) the space of all analytic possibilities. This approach has been labeled multiverse analysis (Steege, Tuerlinckx, Gelman, & Vanpaemel, 2016), as well as other names like sensitivity analysis (Thabane et al., 2013).

In a multiverse analysis, the researcher identifies one or more decision points and lists all possible plans that could be made for each one. For example, a researcher may consider decisions about outlier exclusion rules and identify several defensible thresholds; different scoring schemes for a set of items; different combinations of covariates that could be included in an analysis; different residual covariance structures for a model; etc. The researcher then runs the analysis for all combinations of possible decisions and tabulates or plots the results.

Steege et al. (2016) describe multiverse analysis largely as an interpretive tool, allowing researchers to visualize how robust or fragile results are across different possible decisions. Going a step further, Simonsohn, Simmons, and Nelson (2015) proposed a more formalized approach called specification curve analysis, in which the distribution of results across different

analytic specifications is compared to a null distribution generated by permutations or bootstrapping. This approach allows calculation of a p -value for the multiverse.

When researchers have to make data-dependent decisions, for example because of complications arising that were not anticipated in a preregistration, then a multiverse analysis can be a useful after-the-fact diagnostic for evaluating the consequences of the loss of decision independence. However, multiverse analyses and specification curves can also be planned in advance and preregistered (Steenen et al., 2016). This approach can be effective when researchers do not wish to choose a single plan for one or more decision points; or when they do make a single plan but also want to register a plan to evaluate its robustness.

Conclusion: An Integrative Strategy

The techniques discussed here will continue to be refined as methodologists study them and researchers gain more experience applying them. This list is not exhaustive, and there is much room for more innovation in research workflows and statistical techniques. Researchers should not treat the approaches here as a simple menu. Both the techniques themselves, and the effectiveness of their implementation in each instance, should be evaluated by how well they serve their goal of avoiding capitalizing on chance. The concept of decision independence, and the counterfactual heuristic to evaluate it – if the data had been different, could the decisions have been different? – can provide researchers with a conceptual tool to do that.

For researchers who want to avoid capitalizing on chance in their complicated research, an integrative, multi-strategy approach will often start with preregistration. Attempting to create a decision inventory and a set of plans may seem daunting at first, but it can be a useful planning tool and researchers may be surprised by how much they are able to work out in advance. When researchers are not comfortable choosing a single plan for a decision point, they should consider

incorporating adaptive elements that allow for flexibility but preserve decision independence. They can also document important decisions along the way in interim registrations to create transparency. Researchers can handle unanticipated researcher degrees of freedom with practices that preserve decision independence when possible.

Finally, researchers – and readers, including editors and reviewers – should not let the perfect be the enemy of the good. Almost any researcher can at least create a partial preregistration, which will create transparency. A partial preregistration may safeguard against some kinds of overfitting; it can also facilitate interpretation by clarifying which decisions were not made in advance. Researchers should also make sensible departures from their preregistered plans when they are warranted. In such cases, transparency and careful interpretation are paramount. Because the goal of transparency is to facilitate both verification (Srivastava, 2018) and critique (Vazire, 2018), transparency may lead readers to view some conclusions with greater uncertainty, and they may expect authors to do the same. Decision independence is often both desirable and achievable, but when it is not, open and verifiable research with appropriately calibrated conclusions will still facilitate scientific progress.

References

- Arlot, S., & Celisse, A. (2010). A survey of cross-validation procedures for model selection. *Statistics Surveys*, 4, 40–79. <https://doi.org/10.1214/09-SS054>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>
- Bishop, D. P. (2017, October 1). Pre-registration or replication: the need for new standards in neurogenetic studies. Retrieved November 21, 2018, from <http://deevybee.blogspot.com/2017/10/why-we-need-to-stop-publishing.html>
- Chambers, C. D. (2013). Registered Reports: A new publishing initiative at Cortex. *Cortex*, 49(3), 609–610. <https://doi.org/10.1016/j.cortex.2012.12.016>
- Dawes, M. A. (2010). Triple-Blind Study. In *Encyclopedia of Research Design* (pp. 1542–1543). Thousand Oaks: SAGE Publications, Inc. <https://doi.org/10.4135/9781412961288>
- Dutilh, G., Vandekerckhove, J., Ly, A., Matzke, D., Pedroni, A., Frey, R., ... Wagenmakers, E.-J. (2017). A test of the diffusion model explanation for the worst performance rule using preregistration and blinding. *Attention, Perception, & Psychophysics*, 79(3), 713–725. <https://doi.org/10.3758/s13414-017-1304-y>
- Flake, J. K., Pek, J., & Hehman, E. (2017). Construct validation in social and personality research: Current practice and recommendations. *Social Psychological and Personality Science*, 8(4), 370–378. <https://doi.org/10.1177/1948550617693063>
- Frank, M. (2018, February 26). Babies Learning Language: Mixed effects models: Is it time to go Bayesian by default? Retrieved November 20, 2018, from

