The Epistemic and Pragmatic Function of Dichotomous Claims Based on Statistical Hypothesis Tests

Duygu Uygun Tunç, Mehmet Necip Tunç, Daniël Lakens

Duygu Uygun Tunç is Postdoctoral Researcher at Eindhoven University of Technology, Department of Industrial Engineering and Innovation Sciences, Eindhoven NL 5600 MB (d.uygun.tunc@tue.nl); Mehmet Necip Tunç is PhD Candidate, Tilburg University Social Psychology Department, Tilburg, NL 5000 LE (m.n.tunc@uvt.nl); and Daniël Lakens is Associate Professor, Eindhoven University of Technology, Department of Industrial Engineering and Innovation Sciences, Eindhoven NL 5600 MB (d.lakens@tue.nl). This work was partially funded by VIDI Grant 452-17-013 from the Netherlands Organization for Scientific Research, and by the European Union under the Horizon 2020 Marie Skłodowska-Curie grant 801509 TÜBİTAK Co-Funded Brain Circulation Scheme 2. The authors thank Deborah Mayo, Sander Greenland, Oliver Maclaren, and Thom Baguley for their valuable feedback on an earlier version of this manuscript. All authors have equally contributed to this manuscript.

Abstract

Researchers commonly make dichotomous claims based on continuous test statistics. Many have branded the practice as a misuse of statistics and criticize scientists for the widespread application of hypothesis tests to tentatively reject a hypothesis (or not) depending on whether a p -value is below or above an alpha level. Although dichotomous claims are rarely explicitly defended, we argue they play an important epistemological and pragmatic role in science. The epistemological function of dichotomous claims consists in transforming data into quasi-basic statements, which are tentatively accepted singular facts that can corroborate or falsify theoretical claims. This transformation requires a pre-specified methodological decision procedure such as Neyman-Pearson hypothesis tests. From the perspective of methodological falsificationism these decision procedures are necessary, as probabilistic statements (e.g., continuous test statistics) cannot function as falsifiers of substantive hypotheses. The pragmatic function of dichotomous claims is to facilitate scrutiny and criticism among peers by generating contestable claims, a process referred to by Popper as “conjectures and refutations.” We speculate about how the surprisingly widespread use of a 5% alpha level might have facilitated this pragmatic function. Abandoning dichotomous claims, for example because researchers commonly misuse p -values, would sacrifice their crucial epistemic and pragmatic functions.

Keywords: falsification; inference; philosophy of science; statistical decision; testing

1. INTRODUCTION

The scientific literature is filled with dichotomous empirical claims. Statements such as “X influenced Y” and “X was associated with Y”, or alternatively “X did not influence Y” and “X was not associated with Y” based on the comparison of a test statistic (i.e., a quantity derived from a sample that is used in a hypothesis test, such as a *t*-value) against a criterion value are so common in the empirical scientific literature that it is tempting to unquestioningly accept their role in knowledge generation. The practice of making dichotomous claims is easily taught in introduction to research methods classes or learned by imitating peers. And yet, if scientists were asked to defend why they make dichotomous claims based on probabilistic hypotheses about data in their articles, we expect many of them might not find it easy to provide strong arguments in defense of their own actions.

Some statisticians have strongly argued against the practice of making dichotomous claims altogether. Instead, they have advocated for the practice of interpreting test statistics as continuous measures of discrepancy with a model, or to compute continuous measures of evidence (Gibson, 2020; Greenland & Poole, 2013; Rozeboom, 1960; Wasserstein, Schirm & Lazar, 2019; for discussions, see Hager, 2013; Hubbard & Bayarri, 2011). Listening to their criticisms of the status quo, a researcher might wonder if it is even possible to defend the use of dichotomous claims in science. Perhaps this practice has established itself purely through norms that no one is able to justify? Given the scarcity of principled defenses of the role dichotomous claims play in the scientific literature, should all researchers just admit they suffer from “dichotomania, the compulsion to perceive quantities as dichotomous even when dichotomization is unnecessary and misleading, as in inferences based on whether a P value is ‘statistically significant’” (Greenland, 2017)?

In the present paper, we argue that dichotomous claims have a veritable role to play in science. They are neither inherently problematic, nor are they easily replaceable by alternative approaches, such as estimation or presenting a full probability distribution. We explicate that the

function dichotomous claims serve in science is not primarily a statistical one, but an epistemological and pragmatic one. From a philosophical perspective rooted in methodological falsificationism (Lakatos, 1978; Popper, 1959/2002), dichotomous claims are the outcome of methodological decision procedures that allow scientists to arrive at empirical statements which connect data to phenomena. These empirical statements can have truth values and function as falsifiers or corroborators of theoretical claims about phenomena. In this regard, dichotomous scientific claims based on methodological decision procedures facilitate a process of scientific criticism that contributes to the growth of knowledge, which Popper referred to as “conjectures and refutations.” We believe that a better understanding of the role of dichotomous claims in science will help researchers to judge whether making dichotomous claims is aligned with their philosophy of science, and rationally decide if they want to use tools such as statistical hypothesis tests that are grounded in such a philosophy.

2. THE LOGICAL STRUCTURE OF SCIENTIFIC INFERENCES IN A FALSIFICATIONIST FRAMEWORK

One commonly used approach to scientific inquiry is the hypothetico-deductive method (Hempel, 1966), where a hypothesis is proposed, predictions are deduced and tested, which are subsequently either corroborated or falsified. The hypothetico-deductive method relies on deductive logical relations between *theoretical* and *basic* statements. Simply put, theoretical statements express general conjectures¹ and basic statements express particular observations that can contradict or support these conjectures. Basic statements are necessarily dichotomous in that they affirm or deny the existence of singular facts. These dichotomous empirical claims play a crucial role when testing

¹ These range from claims about patterns or regularities connecting multiple events (i.e., “phenomena” or what are frequently called “effects”) to larger patterns or regularities connecting clusters of events. While the term “theory” is often used to refer to statements about higher-order patterns or regularities, statements about phenomena are also theoretical as these cannot be directly observed. For our purposes in this paper there is no need to introduce a terminological difference between levels of theory. For simplicity, by theoretical statements we often refer to claims about phenomena throughout the paper.

theoretical claims, because it is through establishing singular facts that we assess the empirical merit of theoretical claims.

In his *Logic of Scientific Discovery* Popper proposes a particular hypothetico-deductive approach to theory testing that hinges on the falsifiability of theories through observation, namely methodological falsificationism. The core idea behind falsificationism is that if a specific empirical prediction can be derived from a theory and the prediction turns out to be false, then we can conclude that the theory is also false (although no number of true predictions can assure us that theory is true). Theories, according to Popper, are *universal statements*, which are not restricted to any individual spatio-temporal region.² An example is the second law of thermodynamics. Every universal statement can be expressed in the form of a non-existence statement. For example, the second law of thermodynamics can be expressed as ‘There is no heat flow from a colder object to a hotter object without additional energy supply in that direction’. Universal statements can thereby be falsified by *basic statements*, which are *singular existence statements* that express the occurrence of an event in a certain time and place.³ According to Popper, good scientific theories do not assert that something can be observed, they deny that some things can be observed. If we genuinely managed to create a system that violates the second law of thermodynamics, this fact, expressed in a basic statement like ‘Here is a system that exhibits reverse heat flow with no energy supplied to the

² Logically, falsifiable theoretical statements are strictly *universal* (Popper, 2002). Such statements are either ‘all-statements’ or ‘if-then statements’, such as ‘All swans are white’ or ‘If something is a swan it is white’.

