Null-hacking, a lurking problem

John Protzko

University of California, Santa Barbara

Abstract

Pre-registration of analysis plans involves making data-analysis decisions before the data is run in order to prevent flexibly re-running it until a specific result appears (*p*-hacking). Just because a model and result is pre-registered, however, does not make it reflective of underlying reality. The complement to *p*-hacking, null-hacking, is the use of the same questionable research practices to re-analyze open data to return a null finding. We provide a vocabulary for null-hacking and introduce the threat it poses. Null-hacking forces consideration of model fit to compare pre-registered and ‘alternative’ models. The reason null-hacking cannot be ignored is a null-hacked model can easily provide better fit to the data than a pre-registered one. Model fit, however, is a precarious problem, focusing just on model fit by only selecting a ‘best fitting model’ eliminates pre-registration, while giving default preference to pre-registered results ignores how well our models can represent the data. We provide a beginning solution aimed at retaining the advantage and justifications of pre-registration, while including model fit, and providing protection against null-hacking. We call this Fully-Informed Model Pre-registration and it involves strict supervised machine learning to maximize local model fit within heavily pre-specified decisions. This solution maximizes local model fit, eliminating the only justifications null-hacked results have. It is not yet a complete solution but merely the groundwork for why other approaches may be insufficient and what a future solution may look like.

Keywords: Open Science; Artificial Intelligence; Model Equivalence;

Here we illuminate a lurking problem called null-hacking, the process by which outside researchers, wishing to see a different result from the one published, use open data and unconstrained data analysis strategies to come up with ‘null’ results. Null-hacking is the necessary complement to *p*-hacking (1), but addressing the problem of null-hacking forces us to question the roles of open data, pre-registration, and confirmatory testing.

With the welcome expansion of open science practices including open data sharing, there will come an unprecedented amount of null-hacking. This will occur principally (though not exclusively) not in the peer-reviewed literature but instead in the dark literature, at the reviewer stage, and in independent re-analyses online (e.g. 1-3). This paper aims to bring the phenomenon to light, highlight the unique challenges it creates for open science, and provide a vocabulary for future identification of a solution.

Null-hacking is the process of engaging in questionable research practices (QRPs) with the intent to get a statistical result above some specified threshold (often $p > .1$, but can get above any *p*-value or Bayes Factor). The null-hacking researcher can selectively drop participants, re-parameterize data, add or drop covariates, alter the class of data analysis or estimator chosen, all in the attempts to return a ‘null’ result. As with *p*-hacking (1), the re-analysis may in fact be analyzing the data correctly. It is simply the case that the results that are returned are similarly untrustworthy.

The problem is compounded by the fact that there is a combinatorial explosion of models that can be run, each returning different results. In certain industries such as epidemiology, toxicology, biostatistics—null-hacking occurs at large scales (often through corporate funders) to prevent or delay public health reforms. Studies are re-analyzed to make the risks of smoking, asbestos, indoor tanning beds, leaded paint, benzidine, vinyl chloride monomer, climate

change—disappear (for data on this problem in these fields, see 15-16; and the citations therein; although not all data reuse is done to null-hack (17-18)). Furthermore, most re-analyses do not reach the published literature, but among those that do in clinical trials, over a third (35%) ‘find’ substantively different conclusions than the original authors (19).

There are numerous real-world examples to draw from of null-hacking, from R.A. Fisher’s denial that smoking cigarettes causes lung cancer (20) to debates about whether Asbestos should be considered a toxin (21). The freedom to re-analyze data has led to analyses claiming the climate change is not occurring. Null-hacking data from a 2008 study, researchers were able to argue that there was no change in global temperatures once the ‘correct’ adjustments to the data and statistical analyses were made (22-23). The ability to perform such analyses to ‘find’ null-results in less controversial fields, when the actual ‘truth’ is far from unclear, is equally possible. The question becomes: what is the ‘right’ way to analyze (any) data?

A simpler example from psychology would be the following: The Stroop effect is seemingly one of the most robust findings in psychology. Briefly, people take longer to read words about colors if they are printed in a different color (e.g. the word ‘Blue’ printed in the color green) than if they are printed in the same color (the word “Blue” written in blue). Re-analyzing data from a study using the Stroop and a number of other variables, we were able to show that there was no statistically significant difference between incongruent trials (word and print color mismatch) versus congruent trials (word and color match; see Table 1).

Table 1: Null-hacking the Stroop effect

Raw Incongruent – congruent time	Model adjusted incongruent – congruent time
M = 40.694, se = 1.565, $p < .001$	M = 8.364, se = 5.028, $p > .1$

From this data, should we conclude that the Stroop effect is not real? It seems unlikely, but why? The effect was able to be driven to non-significance simply by the inclusion of a number of covariates, using a procedure no different than that used in *p*-hacking (the covariates were another administration of the Stroop, scores on the trail-making test versions A and B, and sex; it took a total of running four models before finding one that made the Stroop effect disappear). Furthermore, the data was on 33 participants, making statistical power weak and null-hacking noticeably easier. If null-hacking can “eliminate” the Stroop effect, the deleterious effects of smoking, and climate change—then it can be a problem in any field where the ‘truth’ is not known.

One proposed solution to *p*-hacking is pre-registration of data analysis procedures before analyzing (ideally before collecting) the data. A pre-registration includes such factors as procedures, measures, rules for excluding data, plans for data analysis, predictions, state hypotheses and other analytic decisions before results are known (e.g. 14). Our focus here is on the aspects of pre-registration dealing with data and analysis decisions. Pre-registration can be enforced by journals, requiring it for publishing incredible results, changing the culture from pre-registering being rare to researchers having to answer why they did not pre-register. It is a wonderful and welcome addition to should become normal science. As we will see, the problem is much more difficult for null-hacking.