<http://babieslearninglanguage.blogspot.com/2018/02/mixed-effects-models-is-it-time-to-go.html>

Gelman, A., & Loken, E. (2013). The garden of forking paths: Why multiple comparisons can be a problem, even when there is no “fishing expedition” or “p-hacking” and the research hypothesis was posited ahead of time.

Gernsbacher, M. A. (2018). Writing Empirical Articles: Transparency, Reproducibility, Clarity, and Memorability. *Advances in Methods and Practices in Psychological Science*, 1(3), 403–414. <https://doi.org/10.1177/2515245918754485>

Hofer, S. M., & Piccinin, A. M. (2009). Integrative Data Analysis through Coordination of Measurement and Analysis Protocol across Independent Longitudinal Studies. *Psychological Methods*, 14(2), 150–164. <https://doi.org/10.1037/a0015566>

Kaufman, S., Rosset, S., Perlich, C., & Stitelman, O. (2012). Leakage in Data Mining: Formulation, Detection, and Avoidance. *ACM Trans. Knowl. Discov. Data*, 6(4), 15:1–15:21. <https://doi.org/10.1145/2382577.2382579>

Klein, J. R., & Roodman, A. (2005). Blind Analysis in Nuclear and Particle Physics. *Annual Review of Nuclear and Particle Science*, 55(1), 141–163. <https://doi.org/10.1146/annurev.nucl.55.090704.151521>

Larson, S. C. (1931). The shrinkage of the coefficient of multiple correlation. *Journal of Educational Psychology*, 22(1), 45–55. <https://doi.org/10.1037/h0072400>

Lin, W., & Green, D. P. (2016). Standard Operating Procedures: A Safety Net for Pre-Analysis Plans. *PS, Political Science & Politics; Washington*, 49(3), 495–500.

<http://dx.doi.org/10.1017/S1049096516000810>

- Lyons, L. (2008). Open statistical issues in particle physics. <https://doi.org/10.1214/08-AOAS163>
- MacCoun, R. J., & Perlmutter, S. (2017). Blind Analysis as a Correction for Confirmatory Bias in Physics and in Psychology. In *Psychological Science Under Scrutiny* (pp. 295–322). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9781119095910.ch15>
- MacCoun, R., & Perlmutter, S. (2015). Blind analysis: Hide results to seek the truth. *Nature News*, 526(7572), 187. <https://doi.org/10.1038/526187a>
- Nosek, B. A., Ebersole, C. R., DeHaven, A. C., & Mellor, D. T. (2018). The preregistration revolution. *Proceedings of the National Academy of Sciences*, 115(11), 2600–2606. <https://doi.org/10.1073/pnas.1708274114>
- Open Science Framework. (2016). Templates of OSF Registration Forms. <https://doi.org/None>
- Pronin, E. (2007). Perception and misperception of bias in human judgment. *Trends in Cognitive Sciences*, 11(1), 37–43. <https://doi.org/10.1016/j.tics.2006.11.001>
- Santos, H. C., & Grossman, I. (2018). Relationship of Wisdom-related Attitudes and Subjective Well-being over Twenty Years: Application of the Train-Preregister-Test (TPT) Cross-Validation Approach to Longitudinal Data. <https://doi.org/10.31234/osf.io/f4thj>
- Silberzahn, R., Uhlmann, E. L., Martin, D. P., Anselmi, P., Aust, F., Awtrey, E., ... Nosek, B. A. (2018). Many Analysts, One Data Set: Making Transparent How Variations in Analytic Choices Affect Results. *Advances in Methods and Practices in Psychological Science*, 1(3), 337–356. <https://doi.org/10.1177/2515245917747646>
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science*, 22(11), 1359–1366. <https://doi.org/10.1177/0956797611417632>