³ The proper falsifier of a theoretical statement should itself be an empirically testable proposition. Thus, we need a restricted or “singular” existence statement of the form ‘Here is an entity with property p’, or ‘Event X is occurring in space-time region k’, such as ‘There is at least one black swan in this lake’ (see Popper, 1962, pp. 195-196). A basic statement is an observational statement, which can be confirmed or falsified by empirical observation by virtue of being spatiotemporally restricted. At the same time a basic statement contradicts a (strictly universal) theoretical statement by virtue of implying a strict existence statement that is prohibited by it (e.g., ‘There is at least one black swan in this lake’ implies ‘There is at least one black swan’). Thus, basic statements are themselves falsifiable and thereby function as the proper falsifiers of universal theoretical statements.

system from the outside ' would falsify the non-existence statement driven from the second law of thermodynamics, hence the second law itself.

So far, we have simply laid out the logical properties of scientific statements that feature in falsificationist scientific inferences. What is methodological about methodological falsificationism? Science progresses by falsifying theoretical claims expressed in universal statements through observations expressed in basic statements. However, observations do not automatically lead to basic statements. To illustrate why a mere observation cannot falsify a universal statement by itself, let us assume we are testing the universal statement 'All swans are white'. In our search for a non-white swan, we encounter a swan whose feathers are registered by our measurement device as black. How can we be sure that its feathers are congenitally black, and not painted by pranksters or fanatics who oppose the idea that all swans are white? Can we trust our measurement device sufficiently to conclude the swan is really black? If not, how do we decide which equipment provides us with the most accurate assessment of the color of feathers? How do we rule out the possibility that environmental circumstances in the location of observation are atypical and affect the measurement of color? Alternatively, how can we be sure that the animal we are observing is a swan and not a morphologically similar, hitherto undiscovered species? All these questions indicate that when we claim to have observed a black swan, we do so based on "auxiliary" assumptions we make as part of a methodological procedure (see Lakatos, 1978, pp. 23-28; Meehl, 1978; Popper, 2002, sec. 19-20; 1962, pp. 237-239). We can justify our use of a particular methodological procedure to decide upon the color of swan feathers to the extent that our auxiliary assumptions are severely tested (Mayo, 2018, p. 84; Popper, 2002, p. 62). Our basic statement about the existence of a black swan is thus not purely based on empirical observation, but also on a methodological procedure. So, basic statements express singular facts which can be tentatively accepted or rejected via certain *methodological decision procedures* (see Popper, 2002, pp. 65-67, 82-85). The tentative nature of basic statements is important to stress. As Hacking (1965) writes: "Rejection is not refutation. Plenty of rejections must be only tentative." Similarly, Neyman (1960, p. 13) regards

inferential behavior as an “act of will to behave in the future (perhaps until new experiments are performed) in a particular manner, conforming with the outcome of the experiment.”

Universal statements that make up scientific theories are on the level of phenomena, while empirical observations occur on the level of data. One cannot directly falsify the former just using the latter. It is only by transforming the observed data into an appropriate factual claim through a methodological decision procedure (i.e., into a basic statement) that one can potentially falsify a universal statement. Phenomena are relatively stable features of the universe that are not dependent on particular observations, while the specific data that are observed depend on details of the experimental procedure and the measurement device (Woodward, 1989; 2000). To use an example from Frick (1996), that “frustration increases the tendency to aggression” is a statement about a phenomenon, whereas an experiment can at best show that frustration increased (or did not increase) aggression for the subjects in the experiment. Although phenomena are detected by data, the data never directly or necessarily entail claims about phenomena (Bogen & Woodward, 1988), because there is always the possibility of error, for example due to measurement or sampling (Woodward, 1989). That is, the relationship between theories and empirical observations is not strictly logical (Woodward, 2000). To make this relationship logical, we must bridge the gap between universal statements and observed data by transforming the data into basic statements, which consist in *testable claims about the data*. This transformation allows universal statements to be falsified by empirical observations using deductive logical operations.

3. THE EPISTEMIC FUNCTION OF DICHOTOMOUS CLAIMS

So far, our discussion has focused on methodological decisions that link basic statements to observations containing no or negligible error, meaning that no error theory is (or must be) applied to the measurement or analysis of data thus obtained (i.e., “in corrigible data,” Suppes, 1974).⁴

⁴ By negligible error we mean, with Suppes, errors that are not supposed to happen under normal or usual circumstances, such as a lab assistant recording the data wrong, or a solar eclipse suddenly

However, in science such observations are increasingly rare, and we usually have observations that are known to contain random error which is expressed in probabilities (i.e., “corrigible data”, Suppes, 1974).⁵ Because the error associated with these observations is probabilistic, such error-containing observations can only be described in probabilistic terms. This poses an additional challenge for methodological falsificationism, as probabilistic statements (e.g., ‘This is probably a black swan’) are not proper falsifiers for universal statements. In a modus tollens inference (i.e., If P, then Q. Not Q. Therefore, not P) we cannot use a probabilistic proposition (i.e., Probably not Q) as a premise to refute a universal proposition (i.e., If P, then Q).

Regardless of whether data contains negligible error or non-negligible error, philosophically we are faced with the same problem. Researchers need to determine under which conditions an observation would count as a basic statement (i.e., a tentatively accepted singular fact). The solution developed for the former problem also applies to the latter. That is, irrespective of whether data

disturbing a study involving photosynthesis. Suppes (1974) explains further: “When we claim the data are incorrigible, we mean that the data are accepted without challenge, and there is no systematic attempt to estimate error in the measurements or in the recording of the data [under ordinary conditions]...Consequently we would be unlikely to attribute discrepancies between data and theory to errors of observation” (p. 271; 277).