Exploring solutions to this problem of null-hacking brings to light some difficult hurdles. Furthermore, as null-hacking will occur primarily at the reviewer stage and by outside researchers online and not through peer-reviewed journals, requiring null-hackers to pre-register will not be a fruitful solution with no means to enforce it. The solution we present, a mixture of

model-fit optimization methods and pre-registration, we propose, provides the base for a future solution while also integrating pre-registration with model fit.

An Abstraction

Consider a researcher who pre-registers a given study and a single analysis plan (called model A). This researcher runs the study with absolute fidelity, analyzes the data using only their pre-registered analysis plan, and gets the as-predicted result. Now consider a second researcher who uses the open data and runs models B, C, D, E, and F. Most of models B, C, D, E, and F do not show evidence for the effect found in model A. The question is: do we believe in the pre-registered effect found in model A, the null results of models B-F, or something else?

Argument in favor of preferring the null results.

Pre-registering a result that returns as-predicted does not mean it is a true phenomenon in the world; it simply removes one reason for why the results may be false (i.e. the results were not *manufactured* through *p*-hacking). Pre-registering an incorrect data decision can lead to spurious, pre-registered findings that may also replicate (24). One such example would be dichotomizing a continuous variable; such a procedure is never statistically justified (see 25) and can lead to replicable spurious results even if pre-registered. Another example could be creating a causal path from X to Y while the ‘true’ causal path is that X and Y share a common cause. Such a procedure would show an as-predicted causal path even though no such path exists. Just because it is pre-registered does not make it true.

The null-hacking researcher has more models showing a given effect does not exist (see also 26). This can be couched in the term ‘robustness checks’, as in the results from the pre-registered study are not robust to alternate specifications of the model.*

Just as the *p*-hacking researcher can find some justification in the literature for their procedure, likewise the null-hacking researcher can find justification for theirs, so one of their models can be couched as using the ‘correct’ parameterization. Therefore, given more models show no effect and the model showing the effect is not ‘robust’ to alternate models, one can argue we should not believe in the pre-registered results of model A.

Argument in favor of preferring the pre-registered result.

Pre-registration is a protection against using QRPs to create findings. The results found in model A could not have occurred from using QRPs, and therefore the source of uncertainty has been eliminated. In addition, the results of models B, C, D, E, and F have no such protection that casts doubt on how accurately they may be representing the phenomenon.

Furthermore, simple vote counting among models is a specious method of model comparison. Models can be re-fit, parameterized, and adjusted in an exponential number of ways. In addition, the average effect size in psychological science is small (e.g. 27-29). There will always be more models that can show no effect than models can show a small effect. Furthermore, not all models are equal. That different models are better able to represent the underlying data-generating procedure better than others is the core of model fit and its role in statistical modelling and hypothesis comparison (e.g. 30-31).

* Robustness checks make a number of assumptions, the core of which are that all models are equal, and that the path under consideration represents the same structural relation under different specifications. For problems with the second assumption, see (26).

Therefore, as the results from models B, C, D, E, and F were null-hacked, are thus unreliable, and a simple vote-counting method of models is not how hypothesis comparison should be carried out, we should reject the null-hacked models and accept the more reliable pre-registered results of model A.

The tension.

Given the premise above there will always be more models that can show no effect than models can show a small effect, we cannot ignore model fit as it is the central component of deciding between competing models. This point cannot be emphasized enough. Given two competing models, one way to demarcate between them is how well the model fits the underlying data.

Not all models fit the same data equally well. To be in contention for equal consideration against a pre-registered model, a null-hacked model or series of models must show it fits the data at least equally well if not better. A null-hacked model showing inferior fit to the data should be rejected over a pre-registered analysis free from QRPs. It is extremely unlikely, however, that a pre-registered model will be the best fitting one. Given the number of ways models can be adapted, re-parameterized, and given even a few variables in a dataset, millions if not billions of models can be fit to data (see ‘the combinatorial problem’ 32 for this problem without QRPs). The simple likelihood of choosing the best-fitting one *a priori*, through pre-registration, is minute. Therefore, it should be recognized accepting the results from a pre-registered model will likely result in accepting the results from an inferior model compared to the space of all possibilities.

This is one of the strongest ways null-hacking creates a tension for pre-registration, it forces us to acknowledge model fit (for model comparison purposes) which will almost certainly work against the pre-registered results. The null-hacker can simply show the superior fit to the data from any one of thousands to millions to billions of alternate models and have good quantitative grounds for arguing for the veracity of their model over the pre-registered one.

A seemingly attractive solution would be to simply pre-register ‘I will use the results from the best fitting model’. Therefore, the results from such an investigation will be using the best fitting model and pre-registered. This approach, however, does not solve the conundrum of whether to believe the results of model A or models B, C, D, E, or F.

What seems to be an inevitability of all modelling procedures is no single model will be the best fitting but instead a number of models with sometimes contradictory results will all fit the model equally well (33). This phenomenon, called the Rashomon effect (see also ‘model equivalence’), has a number of possible solutions, each of which have implications for the justifications of pre-registering.

One class of solutions is to use *a priori* assumptions to limit the space of possible models that can be accepted. Such assumptions can limit what variables can or cannot be included, disallowing higher-order interactions among covariates, linear opposed to non-linear models, appealing to simpler models. This solution is most in line with the basic philosophy of pre-registering, as these assumptions are built into (sometimes unintentionally or unknowingly) the pre-registered analysis plan. The advantage of this approach is it can bypass overfitting of models to data.

Take a simple model with only one predictor (as seen in Figure 1):

Figure 1: Three Regression lines fit to the same data with different amounts of variance explained by each.