- Simonsohn, U., Simmons, J. P., & Nelson, L. D. (2015). *Specification Curve: Descriptive and Inferential Statistics on All Reasonable Specifications* (SSRN Scholarly Paper No. ID 2694998). Rochester, NY: Social Science Research Network. Retrieved from <https://papers.ssrn.com/abstract=2694998>
- Srivastava, S. (2018). Verifiability is a core principle of science. *Behavioral and Brain Sciences*, 41. <https://doi.org/10.1017/S0140525X18000869>
- Steege, S., Tuerlinckx, F., Gelman, A., & Vanpaemel, W. (2016). Increasing Transparency Through a Multiverse Analysis. *Perspectives on Psychological Science*, 11(5), 702–712. <https://doi.org/10.1177/1745691616658637>
- Strauss, M. E., & Smith, G. T. (2009). Construct Validity: Advances in Theory and Methodology. *Annual Review of Clinical Psychology*, 5(1), 1–25. <https://doi.org/10.1146/annurev.clinpsy.032408.153639>
- Tackett, J. L., Lilienfeld, S. O., Patrick, C. J., Johnson, S. L., Krueger, R. F., Miller, J. D., ... Shrout, P. E. (2017). It's Time to Broaden the Replicability Conversation: Thoughts for and From Clinical Psychological Science. *Perspectives on Psychological Science*, 12(5), 742–756. <https://doi.org/10.1177/1745691617690042>
- Tackman, A. M., Srivastava, S., Pfeifer, J. H., & Dapretto, M. (2017). Development of conscientiousness in childhood and adolescence: Typical trajectories and associations with academic, health, and relationship changes. *Journal of Research in Personality*, 67, 85–96. <https://doi.org/10.1016/j.jrp.2016.05.002>
- Thabane, L., Mbuagbaw, L., Zhang, S., Samaan, Z., Marcucci, M., Ye, C., ... Goldsmith, C. H. (2013). A tutorial on sensitivity analyses in clinical trials: the what, why, when and how.

- BMC Medical Research Methodology*, 13(1), 92. <https://doi.org/10.1186/1471-2288-13-92>
- van 't Veer, A. E., & Giner-Sorolla, R. (2016). Pre-registration in social psychology—A discussion and suggested template. *Journal of Experimental Social Psychology*, 67, 2–12. <https://doi.org/10.1016/j.jesp.2016.03.004>
- Vazire, S. (2018). Implications of the Credibility Revolution for Productivity, Creativity, and Progress. *Perspectives on Psychological Science*, 13(4), 411–417. <https://doi.org/10.1177/1745691617751884>
- Weidman, A. C., Steckler, C. M., & Tracy, J. L. (2017). The jingle and jangle of emotion assessment: Imprecise measurement, casual scale usage, and conceptual fuzziness in emotion research. *Emotion (Washington, D.C.)*, 17(2), 267–295. <https://doi.org/10.1037/emo0000226>
- Weston, S., Ritchie, S., Rohrer, J., & Przybylski, A. (2018). Recommendations for increasing the transparency of analysis of pre-existing datasets. <https://doi.org/10.31234/osf.io/zmt3q>
- Wicherts, J. M., Veldkamp, C. L. S., Augusteijn, H. E. M., Bakker, M., van Aert, R. C. M., & van Assen, M. A. L. M. (2016). Degrees of Freedom in Planning, Running, Analyzing, and Reporting Psychological Studies: A Checklist to Avoid p-Hacking. *Frontiers in Psychology*, 7. <https://doi.org/10.3389/fpsyg.2016.01832>

Table 1

Methods and Strategies That Can Create Decision Independence

Strategy	Description	Can be preregistered	Useful for unanticipated decisions after researchers have started analyzing data
Preregistration	Creating an inventory of all decisions and a plan for each one before collecting or observing data	Yes*	No
Standardization	Creating a default way that decisions will be made in commonly occurring situations	Yes	Yes
Blind analysis	Making decisions based on partial or perturbed data	Yes	Usually no
Data partitioning	Splitting data into a training set to base decisions on and a holdout set to draw inferences from	Yes	No
Supporting studies	Collecting inexpensive data to make decisions that can be applied in an expensive main study	Yes**	Yes
Coordinated analysis	Running analyses across multiple existing datasets	Yes	Yes
Multiverse analysis	Running many analyses with all possible combinations of decisions	Yes	Yes

* A plan to create a future preregistration can itself be preregistered. Interim registrations can be useful in an adaptive plan as a way to document decisions along the way, such as those made after conducting a blind analysis or exploring training set.

** The text emphasizes supporting studies' utility for unanticipated decisions. Researchers can plan them in advance as well, but this begins to blur the line into the more conventional and unremarkable practice of doing programmatic construct validation.