⁵ Examples of non-negligible error can be spurious correlations or deviations in the values or quantities of parameters (see also Mayo, 1996, p. 449). It is worthwhile to note that corrigibility is a property that has essentially to do with the nature of observations, not of theories or hypotheses. As Suppes (1974) emphasizes, the corrigibility/incorrigibility of data does not necessarily follow from the nature of theories (e.g., whether they are probabilistic, deterministic etc.) but from the methodology of testing them. For instance, classical mechanics is a deterministic theory that is applied to corrigible data, thus it has to be combined with a complex analysis of errors in its application to real world observations (unlike quantum mechanics, which is a probabilistic theory that already has a theory of error internal to it). To illustrate further with our previous examples, a hypothesis such as 'All swans are white' is typically applied to incorrigible data, while one such as 'People name the ink color of a color word more quickly when the semantic meaning of the word is congruent with the color it is written in, compared to when the semantic meaning is incongruent with the ink color' is applied to corrigible data – although both are universal statements that express deterministic predictions.

contain negligible or non-negligible error, we need a methodological decision to turn data into a basic statement which can function as a potential falsifier of a theoretical statement. In the case of data with non-negligible error, we just need additional methodological decision procedures, which allow us to tentatively accept the claim that the data expressed in probabilistic terms indicates a singular fact (Popper, 2002, p. 260; Lakatos, 1978, p. 25). As in other methodological decisions (such as accepting the validity of a measurement tool, or assuming that all the relevant covariates are controlled for in a model), the crux is that the decision procedure should be open to scrutiny and be specified in advance.⁶

Most of the data that scientists collect contain non-negligible random error, and these data can be transformed into potential falsifiers only through statistical inference. For data containing non-negligible random error, we need a reliable methodological decision procedure such as statistical hypothesis testing that effectively distinguishes true signals from noise (at least in the long run). It is always possible that some observations are not in line with our predictions *only* due to random variation. These observations represent error, and as such, should not be regarded as proper falsifiers of theoretical predictions. Thus, we need a methodological decision procedure that allows us to distinguish random error from a signal in the data with some known reliability, whose

⁶ Some critics have argued that one of the statistical assumptions about the data model related to the mechanism of data generation, namely random selection from an underlying population, rarely occurs, especially in the social sciences, and therefore some statistics (such as *p*-values) based on this assumption are always wrong (see Trafimow, 2019). This criticism can be damning if test statistics are calculated with the intention of generalizing to a well-defined population rather than testing hypotheses. However, when the intention is to test a hypothesis, the goal is to find a disconfirming incidence rather than to generalize the findings to a well-defined population (see Hager, 2000 for the difference between statistical auxiliary hypotheses without and with empirical content). Thus, as is often the case in the social and behavioral sciences, when the data generated by non-random (e.g. convenience) sampling, the statistics will refer to a hypothetical population (i.e. a population of respondents that would result if the same sampling strategy were followed infinitely many times). Then, the calculated statistics can be interpreted as disconfirming incidences which serve to falsify the test hypothesis. Thus, random sampling from a well-defined population is not a necessary assumption if statistics such as the *p*-value describe a specific empirical situation related to the basic statements associated with the test hypothesis.

performance is itself open to empirical evaluation. Statistical hypothesis testing is such a methodological decision procedure. For example, in error statistical approaches which use error frequencies or error probabilities (see Mayo & Spanos, 2011; Mayo, 2018) the goal is to design a methodological procedure that, with a low enough probability of leading to incorrect conclusions, allows us to infer that we have made an observation that is consistent or inconsistent with our theory. What error probability is deemed “low enough” is determined through pre-specified decision criteria, which are integral to the methodological procedure. By making dichotomous claims based on a pre-specified statistical rejection rule, one arrives at what we propose to call *quasi-basic* statements. Quasi-basic statements differ from basic statements we discussed previously due to their inferential nature, but they fulfill a very similar epistemic function. Quasi-basic statements express (with a prespecified maximal error probability) statistical inferences that can corroborate or falsify theoretical claims. Thus, the core epistemic function of dichotomous statistical decisions based on test statistics is to arrive at a quasi-basic statement that can corroborate or falsify a theoretical statement in accordance with prespecified methodological criteria.

For example, a researcher might be interested in the relationship between cognitive processing speed and the congruence of information. Suppose that the researcher designs a test to see if on average people name the ink color of a color word more quickly when the semantic meaning of the word is congruent with the color it is written in, compared to when the semantic meaning is incongruent with the ink color (i.e., the *theoretical statement*).⁷ The data show that the time required to name the ink color varies considerably across participants and trials, but the response times are on average much slower in the incongruent condition than in the congruent condition (i.e., the *empirical observation*). Suppose that the researcher specified in advance that they would consider their prediction corroborated if on average response times in the congruent condition were faster than in the incongruent condition, and specifically when the *p*-value from a

⁷ Note that if the theoretical statement was about “all people,” the associated basic statement should not have been put into test with a sharp null hypothesis.

paired *t*-test was lower than an alpha level of 0.05 (e.g., the *statistical hypothesis*). After performing the statistical test, the researcher finds a *p*-value smaller than the alpha level, and states that the participants in their sample required more time to name the ink color for color words with an incongruent semantic meaning (i.e., a *quasi-basic statement*). They claim their prediction is corroborated and conclude that there is at least one sample that indicates a population with the property that the response times are longer when the color word is incongruent with ink color than when it is congruent. Of course, the extent to which this empirical claim is warranted does not depend solely on the data, but also on how strongly supported all auxiliary hypotheses in the experiment were (e.g., participants were not colorblind, they understood and followed the instructions, etc.). Moreover, to what extent this empirical claim corroborates a universal theoretical claim, such as the universal relationship between reaction time and stimuli congruence/incongruence, depends also on the validity of other theoretical assumptions (e.g., background theories about human cognitive capacities, possible boundary conditions regarding sensory modalities, number of stimuli etc.). We elaborate in the next section on these additional sources of error and uncertainty in scientific inference.

3.1 Claims about data vs. claims about phenomena: From basic to theoretical statements

Statistical error control is only one step in the management of uncertainties inherent to theory testing (Mayo, 2018; Spanos & Mayo, 2015). When we set out to empirically evaluate theoretical claims, we must make methodological decisions concerning uncertainties at every step of the derivational chain that links the conceptual core of a theory to concrete singular observations. Thus, the epistemic function we assign to dichotomous statistical decisions is best understood as one step in a broader set of methodological decisions scientists must make if they want to test a theory. Drawing mainly on work by Lakatos (1978), Mayo (1996), Meehl (1978; 1990) and Suppes (1962), Figure 1 presents a general outline of a model-theoretical representation of the interrelations between descriptive statistical statements (at the level of data), statistical inference (at the level of

statistical hypotheses), empirical predictions (at the level of substantive hypotheses) and theoretical claims about phenomena.

Figure 1: Derivational chains linking different levels of scientific models

The theoretical model is used to specify a state of the world in which the predictions of the theory would hold. Besides the axioms or the theoretical core of the theory, theoretical models typically require researchers to decide to accept other theories (e.g., a theory about human memory might require accepting a theory about perception or attention) to derive theoretical predictions. Theoretical models might also require certain factual assumptions, for example when specifying boundary conditions of the theory. The experimental model comprises criteria for the design and analysis of experiments. At this level, researchers need to make additional auxiliary assumptions concerning theories of measurement, the absence of confounding factors, and *ceteris paribus* conditions. Finally, the data model comes with additional, statistical auxiliary assumptions that researchers need to make, such as assumptions about how the data might be realized (e.g., normally distributed data). The epistemic function of quasi-basic statements is to establish an inferential connection between data and substantive hypotheses through the tests of statistical hypotheses.