Three regression lines are fit to this simple model, a linear, a quadratic, and a 5th order polynomial. A linear growth model (solid line) is most people's default assumption, and in this example accounts for 37% of the variance in scores. A quadratic model (dotted line) fits the data better than the linear model (accounting for 40% of the variance) and makes nearly the same predicted values. The 5th-order polynomial, however, provides substantially better fit to the data (accounting for 55% of the variance) than either linear or quadratic models. Given any analysis, a non-linear model will always fit the data better than a linear model. The appeal to linearity is a statistical and philosophical *assumption*. We may feel confident in rejecting the 5th-order polynomial model because of overfitting and (we believe) it is thus unlikely to reflect the true relationship. This example, however, should make the disadvantage of using *a priori* assumptions (e.g. linearity) clear, it is unlikely to return the overall best fitting model.

The second class of solutions is to allow the competing models to co-exist, testing them in future data (e.g. 34). This pluralism provides a solution through a non-solution. It argues our degrees of belief should be spread across and updated among all of the different potential models. In a Bayesian context, this would correspond to all models in the ‘best fitting’ class being updated equally, effectively preserving the rank-order of the priors.

The advantage of this approach is it side steps the problem of plurality of best-fitting models by making it a feature. The disadvantage is it questions the entire point of pre-registering. If the models are equally well fitting between pre-registered model A and any of the null-hacked models B, C, D, E, and F, the solution is to believe the most in what you believed in before the results were shown (preserving rank order of your priors). Here there is little point to pre-registering as it neither alters prior rank orders nor model fit. If the models are not all equally well fitting, the best will likely not be the pre-registered model. As discussed previously, the pre-registered model is almost assuredly not the overall best fitting model. Thus, there would be no point in pre-registering a given data analysis. Instead, the only thing should be pre-registered is: “we will fit a series of models and accept the results of all of the models in the ‘best-fitting’ class”. This obviously creates a problem when results differ across different models (e.g. does X correlate with Y if one model shows it does and another, equally well-fitting, shows it does not?). The solution to whether believe the pre-registered model A or any of the null-hacked models is to believe them all according to how well the data fits, so null-hacking will likely be a fruitful endeavor under this approach.

Furthermore, not all research involves the administration of short studies online inexpensively to thousands of participants collected in a very short amount of time. Some of science’s biggest studies simply cannot be replicated (large natural experiments as in famine

effects in the Netherlands at the hands of the Nazis, for example). Thus, the pluralistic-and-wait approach also cannot provide a solution in unreplicable studies as no future study could resolve the competing accounts.

The third class of solutions to the Rashomon effect/model equivalence is to use unsupervised machine learning optimization methods (sometimes-called non-directed data mining; 35) to search for the single best model. This class of solutions often involves constructing a model or series of models either on independent historical data or on a test-set of the primary data (e.g. random 10-20%). These models are then verified on the new data or the rest of the old data. Numerous techniques exist and are being expanded such as neural nets, random forests, support vector networks, and a host of yet-to-be-invented methods. A full treatment of this class of models is beyond the point of this paper, but they provide a promising way of optimizing models to data (see 36; for an example of this). Many times these procedures still produce a plurality of models but work is ongoing to bypass this problem with new techniques (e.g. 37-38). The disadvantage they provide for the argument about null-hacking is there would no longer be a point to pre-registering a data analysis plan. Instead, the pre-registrations can simply read: ‘we will use procedure x to find the best fitting model’. This then turns the entirety of scientific inquiry into an exercise in *unsupervised* machine learning free from hypothesis or confirmatory testing.

A Solution to Null-hacking?

Null-hacking creates a threat to the purpose of pre-registration and open science. Given open data, the null-hacking researcher will be able to fit a plurality of models to a given dataset in the attempt to dispute a finding. The null-hacking researcher will likely be able to find more models and parameterizations showing an effect does not exist than the pre-registering researcher

will be able to show it does, as effect sizes are small. Therefore, model fit must be taken into account to counter the simple vote counting of competing models. The pre-registered model, however, is unlikely to ever be the best fitting one. We are thus forced to accept that a null-hacked model with superior fit should not be accepted while a pre-registered model with inferior fit should. We cannot open the door to argument by model fit when it suits us but then try to close it to competitors when it goes against our results.

Take, for example, the eliminating the Stroop effect example from earlier. The base model with only one predictor (trail-making test A) still shows evidence for the Stroop effect (mean time difference 42.94 seconds, $se = 4.944$, $p < .001$). It also accounts for .74% of the variance and shows acceptable fit to the data (e.g. non-significant χ^2). The null-hacked model, on the other hand, accounts for 73.4% of the data and shows superior fit (likelihood-ratio test of model comparison $\chi^2(3) = 43.46$, $p < .001$). Therefore, an argument can be made *solely* on model fit that the null-hacked model is indeed the 'right' one. Even if the original model was pre-registered, the large differences in model 'quality' can be used to leverage the argument in favor of null-hacking. If pre-registering would give us reason to believe the returned results were true, the problem would be much less severe. Pre-registration, however, can only eliminate one reason why the results may be spurious and acceptance of the results cannot be justified on only these grounds.

As we posit, the majority (but not entirety) of null-hacking will occur in the reviewer stage and by outside researchers online, with no way to enforce such pre-registration at a policy level. Requiring null-hacking researchers to pre-register their re-analysis plans before having access to open data is also nearly impossible to enforce. Pre-registration can be enforced outside the dark literature through preventing publication in peer-reviewed journals unless pre-

registration is used. When the re-analysis is done on a private blog or other outlet with no editorial oversight, the only way would be for the owners of the open data to act as gatekeepers. With the incentive to not have one's research challenged, however, allowing those who own the data to act as gatekeepers is undermined by a strong conflict of interest (see 39). Furthermore, such procedures cannot be enforced when the original researchers cannot be contacted. Going one-step further, simply closing data and not allowing it to be accessed by other researchers will absolutely solve the null-hacking problem but is antithetical to the open science movement.