In order for dichotomous empirical claims based on statistical tests (i.e., quasi-basic statements) to have implications for theoretical claims, such as corroboration or falsification, all auxiliary assumptions that are relied on in the theoretical, experimental and data models must be sufficiently valid. That is to say, if we want to put forward any scientific claim about phenomena (rather than data), we must specify our theoretical and experimental models, and aim for sufficiently high test severity and error control at each of the levels in Figure 1. As Mayo (2019, p. 80) explains, we want the methods we use to have the capability to uncover incorrect claims, and we systematically apply these methods at different levels to severely test auxiliary hypotheses (Uygun Tunç & Tunç, *in press*). Otherwise, we should not conclude anything beyond a very concrete, low level statistical finding (e.g., our conclusions are limited to the observed effect). In other words, the quasi-basic statement we formulated cannot have further consequences for the theoretical statement(s) under test.⁸

Since in the so-called soft sciences the derivational chains connecting various levels of scientific statements are relatively weak, the construction and testing of model hierarchies require much higher caution with auxiliary assumptions and stricter error-control at each step. In Lakatos' terms, to the extent that the derivational chains are weak, the "protective belt" of auxiliary assumptions becomes nebulous and thus absorbs the impact of disconfirming evidence. That is, when the auxiliary assumptions are not sufficiently valid, researchers can easily blame some of these *ad hoc* as alternative explanations of undesirable results of tests, rather than admitting that the theoretical prediction was wrong. To the extent that the impact of disconfirming evidence can be thus absorbed, basic statements have only limited and possibly misleading implications for the

⁸ To refer back to our previous example of the Stroop effect, a single observation of a delay in reaction time for a given group of subjects when responding to trials where the word color and word meaning are incongruent, compared to congruent, could offer only a very weak support for a universal relationship between cognitive processing speed and information congruence. Moreover, it could not by itself decide between alternative theoretical explanations of the phenomenon: e.g., whether the difference in reaction time is due to a difference in the stimulus identification, response selection, or response performance.

substantive hypothesis. The results of a statistical hypothesis test are never the final arbiter of scientific truth, as the theoretical and methodological assumptions that lead to these results should also be scrutinized in the evaluation of a scientific claim. For these reasons, a statistical hypothesis test is never the first step in a research line. Rather, it must be preceded by studies that increase the validity of the theoretical and methodological assumptions to be made in the hypothesis test, and thereby strengthen the derivational chain (Scheel et al., 2020).

3.2 Neyman-Pearson Hypothesis Tests as Methodological Decision Procedures

A number of statistical approaches to statistical inferences rely on methodological decision procedures, starting with Wald's seminal work on statistical decision functions (Wald, 1950), and culminating in Bayesian decision analysis. Although the epistemic function of dichotomous claims applies to all statistical decision procedures, we will illustrate this function in the context of the Neyman-Pearson approach to statistical inferences, which is the most widely used statistical decision procedure to make dichotomous decisions in science. In a Neyman-Pearson approach data are used to make decisions about how to act (Neyman & Pearson, 1933). Neyman (1957) calls this approach *inductive behavior*. The Neyman-Pearson approach to statistical inferences is part of statistical decision theory (Birnbaum, 1977), and a researcher who uses a Neyman-Pearson approach should be interested in deciding to at least tentatively act as if one of multiple possible hypotheses is true. Due to random error in the data, it is possible that researchers make an incorrect decision (either a false positive or a false negative). The probability of incorrect decisions can never be eliminated completely, and the Type 1 and Type 2 error rate are typically non-negligible due to resource constraints. In a Neyman-Pearson approach to statistical inferences the goal is to design studies that, in the long run, make these decision errors sufficiently rare.

It is important to distinguish the Neyman-Pearson approach to hypothesis testing that we discuss in this manuscript from the Fisherian approach to significance testing (Hubbard, 2004). According to Fisher (1956, p. 44) statistical tests lead to “a rational and well-defined measure of reluctance to the acceptance of the hypotheses they test”. The lower the p -value, the greater the

reluctance to accept the null-hypothesis. Spanos (1999) refers to the Fisherian approach as misspecification testing, as used when testing model assumptions. In a misspecification test (such as measures of goodness of fit) the null hypothesis is tested against any unspecified alternative hypothesis. Furthermore, in the Fisherian approach the p -value is not compared against an alpha level that controls Type 1 errors. Without a specified alternative hypothesis, the p -value can only be used to describe the probability of the observed or more extreme data under the null model. It is not possible to interpret p -values as a continuous measure of evidence against the null hypothesis (Johansson, 2011; Lakens, 2022), and attempts to develop the interpretation of p -values into a form of inductive reasoning known as the fiducial argument has largely failed (Zabell, 1992). Without a prespecified cut-off to reject the null-hypothesis any researcher can simply dismiss findings that falsify their predictions by arguing the p -value is not low enough, and there is no mechanism in place to keep researcher's confirmation bias in check by controlling the probability of false positive claims. According to Spanos (1999, p. xxiii), after the model assumptions have been checked, when "testing theoretical restrictions within a statistically adequate model, however, the Neyman–Pearson approach becomes the procedure of choice." As our focus lies on hypothesis tests with the goal to distinguish between a null hypothesis and a specified alternative hypothesis, in the remainder of this text we will focus on the Neyman-Pearson approach to hypothesis testing.

Because a Neyman-Pearson hypothesis test is a methodological rule that can be used to make decisions while controlling error rates in the long run, there is an intuitive connection to Popper's methodological falsificationism. Lakatos insightfully observed that statistical decision criteria have a function remarkably similar to Popperian methodological decisions; namely, they enable severe empirical tests, and thereby the falsification of statistical hypotheses. Lakatos (1978, p. 25) writes: Popper's "methodological falsificationism is the philosophical basis of some of the most interesting developments in modern statistics. The Neyman-Pearson approach rests completely on methodological falsificationism." Popper had already maintained that probabilistic (in this case, statistical) hypotheses are falsifiable if "we decide to make them falsifiable by accepting a

methodological rule” (Popper, 2002, p. 260). An important difference between methodological falsificationism compared to naïve or dogmatic falsificationism is that the first acknowledges the necessity of certain conventions before a prediction is tested, such as pre-specifying the conditions under which an empirical test would lead to the acceptance of a basic statement (Mayo, 2018; Popper, 2002; Lakatos, 1978; Uygun Tunç & Tunç, *in press*). This convention should be justified well enough to convince peers. An example of such a convention is to design a study with a desired statistical power given a pre-specified alpha level and deciding that a prediction is confirmed only when the p -value is smaller than the alpha level. Once the decision rule has been pre-specified there is no element of subjective opinion or arbitrariness in accepting or rejecting a hypothesis, and the results of empirical tests determine whether researchers should tentatively conclude if their prediction is corroborated or not.