Pre-registration and model fit

Based on the arguments above, we argue pre-registration and model fit lie in a space where the two can easily be in opposition to one another. The pre-registered model is unlikely to ever be the overall best fitting model (due to the exponentially large number of ways a given effect can be parameterized).

One solution to the null-hacking problem is to focus entirely on model fit. This solution will ensure either a small series of models will be fit to the data. Thus, the conflict between the pre-registering researcher and the null-hacking researcher is resolved if none of their models are to be accepted. Instead, the overall best fitting model shall be the one to be interpreted. The problem with this solution is it turns all of scientific inquiry into machine learning optimization algorithms with no longer any role for pre-registering analysis plans outside of 'we will use XXX procedure to return the best fitting model'. Furthermore, such algorithms have been criticized that while they can be unparalleled for simple prediction, the overfitting they can return can give incorrect (non-veridical, or not matching the truth of the world) causal results (e.g. 33*Error! Bookmark not defined.*). Focusing solely on model fit also eliminates

confirmatory testing, which should also be preserved. Therefore, we cannot focus entirely on maximizing model fit.

A second solution to the null hacking problem is to give unquestioned priority to pre-registered studies. When faced with a pre-registered model versus an alternative one, the pre-registered model takes precedence regardless of fit. This allows us to retain pre-registration of analysis plans and the confidence they bring. It seems unlikely, however, this approach will be in the best interest to science. Ignoring fit entirely in preference of an *a priori* specified model puts the cart before the horse. The majority of science is done in the pursuit of understanding the world and its underlying truths (e.g. 40). To ignore data, the way the world looks, and instead give priority to our preferred representation of the world (the model) is hard to justify from a realist perspective.

Take a simple example: the model in Figure 2.

Figure 2: Pre-registered linear model imposed on a 'true' curvilinear one.

The 'true' relationship between these two variables is likely curvilinear. Without knowing this ahead of time, one would most likely impose a simple linear model on the data (because many scientists, through simplicity or training, default to linear over non-linear models). The problem is: a pre-registered linear regression would still 'return' a positive, statistically-significant linear slope. Thus, the model was pre-registered, the results were 'confirmed', but the basic functional form is completely wrong (see 41). We would be grossly misrepresenting the world. This further shows the danger of a blank acceptance of a pre-registered model without regarding the possible underlying 'truths'.

The solution to the problem of null-hacking that is likely to be of the greatest benefit lies somewhere in the subspace between entirely focusing on model fit and giving an unquestioning

priority to pre-registration. The solution we introduce here can be called Fully-Informed Model Pre-registration and represents the intersection of the pre-registration and model fit optimization approaches. The approach involves taking the space of all possible models that can be fit to the data and reducing the possible models to a narrow ‘sub-space’ of possible models based on pre-registered decisions. Then, focusing on model fit optimization through the use of more advanced estimators and machine learning, we can return the locally best fitting model given tight constraints.

This solution is different from pre-registering an analysis plan, in that it does not specify the exact model to be tested—the gold standard of pre-registration of analysis plans. Instead, it allows for flexible analysis plans to be implemented *automatically*, without influence from the researcher.

This beginning of a solution is similar in form to other solutions in that the space of all possible models is reduced to a small subspace (e.g. 42). What null-hacking forces us to address, however, is model fit. Not all models fit data equally well, and solutions that do not take this fact into account are susceptible to being dismissed as weakening the investigative process by the inclusion of ill-fitting models. These assumptions can often be unexamined, such as when model robustness checks are equally averaged over specifications (e.g. 43). Such simple averaging assumes equal weight among all models, in short, making the assumptions that all models are equal (see also 44). Furthermore, such approaches not only treat all models as equal but all models as equally valid. Another solution could also be to include many variations on data decisions (not model decisions); such forms of analyses can allow others to see how seemingly arbitrary decisions could alter a pattern of results (e.g. 45). Similarly, however, such models treat all data decisions as equally valid through the unweighted presentation and interpretation of

results. Many models can be said to be invalid based on seemingly innocuous data decisions that should not be used for inference. For some examples: listwise deletion of noncompliers in randomized trials breaks randomization and ability to infer causation (see.46); dichotomizing continuous variables (24, 47); conditioning on mediators or colliders (48-49); mixing exogeneous and endogeneous variables with the OLS estimator (50) . Not all models are equally valid.

Using one administration of the Stroop as a covariate for a second administration of the Stroop to show the Stroop effect ‘doesn’t exist’ (Table 1 above) is a clear example of not all models being equally valid. Thus, the blind inclusion of ‘all possible models’ for weighting makes a number of invalid assumptions. What null-hacking forces us to do instead is recognize that all models are not equal. The goal of any solution would be to maximize local model fit within the model subspace that is pre-specified.

In this introductory solution, a number of assumptions underlying the data are pre-registered before a chosen model-fit optimization procedure is taken. These assumptions involve decisions about the distribution of the data (normal, Poisson, binomial), functional form of relationships (linear, non-linear etc.), the forced inclusion of any variables (e.g. not accepting a model that does not return with certain variables included), the allowing or disallowing of higher-order interactions etc. Then, with these assumptions built in, a pre-registered given procedure (e.g. gradient boosting tree, random forests, Tabu search, Bayesian additive regression trees, LASSO optimization, greedy search selection, etc.) is used to return the best fitting model while conforming to the constraints. This is similar to supervised machine learning (e.g. 35) and adds in pre-registration of the strict decisions to limit the sub-space of models the algorithms will

search.² This allows the researcher to find the best fitting model given the constraints, eliminate model-based p -hacking, and preserve a role for pre-registration.

The approach here does not need to only concern complicated models. Even models as simple as the humble t -test have a number of ways to be null-hacked. The most obvious is through testing model assumptions. Testing whether, for example, errors are heteroskedastic or normally distributed or groups have equivalent variances is a core part of statistical education. Often the assumptions are tested and then the ‘appropriate’ model is chosen afterwards. This two-stage procedure in fact has deleterious effects on the conclusions that can be drawn from the resulting models (see 51-60, for examples). These assumption-testing tests have their own assumptions whose tests have their own assumptions—it is statistical assumptions all the way down. Testing whether to use a separate-variance t -test over a pooled-variance t -test, for example, creates biased inference in the resulting estimation. Though these problems are too large to discuss here, the common thread seems to be statistical assumption tests are often underpowered, based on their own assumptions (which can often be violated), and resulting inference more often than not relies on accepting the null hypothesis from these statistical tests. Though the majority of the work on the deleterious effect on inference after assumption checking has been done in relation to t -tests and ANOVAs, similar problems arise when using multi-step model tests in OLS regression (see 61, for example). In short, the conclusion from this literature is to avoid testing assumptions or using tests for covariate inclusion/exclusion, instead defaulting

² It is important to point out that all of our estimators are at their heart, model fit maximization. The simple OLS regression is a model fit maximization procedure, which does so through the minimization of mean squared errors. More complicated procedures we suggest here simply allow for more flexibility and maximize model fit through other criteria such as gradient descent or mean prediction error. Even the t -test uses comparisons of means (not medians or some other aggregator) via the assumption of the mean being the most informative estimator from a normal distribution.

to using the robust estimator from the beginning with all covariates included. This approach, bolstered by the growing literature on the problems of assumption testing, can help protect against one class of null-hacking, but not alternately formed models which may provide superior fit to the data.

Another position that has been advocated is a ‘report both ways’ approach, where the researcher reports the results both with and without their data decision (like dropped outliers). This way, the researcher can report the results for both the model with data decision A and data decision B. This is helpful if results from data decision A and B come to the same conclusion, but what if A and B come to different conclusions? The way to demarcate between them is, again, a function of model fit, which puts us right back where we started with having to take into account fit. Furthermore, there is no reason why we should only interpret the results from data decisions A and B while numerous other data decisions exist. What justification is there to be limited to the possible data decisions based only on the creativity of the original investigator? Therefore, ‘report both ways’ turns into ‘report a large number of ways’ and we are left wondering which to believe.

Overall, when models start to be compared against each other, model fit must be taken into account. The easiest way to avoid this problem within one’s own research is to provide one and only one confirmatory model for testing. This should be done using robust estimators and have only collected data and covariates will be included in the final model. By using this confirmatory testing (instead of testing assumptions or a ‘report k ways’ approach), the researcher can avoid model fit comparisons within their own data.

This approach does not help, however, when a reviewer or someone else wishes to use the open data and provide an alternate model. As mentioned before, the confirmatory, pre-

registered model is almost guaranteed to not be the best fitting model. If, however, the Fully-Informed Model Pre-registration approach is used, then we can be sure any alternate model fit to the data will be at best of equal fit and more likely of inferior fit. Thus, we can have a better defense of the pre-registered model than “but I pre-registered it”, which is often the only defense that can be mustered. We can acknowledge aspects like model fit, and preserve pre-registration.

An example.

A simple example would be the following: Suppose you were a researcher trying to test which applicants to a job were given a callback from a recruiter. At your disposal is a large number of variables that can be parameterized in different ways about the applicant: education, ethnicity, sex, years of experience, currently employed [yes/no], time since last employment if unemployed, number of spelling mistakes in the cover letter, whether referred by a current employee, length of resume in pages, percentage of job ad keywords appearing in CV or application, and potentially many others.

The study could be a simple prediction, or it could be an experiment, where resumes are submitted with sex of the applicant randomly assigned. How do we analyze such a dataset? In the methods before pre-registration, many different models and parameterizations of the variables could be run on the data to show that male applicants are more likely to get a call back than a female applicant (or vice versa, depending on what the researcher ‘wants’ to show).

With pre-registration, how each of the variables is coded, which ones are included in the confirmatory model, and a finalized analysis plan is (ideally) locked in and then that model is run. As we discussed above, however, that model is nearly guaranteed to not be even close to among the best-fitting models. Non-linearities, moderations, and other effects are unlikely to be

included which may better capture the phenomenon of ‘being called back’. One model is chosen, and that is it.

With Fully-Informed Model Pre-registration, however, the situation is different. First, the parameterization of each of the variables is pre-registered ahead of time, including whether they would be allowed to interact, linearities, and which covariates *might* be allowed to be in the procedure. Then, the data is applied to an estimator that can use either holdout samples or cross-validation to maximize model fit within the space of possible models. There are a quite a few such procedures and their number continues to grow. Suppose one chooses LASSO, an estimator that optimizes prediction through evaluation of added information of each individual predictor.³ Including redundant predictors would hurt the efficiency of our prediction equation, so an idealized subset of the possible variables could be predicted automatically, regardless of the performance of our indicator of inquiry (sex). To prevent overfitting, a penalty function is applied and pre-registered so the resulting selection procedure balances decreased error without overfitting. When the data is analyzed, the resulting model will return not all of the models and interactions that were present in the pre-registration analysis plan (as in traditional pre-registration), but instead an idealized subset containing enough information to overcome the penalty function and provide an optimal fit model. This model will be better fitting than a traditionally pre-registered model. In fact, by virtue of the estimation procedure, a better model is unlikely to be able to be found. If further protections against overfitting (such as cross-validation) are also implemented, then the results are one from numerous independent confirmations in the same dataset. Thus, we retain pre-registration, and through the local

³ LASSO is but one of the many types of variable selection machine learning algorithms mentioned in this paper. It is only used here for illustration purposes as it is frequently posed as a beginner’s level introduction to Machine Learning.

optimization of model fit, any further model ‘null-hacked’ will be less well fitting and can be rejected on those empirical grounds.