Both Popper and Neyman are interested in determining a methodological rule that allows researchers to make decisions that will render statistical statements falsifiable in a way that is unlikely to lead them astray too often. When we “accept” or “reject” a hypothesis in a Neyman-Pearson approach to statistical inferences, we do not communicate any belief or conclusion about the substantive hypothesis. We utter a basic statement, based on a prespecified decision rule and given background assumptions, that the observed data reflect a certain state of the world. This basic statement corroborates our prediction, or not. A core feature of a basic statement is that it has a *truth value* and thus can itself be falsified by future observations. The same is true of a quasi-basic statement: this is an essential aspect of what it means to *test* a statistical hypothesis. As such, a Neyman-Pearson hypothesis test is an implementation of a methodological procedure that allows us to utter quasi-basic statements on the basis of probabilistic hypotheses about data. It is important to keep in mind that by uttering these quasi-basic statements we do not express a belief in their truth.

It is important to point out that the interpretation of an “act” by Neyman is much broader than one might intuitively think. He writes (1957, p. 10) that the concluding phase of a study involves “an act of will or a decision to take a particular action, perhaps to assume a particular

attitude towards the various sets of hypotheses.” In this sense, “acts” are defined in the Neyman-Pearson approach in such a way that they closely correspond to the decisional element concerning what will be considered as a falsifying instance in methodological falsificationism (Lakatos, 1978; Popper, 2002). Cox (1958, p. 359) writes: “it might be argued that in making an inference we are 'deciding' to make a statement of a certain type about the populations and that therefore, provided that the word decision is not interpreted too narrowly, the study of statistical decisions embraces that of inference.” Thus, dichotomous decisions in this context are not decisions to take practical action or even verdicts about the truth-value of scientific theories. They are employed simply as part of a complex methodological procedure to arrive at observational statements that can potentially be introduced to what Frick (1996) calls a “corpus of findings.” Some researchers have suggested the goal of science is not to establish claims, but to update beliefs about claims. Frick (1996) discusses how this might be feasible when the number of claims is limited, but argues it is not feasible for researchers to continuously track and rationally update beliefs about claims. Instead, findings tentatively enter the corpus of claims, until future experiments and meta-analyses suggest otherwise, based on a categorical decision following a hypothesis test.

4. THE PRAGMATIC FUNCTION OF DICHOTOMOUS CLAIMS

Just because the practice of making dichotomous claims based on statistical tests is a coherent approach to scientific inquiry from the perspective of a hypothetico-deductive philosophy of science (and specifically in methodological falsificationism) does not mean that the scientific community should adopt it. It is also important to reflect on the pragmatic aspects of a method of scientific inferences. Although the following observations are largely speculative due to a lack of empirical data, we believe dichotomous decisions have pragmatic functions that make methodological falsificationism an attractive approach to scientific inquiry, at least for certain scientific questions.

One desirable feature of cumulative science is the intersubjective testability of scientific claims through repeatable observations. Through repeated tests of hypotheses, researchers can

establish if predictions are confirmed in hypothesis tests with an expected frequency, after which the claims about observations can be (always tentatively) accepted. The main idea behind formulating statistical hypotheses in a form that can be accepted or rejected by others is that the outcome of the hypothesis test is a quasi-basic statement that can have implications for substantive hypotheses and scientific theories. A core assumption in this philosophy of science is that scientific objectivity is established and maintained only within a context of critical discussion, where individuals put forward claims that can be proven wrong. A (quasi-)basic statement is an empirical claim that others can scrutinize, dispute, challenge, and falsify. Through its dichotomous nature, a (quasi-)basic statement forces peers to take a stance: Will they accept the observation, or do they choose to reject it? If they reject the claim (assuming the research was performed competently and without bias), peers will need to justify their rejection by proposing a hypothesis that should itself be testable (where one alternative hypothesis is that the previous observation was a false positive).

This back and forth between what Popper refers to as 'conjectures and refutations' is a way to decide which of the existing theoretical accounts withstands scrutiny the best, and whether any changes to theories to accommodate new observations lead to progressive or degenerative research lines (Lakatos 1978). As long as claims withstand scrutiny, we can assume they might have some verisimilitude, or "truthlikeness" (Popper, 1962, pp. 233-234; Niiniluoto, 1987). This unique role of tests of hypotheses with respect to establishing intersubjectively agreed upon facts sets it apart from other approaches to statistical inferences, such as estimation approaches. Intersubjective agreement on a scientific fact (however tentative and fallible) requires a binary decision to reject, until there is persuasive evidence to the contrary, that something is not a fact. Estimates (such as an effect size estimate) cannot be accepted or contested, since they are not quasi-basic statements that can have truth values (although they can be transformed into claims about point predictions or range predictions that can be true or false, and which can then be tested).

One might wonder if we need dichotomous "accept" or "reject" claims that force peers to accept a statement about a phenomenon, or reject it, to have scientific progress. Would it be possible

to organize science in a way that relies less on tests of competing theories to arrive at intersubjectively established facts about phenomena? According to Tukey (1960), in order to replace dichotomous scientific claims “as the basic means of communication” between scientists, “it would be necessary to rearrange and replan the entire fabric of science.” Alternative approaches to how to advance scientific knowledge seem feasible if the broader scientific community agrees on the research questions that need to be investigated, and methods to be utilized, and coordinate their research efforts. We think that such a level of non-adversarial collaboration is a nearly unattainable goal for any scientific discipline, because even among scientists working within well-established and widely shared paradigms there will often be disagreement about contentious theoretical and methodological issues (Kuhn, 1996, pp. 29-30), and researchers will often need to convince peers of the value of their hypotheses or theoretical models, and of the validity of their auxiliary assumptions. Note that our point is not limited to individual-oriented organizations of scientific activity or to pre-paradigmatic sciences characterized by theoretical dissensus, but for example extends to collaborative experiments to test theories formulated within the standard model in high energy physics.

One pragmatic consequence of methodological falsificationism, and the practice of making falsifiable scientific claims, might be that it motivates other scholars to engage critically with these claims. Researchers who want their claims to be taken seriously depend, in part, on independent replication and extension studies by peers. Such studies reveal whether claims can withstand falsification attempts and are especially convincing when claims withstand falsification by skeptics. Because our peers have limited resources, they will need some motivation to engage with our claims, and one way such motivation might emerge is through a desire to prove someone else wrong. We feel an analogy to Cunningham's Law applies in science, in that just as how “the best way to get the right answer on the internet is not to ask a question; it's to post the wrong answer,” due to the current lack of coordination in science “the best way to generate robust scientific knowledge is not to ask a question; it's to publish a claim other researchers think is wrong.”

Dichotomous decisions based on statistical tests thus allow the scientific community to intersubjectively test, agree upon or reject scientific claims that are based on data involving non-negligible random error. Thereby researchers can reach intersubjective agreement on scientific facts through practicing mutual criticism, which cannot be achieved by merely publicly sharing probabilistic statements. One core ingredient of this process of mutual criticism is the existence of justifiable decision criteria to be used in statistical inference. Since all rational decisions require some decision criteria based on certain standards, rules or norms, statistical inference similarly involves decision criteria based on thresholds. In practice, most researchers use an alpha level of 0.05 as a decision threshold to conclude a prediction is corroborated, or not. Nothing we have said so far points us towards how and where to set such a decision threshold, beyond explaining the general rationale for the existence of decision criteria in statistical inference. So, in the next section, it might be worthwhile to explore how the highly disputed significance level of 0.05 might have survived until today to fulfill the epistemic and pragmatic functions of dichotomous decisions in statistical tests.

5. THE CONVENTION OF AN ALPHA LEVEL OF 0.05

Although the widespread use of an alpha level of 0.05 is largely a historical accident (Kennedy-Shaffer, 2019), it has led to an almost miraculous level of agreement among scientists about the Type 1 error rate that is deemed acceptable when making scientific claims. Such an institutional norm is not necessary for the epistemological or pragmatic function of dichotomous claims. Epistemologically, probabilistic claims need to be made falsifiable by drawing the line somewhere, and as long as the decision for the error rate is specified in advance, every peer can evaluate how risky the prediction was (i.e., an alpha level of 0.5 is much less risky than an alpha level of 0.001). Ideally, an alpha level is justified on the basis of a cost-benefit analysis (Lakens et al., 2018; Neyman & Pearson, 1933). However, scientific claims often (although not always) have little direct impact beyond the costs and benefits for scientists engaging with a scientific claim (Tukey, 1960), which limits the situations in which the alpha level can be specified based on a cost-

benefit justification. Although alternative approaches to justify the alpha level exist, such as minimizing combined error rates (Mudge, Baker, Edge & Houlahan, 2012) or compromise power analysis (Erdfelder, Faul & Buchner, 1996), the default alpha level of 0.05 can be seen as an institutional norm that is a “convenient convention” (Fisher, 1971, p. 13) for pre-specifying the error rate at which we accept scientific claims by peers. Although there is no reason why a default alpha level should be 0.05 (and there have been fields where conventions about the default alpha level have changed, such as the 5 sigma rule in some research lines in physics), we might speculate why the default 5% alpha level has survived as a practice for nearly a century.

One possible reason is that, as far as conventions go, an alpha level of 5% might be low enough such that peers take any claims made with this error rate seriously, while at the same time being high enough such that peers will be motivated to perform an independent replication study to increase or decrease our confidence in the claim. Although lower error rates would establish claims more convincingly, this would also require more resources. One might speculate that in research areas where not every claim is important enough to warrant investing the resources required to establish claims with low error rates (Isager et al., 2021), an alpha level of 5% has a pragmatic function in facilitating conjectures and refutations in fields that otherwise lack a coordinated approach to knowledge generation but are faced with limited resources.

The past conventional use of an alpha level of 0.05 can also be explained by the requirement in methodological falsificationism that statistical decision procedures are specified before the data is analyzed (see Popper, 2002, sec. 19-20; Lakatos, 1978, pp. 23-28). If researchers are allowed to set the alpha level after looking at the data, there is a possibility that confirmation bias (or more intentional falsification-deflecting strategies) influences the choice of an alpha level (Gosset, 1904). An additional reason for a conventional alpha level of 0.05 is that before the rise of the internet it was difficult to transparently communicate the pre-specified alpha level for any individual test to peers. The use of a default alpha level therefore effectively functioned as a pre-specification. For a convention to work as a universal pre-specification, it must be accepted by nearly everyone, and be

extremely resistant to change. If more than a single conventional alpha level exists, this introduces the risk that confirmation bias influences the choice of an alpha level. Now that researchers can easily preregister their statistical analysis plan, the possibility of confirmation bias can be mitigated, which makes it possible to use well-justified non-conventional alpha levels in the future (Maier & Lakens, 2022).

6. DISCUSSION

From a purely statistical perspective, any dichotomous decision based on continuous test statistics throws away information, and if we would solely focus on statistical inferences an argument can be made to “accept uncertainty” and “be thoughtful”, and state that results should not be “categorized into any number of groups, based on arbitrary p -value thresholds” (Wasserstein, Schirm & Lazar, 2019). Dichotomous claims, however, fulfil a crucial epistemic function as well as a pragmatic function, by allowing researchers to make existence claims on the basis of data involving non-negligible random error. These quasi-basic statements play an important role in creating and testing scientific theories, and invite refutation attempts by peers. Statistical arguments against dichotomous conclusions are largely irrelevant when we evaluate the epistemic and pragmatic functions of dichotomous claims in science.

Extant arguments for the practice of making dichotomous claims in science following a methodological decision procedure, such as null-hypothesis significance testing, all share the core idea that these methodological procedures allow researchers to make ordinal claims about the direction of an effect while controlling error rates (e.g., Abelson, 1997; Chow, 1988; Hagen, 1997; Meehl, 1990; Kline, 2004; Nickerson, 2000). As Frick (1996) writes “An ordinal claim can be defined as one that does not specify the size of effect; alternatively, it could be defined as a claim that specifies only the order of conditions, the order of effects, or the direction of a correlation.” If we look at the use of hypothesis tests in the scientific literature, researchers often seem to be interested in making ordinal claims based on hypothesis tests (Lakens, 2021). These defenses of statistical hypothesis tests are perfectly aligned with their epistemic function in a methodological

falsificationist framework, which consists in enabling researchers to transform data into quasi-basic statements about whether a phenomenon has been observed (with a certain error rate, based on a methodological decision procedure).

Researchers have discussed the strengths and weakness of dichotomous inferences based on p-values for decades. A recent special issue in *The American Statistician* contained 43 contributions on how to 'move beyond $p < .05$ '. An evaluation of the ban of NHST in the journal *Basic and Applied Social Psychology* (Fricker, Burke, Han, & Woodall, 2019) revealed it reduced the rigor with which claims are made in the scientific literature, leading the authors to conclude "the ban seems to be allowing authors to make less substantiated claims, the net result of which could be misleading to readers and could well result in a body of literature within which there are a greater number of irreproducible effects." The most common recommendation in the special issue (present in five articles) is to use interval hypothesis tests instead of null hypothesis tests (e.g., Anderson, 2019; Betensky, 2019; Pogrow, 2019). These tests incorporate which effect sizes are considered meaningful into the statistical inference process, and importantly, still lead to dichotomous claims (see also Lakens, 2021). Goodman, Spruill, and Komaroff (2019) examine an effect size plus p-value decision criterion, and write (p. 178): "the authors believe that the NHST model can still have a place in scientific research" but argue that "If the p -value criterion is met, it should *also* be assessed whether a meaningful minimum effect size was observed." Blume and colleagues (2019) state that p-values might be criticized in the literature, but "Having a gross indicator for when a set of data are sufficient to separate signal from noise is not a bad idea" and propose an interval hypothesis test that is conceptually similar to an equivalence test (see Lakens & Delacre, 2020).