Implications for the Future

We have not provided a complete solution, just the blueprints and vocabulary to begin discussing this problem. Acknowledging that null-hacking exists and will continue to grow in psychology as it has in other fields (e.g. toxicology, epidemiology) means we need to find justification for believing one model over another. The solution presented here, Fully-Informed Model Pre-registration, integrating model fit maximization into heavily pre-registered data decisions, provides a partial solution. It comes with some costs and implications, however.

When collecting new data, each new decision to collect another variable should be rigorously examined. Often we have found researchers like to collect multiple, unreported variables ‘just to see what happens’. Such collection of additional dependent variables or covariates (e.g. collecting extensive demographics by default) provides exponentially more models that can be run. Without a specific reason for including a specific variable in the model, it should not be collected. Unfortunately, not all research is able to accommodate this. Research using large existing datasets (e.g. UK biobank, General Social Survey, World Values Survey, etc.) must make do with what was collected. There are problems with a complete focus on model fit as theories with odd predictions having stricter tests, focusing on falsifiability, and avoiding affirming the consequent can become problems (e.g. 62). These decisions, however, are more often in the process of generating the study and tests rather than the specifics of the model after the data is collected. Therefore, there should be very strict decisions made before data is collected to account for all of the potential problems described in this paper.

In addition, researchers should become more comfortable with model-fit maximization procedures. While we mentioned a couple of procedures often under the umbrella term ‘machine learning’, we do not endorse any specific approach (see also Bayesian Model Averaging, 63). What the conflict between null-hacking and pre-registration forces us to do is to acknowledge model fit. Procedures that treat all models as equal or all models as equally valid will not provide a solution to these problems. Many researchers are not familiar with this concept beyond an ‘amount of variance accounted for’ understanding. Hopefully, future research can help simplify methods for maximizing local model fit within a subspace of possible models based on pre-registered decisions. Our solution, however, also forces us to acknowledge we will unlikely be fitting the globally best fitting model. This is acceptable, as scientific hypotheses must do more than simply ‘fit’ data well. But the tension between pre-registration and a complete focus on model fit cannot be resolved by choosing only one path. Both must be better integrated into future research.

Conclusion

Null-hacking, using QRPs to ‘find’ results that conflict with the original reported results, will happen. Given the average effect size in many of the sciences is often small; it will be easier to produce a model showing an effect does not exist than one that it does. Therefore, simple model counting cannot resolve which results to believe. Furthermore, as in *p*-hacking, null-hacked results are not by definition wrong, they are simply untrustworthy. When faced with a pre-registered model showing a result and a competitor model from the same data showing no result, what do we believe? As not all models fit data equally well, model fit cannot be ignored. A pre-registered model is unlikely to ever be the best fitting model, even the locally best-fitting one. Fully optimizing model fit, however, leaves no room for the purpose of pre-registration. The

proposed solution here involves restricting the subspace of all possible models based on strict a priori criteria. Then, through supervised machine learning, the best fitting model(s) within those constraints can be found. This allows for accounting for model fit (which null-hacking forces us to acknowledge) while at the same time allowing a role for pre-registration. We must simply accept this blend of pre-registration and model fit will not fit the globally best fitting model. Many already tacitly make this assumption, for example, in prioritizing linear models over higher-order polynomial models (as in Figure 1).

The heart of the problem enumerated here is best summarized by the following: (32, p. 53-54):

“Unless all but a handful of alternative hypotheses have been ruled out, the fact that a model passes a statistical test provides almost no reason to give it credence, for there may well be billions of alternative models of the same data that were they to be subjected to statistical testing, would do as well as or better than the particular model under consideration. We cannot rationally pretend that the alternatives don't exist or that they are inconceivable, or unconceived. If we have no knowledge as to whether or not better theories may be hiding among the unexamined multitudes, how can it be reasonable to give credence to the theory we happened to think of and to examine statistically?”

Pre-registration and open science should be preserved, promoted, and used. It may well be the new era of science, protecting it is of the utmost importance. We hope the vocabulary introduced here and themes strung together can be put forward to better the future of open science.

Acknowledgments: We would like to thank Hunter Gelbach; Jon Krosnick; Lee Jussim; Jonathan Schooler; Dan Simons for their helpful comments on this paper.

References

1. 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>
2. Blatt, M. R. (2015). Vigilante science. *Plant Physiology*, 169, 907-909.
<https://doi.org/10.1104/pp.15.01443>
3. Fiske, S. T. (2016). Mob Rule or Wisdom of Crowds. *APS Observer*
4. Simonsohn, U. (2014). Posterior-hacking: Selective reporting invalidates Bayesian results also. <https://doi.org/10.2139/ssrn.2374040>
5. Cook, R. D., & Weisberg, S. (1982). *Residuals and influence in regression*. New York: Chapman and Hall.
6. Anderson, C. A., Lepper, M. R., & Ross, L. (1980). Perseverance of social theories: The role of explanation in the persistence of discredited information. *Journal of Personality and Social Psychology*, 39(6), 1037-1049.<https://doi.org/10.1037/h0077720>
7. Gehlbach, H., & Robinson, C. D. (2018). Mitigating illusory results through preregistration in education. *Journal of Research on Educational Effectiveness*, 11(2), 296-315.
doi:10.1080/19345747.2017.1387950
8. Zhao, N. (2009, March 23). The Minimum Sample Size in Factor Analysis. Retrieved from <https://www.encorewiki.org/display/~nzhao/The+Minimum+Sample+Size+in+Factor+Analysis>
9. Kline, P. (1979). *Psychometrics and psychology*. London: Academic Press.

10. Comrey, A. L., & Lee, H. B. (1992). *A first Course in Factor Analysis*. Hillsdale, NJ: Erlbaum.
11. Gelman, A., & Loken, E. (2014). The Statistical Crisis in Science Data-dependent analysis-a "garden of forking paths"-explains why many statistically significant comparisons don't hold up. *American Scientist*, 102(6), 460. <https://doi.org/10.1511/2014.111.460>
12. 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, 337-356. <https://doi.org/10.1177/2515245917747646>
13. John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological science*, 23(5), 524-532. <https://doi.org/10.1177/0956797611430953>
14. Lindsay, D. S., Simons, D. J., & Lilienfeld, S. O. (2016). Research preregistration 101. *APS Observer*, 29(10).
15. Michaels, D., & Monforton, C. (2005). Manufacturing uncertainty: contested science and the protection of the public's health and environment. *American journal of public health*, 95(S1), S39-S48. <https://doi.org/10.2105/AJPH.2004.043059>
16. Levy, K. E., & Johns, D. M. (2016). When open data is a Trojan Horse: The weaponization of transparency in science and governance. *Big Data & Society*, 3(1), 2053951715621568. <https://doi.org/10.1177/2053951715621568>
17. Curty, R. G., Crowston, K., Specht, A., Grant, B. W., & Dalton, E. D. (2017). Attitudes and norms affecting scientists' data reuse. *PloS one*, 12(12), e0189288. <https://doi.org/10.1371/journal.pone.0189288>

18. Pearce, N., & Smith, A. H. (2011). Data sharing: not as simple as it seems. *Environmental Health*, 10(1), 107. <https://doi.org/10.1186/1476-069X-10-107>
19. Ebrahim, S., Sohani, Z. N., Montoya, L., Agarwal, A., Thorlund, K., Mills, E. J., & Ioannidis, J. P. (2014). Reanalyses of randomized clinical trial data. *Jama*, 312(10), 1024-1032. <https://doi.org/10.1001/jama.2014.9646>
20. Fisher, S. R. A. (1959). *Smoking: the cancer controversy: some attempts to assess the evidence*. Edinburgh: Oliver and Boyd.
21. Hart, L., (2001). Summary Minutes of the National Toxicology Program Board of Scientific Counselors Report on Carcinogens Subcommittee Meeting December 13 - 15, 2000. <https://ntp.niehs.nih.gov/ntp/roc/twelfth/draftbackgrounddocs/minutes20001213.pdf>
22. Wallace III, J. P., D'Aleo, J. P. & Idso, C. D. On the Validity of NOAA, NASA and Hadley CRU Global Average Surface Temperature Data & The Validity of EPA's CO2 Endangerment Finding Abridged Research Report. Accessed June 21 at <https://archive.is/69vfk#selection-85.0-173.9>
23. Kasprak, A. (2017, July 14). Peer-Reviewed Study Proves All Recent Global Warming Fabricated by Climatologists? Retrieved from <https://www.snopes.com/fact-check/climatology-fraud-global-warming/>
24. Mayo, D. G. (1996). *Error and the growth of experimental knowledge*. University of Chicago Press. <https://doi.org/10.7208/chicago/9780226511993.001.0001>
25. DeCoster, J., Iselin, A. M. R., & Gallucci, M. (2009). A conceptual and empirical examination of justifications for dichotomization. *Psychological methods*, 14(4), 349. <https://doi.org/10.1037/a0016956>

26. Patel, C. J., Burford, B., & Ioannidis, J. P. (2015). Assessment of vibration of effects due to model specification can demonstrate the instability of observational associations. *Journal of clinical epidemiology*, 68(9), 1046-1058. <https://doi.org/10.1016/j.jclinepi.2015.05.029>
27. Richard, F.D., Bond, C.F., Jr., & Stokes-Zoota, J.J. (2003). One hundred years of social psychology quantitatively described. *Review of General Psychology*, 7, 331-363. <https://doi.org/10.1037/1089-2680.7.4.331>
28. Ozer, D. J. (2007). *Evaluating effect size in personality research*. Guilford Press.
29. Szucs, D., & Ioannidis, J. P. (2017). Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLoS biology*, 15(3), e2000797. <https://doi.org/10.1371/journal.pbio.2000797>
30. Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association*, 71(356), 791-799. <https://doi.org/10.1080/01621459.1976.10480949>
31. Gelman, A. (2011). Induction and deduction in Bayesian data analysis. *Rationality, Markets and Morals*, 2(67-78), 1999.
32. Glymour, C., Scheines, R., & Spirtes, P. (1987). *Discovering causal structure: Artificial intelligence, philosophy of science, and statistical modeling*. Academic Press. <https://doi.org/10.1016/B978-0-12-286961-7.50010-X>
33. Breiman, L. (2001). Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science*, 16(3), 199-231. <https://doi.org/10.1214/ss/1009213726>
34. Chang, H. (2012). *Is water H₂O?: Evidence, realism and pluralism* (Vol. 293). Springer Science & Business Media.
35. Marcoulides, G. A. (2005). *Discovering Knowledge in Data: an Introduction to Data Mining*. <https://doi.org/10.1198/jasa.2005.s61>