Several other papers support the use of dichotomous claims. For example, Krueger and Heck (2019) write "As experimentalists, we are reluctant to relinquish dichotomous decision-making entirely." Colquhoun (2019, p. 193) writes "Users will continue to use statistical tests, with the aim of trying to avoid making fools of themselves by publishing false positive results. That is entirely laudable" but recommends scientists to supplement p-values with an evaluation of the false

positive risk. Matthews (2019) suggests the use of Bayesian credible intervals instead of p-values, which still leads to a dichotomous inference, and writes (p. 208): "it should be stressed that dichotomization per se has never been the problem with NHST; it is the actions that flow from it." Gannon, de Bragança Pereira, and Polpo (2019) propose changing the alpha level as a function of the sample size, which improves the current use of hypothesis tests, but that (p. 213) "doesn't 'throw out the baby with the bath water,' retaining the useful concept of statistical significance and the same operational procedures as currently used tests". Billheimer (2019) argues for a focus on prediction, but explicitly does so within a dichotomous decision making framework where future predictions can be corroborated or refuted. Manski (2019) argues for the adoption of statistical decision theory. It might be considered surprising that so many contributions to this special issue on 'moving beyond $p < .05$ ' explicitly support the use of dichotomous statistical inferences (especially given that many contributions in the special issue did not focus on applied statistics, but on topics such as history, education, or reward structures in science). It reveals that the use of statistical tests to make dichotomous claims might be criticized by some, but at the same time continues to be supported by many.

Having said that, it is important to stress that dichotomous decisions have a very specific function in tests of predictions from a falsificationist point of view. Researchers should realize that if a test is not built on a strong derivational chain, falsifying a prediction does little to inform us about the substantive hypothesis and the results may actually be uninterpretable from the point of theory evaluation. A misunderstanding of their specific function might cause the layman and novice to assign dichotomous claims an unwarranted kind and degree of confidence. Our hope is that our defense of the role of dichotomous claims in statistical inferences also highlights their limitation when making theoretical inferences. Because researchers are rarely taught about the epistemic function of dichotomous claims within scientific reasoning, misapplications of statistical decision procedures are common. Misapplications of statistical inferences are not unique to dichotomous

claims based on statistical hypothesis tests, but as statistical hypothesis tests are so widely used in science, improving education about the correct use of these procedures is increasingly important.

Dichotomous claims are coherent, legitimate, and necessary from a methodological falsificationist perspective on scientific reasoning. However, researchers do not have to adopt methodological falsificationism. It is not surprising that criticisms of the common practice to make dichotomous claims are often accompanied by a philosophy of science that is not built on falsificationist theory testing, but on more descriptive or inductive approaches to science. For example, Rozeboom (1960) writes: "But the primary aim of a scientific experiment is not to precipitate decisions, but to make an appropriate adjustment in the degree to which one accepts, or believes, the hypothesis or hypotheses being tested." Compare this to Frick (1996), who in his article "The appropriate use of null hypothesis testing" writes: "scientists act like the goal of science is to collect a corpus of claims that are considered to be established." Both positions are logically coherent, depending on the approach to scientific inquiry one adopts. There is room to choose, and methodological falsificationism is only one of many possible approaches to scientific reasoning.

We believe it is problematic to assume that abolishing dichotomous claims will be the sword of Alexander that can cut the Gordian knot of weak scientific inferences. Dichotomous claims based on statistical hypothesis testing have been criticized since the days they were proposed (see Boring, 1919), and one of the main arguments to abandon dichotomous inferences is that it leads to widespread misuse of statistical tests. Although we acknowledge that misuse is problematic, it is not a sufficient basis for denying that dichotomous claims perform important epistemic and pragmatic functions in scientific inference. Thus, particularly if a researcher wants to deductively test theory-driven hypotheses based on data containing non-negligible random error, abandoning dichotomous claims on the odd chance that it will magically solve the misuse of statistical tests would amount to throwing the baby out with the bathwater (Hand, 2021; Mayo, 2018). Too often, arguments in favor of abandoning dichotomous decisions in hypothesis testing resemble a call for a ritual sacrifice. What seems to be expected from such a sacrifice is to trigger a fundamental perspective shift on

how various uncertainties inherent to theory evaluation are handled by researchers. However, abandoning dichotomous claims in the hope that doing so will improve scientific inferences is an overly simplistic solution that overlooks the important epistemic and pragmatic functions that these claims play in investigating scientific questions across disciplines.

References

- Abelson, R. P. (1997). On the surprising longevity of flogged horses: Why there is a case for the significance test. *Psychological Science*, 8, 12–15. DOI: [10.1111/j.1467-9280.1997.tb00536.x](https://doi.org/10.1111/j.1467-9280.1997.tb00536.x)
- Anderson, A. A. (2019). Assessing Statistical Results: Magnitude, Precision, and Model Uncertainty. *The American Statistician*, 73(sup1), 118–121. <https://doi.org/10.1080/00031305.2018.1537889>
- Betensky, R. A. (2019). The p-Value Requires Context, Not a Threshold. *The American Statistician*, 73(sup1), 115–117. <https://doi.org/10.1080/00031305.2018.1529624>
- Billheimer, D. (2019). Predictive Inference and Scientific Reproducibility. *The American Statistician*, 73(sup1), 291–295. <https://doi.org/10.1080/00031305.2018.1518270>
- Birnbaum, A. (1977). The Neyman-Pearson theory as decision theory, and as inference theory; with a criticism of the Lindley-Savage argument for Bayesian theory. *Synthese*, 36, 19–49.
- Blume, J. D., Greevy, R. A., Welty, V. F., Smith, J. R., & Dupont, W. D. (2019). An Introduction to Second-Generation p-Values. *The American Statistician*, 73(sup1), 157–167. <https://doi.org/10.1080/00031305.2018.1537893>
- Bogen, J., and Woodward, J. (1988). Saving the phenomena. *The Philosophical Review*, 97, 303–352. DOI: [10.2307/2185445](https://doi.org/10.2307/2185445)
- Boring, E. G. (1919). Mathematical vs. scientific significance. *Psychological Bulletin*, 16, 335–338. DOI: [10.1037/h0074554](https://doi.org/10.1037/h0074554)
- Chow, S. L. (1988). Significance test or effect size? *Psychological Bulletin*, 103, 105–110. DOI: [10.1037/0033-2909.103.1.105](https://doi.org/10.1037/0033-2909.103.1.105)
- Colquhoun, D. (2019). The False Positive Risk: A Proposal Concerning What to Do About p-Values. *The American Statistician*, 73(sup1), 192–201. <https://doi.org/10.1080/00031305.2018.1529622>