36. Taylor, M., Kalbach, C., & Rose, D. (2019). Teleology and personal identity. Unpublished manuscript.
37. Tulabandhula, T., & Rudin, C. (2014). Robust optimization using machine learning for uncertainty sets. arXiv preprint arXiv:1407.1097.
38. Bertsimas, D., Gupta, V., & Paschalidis, I. C. (2015). Data-driven estimation in equilibrium using inverse optimization. *Mathematical Programming*, 153(2), 595-633.
<https://doi.org/10.1007/s10107-014-0819-4>
39. Doshi, P., Goodman, S. N., & Ioannidis, J. P. (2013). Raw data from clinical trials: within reach?. *Trends in pharmacological sciences*, 34(12), 645-647.
<https://doi.org/10.1016/j.tips.2013.10.006>
40. Niiniluoto, I. (2017). Optimistic realism about scientific progress. *Synthese*, 194(9), 3291-3309. <https://doi.org/10.1007/s11229-015-0974-z>
41. Tukey, J. W. (1977). *Exploratory data analysis* (Vol. 2).
42. Simonsohn, U., Simmons, J. P., & Nelson, L. D. (2015). "Specification Curve: Descriptive and Inferential Statistics on all reasonable Specification"
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2694998
<https://doi.org/10.2139/ssrn.2694998>
43. Young, C., & Holsteen, K. (2017). Model Uncertainty and Robustness: A Computational Framework for Multimodel Analysis. *Sociological Methods & Research*, 46(1), 3-40.
<https://doi.org/10.1177/0049124115610347>
44. Slez, A. (2017). The Difference Between Instability and Uncertainty: Comment on Young and Holsteen (2017). *Sociological Methods & Research*,
<https://doi.org/10.1177/0049124117729704>

45. Steegen, 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>
46. Angrist, J. D., Imbens, G. W., & Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434), 444-455.
<https://doi.org/10.1080/01621459.1996.10476902>
47. MacCallum, R. C., Zhang, S., Preacher, K. J., & Rucker, D. D. (2002). On the practice of dichotomization of quantitative variables. *Psychological Methods*, 7, 19-40.
<https://doi.org/10.1037/1082-989X.7.1.19>
48. Schisterman, E. F., Cole, S. R., & Platt, R. W. (2009). Overadjustment bias and unnecessary adjustment in epidemiologic studies. *Epidemiology (Cambridge, Mass.)*, 20(4), 488.
<https://doi.org/10.1097/EDE.0b013e3181a819a1>
49. Rohrer, J. M. (2018). Thinking clearly about correlations and causation: Graphical causal models for observational data. *Advances in Methods and Practices in Psychological Science*, 1(1), 27-42. <https://doi.org/10.1177/2515245917745629>
50. Freedman, D. A. (2008). On regression adjustments to experimental data. *Advances in Applied Mathematics*, 40(2), 180-193. <https://doi.org/10.1016/j.aam.2006.12.003>
51. Easterling, R. G., & Anderson, H. E. (1978). The effect of preliminary normality goodness of fit tests on subsequent inference. *Journal of Statistical Computation and Simulation*, 8(1), 1-11. <https://doi.org/10.1080/00949657808810243>
52. Zimmerman, D. W. (2004). A note on preliminary tests of equality of variances. *British Journal of Mathematical and Statistical Psychology*, 57(1), 173-181.
<https://doi.org/10.1348/000711004849222>

53. Shuster, J. J. (2005). Diagnostics for assumptions in moderate to large simple clinical trials: do they really help?. *Statistics in medicine*, 24(16), 2431-2438.
<https://doi.org/10.1002/sim.2175>
54. Rasch, D., Kubinger, K. D., & Moder, K. (2011). The two-sample t test: pre-testing its assumptions does not pay off. *Statistical papers*, 52(1), 219-231.
<https://doi.org/10.1007/s00362-009-0224-x>
55. García-Pérez, M. A. (2012). Statistical conclusion validity: Some common threats and simple remedies. *Frontiers in psychology*, 3, 325. <https://doi.org/10.3389/fpsyg.2012.00325>
56. García-Pérez, M. A., Núñez-Antón, V., & Alcalá-Quintana, R. (2015). Analysis of residuals in contingency tables: another nail in the coffin of conditional approaches to significance testing. *Behavior research methods*, 47(1), 147-161. <https://doi.org/10.3758/s13428-014-0472-0>
57. Rochon, J., Gondan, M., & Kieser, M. (2012). To test or not to test: Preliminary assessment of normality when comparing two independent samples. *BMC medical research methodology*, 12(1), 81. <https://doi.org/10.1186/1471-2288-12-81>
58. Lantz, B., Andersson, R., & Manfredsson, P. (2016). Preliminary tests of normality when comparing three independent samples. *Journal of Modern Applied Statistical Methods*, 15(2), 11. <https://doi.org/10.22237/jmasm/1478002140>
59. Parra-Frutos, I. (2016). Preliminary tests when comparing means. *Computational Statistics*, 31(4), 1607-1631. <https://doi.org/10.1007/s00180-016-0656-4>
60. Delacre, M., Lakens, D., & Leys, C. (2017). Why Psychologists Should by Default Use Welch's t-test Instead of Student's t-test. *International Review of Social Psychology*, 30(1).
<https://doi.org/10.5334/irsp.82>

61. Leeb, H., & Pötscher, B. M. (2005). Model selection and inference: Facts and fiction. *Econometric Theory*, 21(1), 21-59. <https://doi.org/10.1017/S0266466605050036>
62. Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological review*, 107(2), 358-367. <https://doi.org/10.1037/0033-295X.107.2.358>
63. Montgomery, J. M., & Nyhan, B. (2010). Bayesian model averaging: Theoretical developments and practical applications. *Political Analysis*, 18(2), 245-270. <https://doi.org/10.1093/pan/mpq001>