- Cox, D. R. (1958). Some problems connected with statistical inference. *Annals of Mathematical Statistics*, 29, 357–372. DOI: [10.1214/aoms/1177706618](https://doi.org/10.1214/aoms/1177706618)
- Erdfelder, E., Faul, F., and Buchner, A. (1996). G*POWER: A general power analysis program. *Behavior Research Methods, Instruments, & Computers*, 28, 1–11. DOI: [10.3758/BF03203630](https://doi.org/10.3758/BF03203630)
- Fisher, R. A. (1956). *Statistical methods and scientific inference* (Vol. viii). Hafner Publishing Co.
- Fisher, R. A. (1971). *The Design of Experiments* (9th ed.). Macmillan Pub Co.
- Frick, R. W. (1996). The appropriate use of null hypothesis testing. *Psychological Methods*, 1, 379–390. DOI: [10.1037/1082-989X.1.4.379](https://doi.org/10.1037/1082-989X.1.4.379)
- Fricker, R. D., Burke, K., Han, X., & Woodall, W. H. (2019). Assessing the Statistical Analyses Used in Basic and Applied Social Psychology After Their p-Value Ban. *The American Statistician*, 73(sup1), 374–384. <https://doi.org/10.1080/00031305.2018.1537892>
- Gannon, M. A., de Bragança Pereira, C. A., & Polpo, A. (2019). Blending Bayesian and Classical Tools to Define Optimal Sample-Size-Dependent Significance Levels. *The American Statistician*, 73(sup1), 213–222. <https://doi.org/10.1080/00031305.2018.1518268>
- Gosset, W. S. (1904). The application of the “Law of Error” to the work of the brewery (1 vol 8; pp. 3–16). Arthur Guinness & Son, Ltd.
- Gibson, E. W. (2020). The role of p -values in judging the strength of evidence and realistic replication expectations. *Statistics in Biopharmaceutical Research*, 13, 6–18 DOI: [10.1080/19466315.2020.1724560](https://doi.org/10.1080/19466315.2020.1724560)
- Goodman, W. M., Spruill, S. E., & Komaroff, E. (2019). A Proposed Hybrid Effect Size Plus p -Value Criterion: Empirical Evidence Supporting its Use. *The American Statistician*, 73(sup1), 168–185. <https://doi.org/10.1080/00031305.2018.1564697>
- Greenland, S. (2017). Invited commentary: The need for cognitive science in methodology. *American Journal of Epidemiology*, 186, 639–645.

- Greenland, S., and Poole, C. (2013). Living with p -Values: Resurrecting a Bayesian Perspective on Frequentist Statistics. *Epidemiology*, 24, 62–68, DOI: [10.1097/EDE.0b013e3182785741](https://doi.org/10.1097/EDE.0b013e3182785741)
- Hacking, I. (1965). *Logic of Statistical Inference*. Cambridge University Press.
- Hagen, R. L. (1997). In praise of the null hypothesis statistical test. *The American Psychologist*, 52, 15–24.
- Hager, W. (2000). About some misconceptions and the discontent with statistical tests in psychology. *Methods of Psychological Research Online*, 5(1), 31.
- Hager, W. (2013). The statistical theories of Fisher and of Neyman and Pearson: A methodological perspective. *Theory & Psychology*, 23, 251–270, DOI: [10.1177/0959354312465483](https://doi.org/10.1177/0959354312465483)
- Hand, D. J. (2021). Trustworthiness of statistical inference. Preprint, ArXiv:2105.03948 [Stat]. Available at <http://arxiv.org/abs/2105.03948>
- Hempel, C. (1966). *Philosophy of Natural Science* (1st ed.). Pearson.
- Hubbard, R. (2004). Alphabet Soup: Blurring the Distinctions Between p 's and a 's in Psychological Research. *Theory & Psychology*, 14(3), 295–327. <https://doi.org/10.1177/0959354304043638>
- Hubbard, R., and Bayarri, M. J. (2003). Confusion over measures of evidence (p 's) versus errors (α 's) in classical statistical testing. *The American Statistician*, 57, 171–178, DOI: [10.1198/0003130031856](https://doi.org/10.1198/0003130031856)
- Isager, P. M., van Aert, R. C. M., Bahník, Š., Brandt, M. J., DeSoto, K. A., Giner-Sorolla, R., Krueger, J. I., Perugini, M., Ropovik, I., van 't Veer, A. E., Vranka, M., & Lakens, D. (2021). Deciding what to replicate: A decision model for replication study selection under resource and knowledge constraints. *Psychological Methods*. <https://doi.org/10.1037/met0000438>
- Johansson, T. (2011). Hail the impossible: P -values, evidence, and likelihood. *Scandinavian Journal of Psychology*, 52(2), 113–125. <https://doi.org/10.1111/j.1467-9450.2010.00852.x>

- Kennedy-Shaffer, L. (2019). Before $p < 0.05$ to beyond $p < 0.05$: Using history to contextualize p -values and significance testing. *The American Statistician*, 73, 82–90, DOI: [10.1080/00031305.2018.1537891](https://doi.org/10.1080/00031305.2018.1537891)
- Kline, R. B. (2004). *Beyond Significance Testing: Reforming Data Analysis Methods in Behavioral Research* (1st ed). American Psychological Association.
- Kuhn, T. S. (1996). *The Structure of Scientific Revolutions* (3rd ed.). University of Chicago Press.
- Krueger, J. I., & Heck, P. R. (2019). Putting the P-Value in its Place. *The American Statistician*, 73(sup1), 122–128. <https://doi.org/10.1080/00031305.2018.1470033>
- Lakatos, I. (1978). *The Methodology of Scientific Research Programmes*. Cambridge University Press.
- Lakens, D. (2021). The practical alternative to the p value is the correctly used p value. *Perspectives on Psychological Science*, 16(3), 639–648. <https://doi.org/10.1177/1745691620958012>
- Lakens, D. (2022). Why P values are not measures of evidence. *Trends in Ecology & Evolution*. <https://doi.org/10.1016/j.tree.2021.12.006>
- Lakens, D., & Delacre, M. (2020). Equivalence Testing and the Second Generation P -Value. *Meta-Psychology*, 4, 1–11. <https://doi.org/10.15626/MP.2018.933>
- Lakens, D., Adolphi, F. G., Albers, C. J., Anvari, F., Apps, M. A. J., Argamon, S. E., Baguley, T., Becker, R. B., Benning, S. D., Bradford, D. E., Buchanan, E. M., Caldwell, A. R., Calster, B., Carlsson, R., Chen, S.-C., Chung, B., Colling, L. J., Collins, G. S., Crook, Z., ... Zwaan, R. A. (2018). Justify your alpha. *Nature Human Behaviour*, 2, 168–171, DOI: [10.1038/s41562-018-0311-x](https://doi.org/10.1038/s41562-018-0311-x)
- Maier, M., & Lakens, D. (2022). Justify Your Alpha: A Primer on Two Practical Approaches. *Advances in Methods and Practices in Psychological Science*, 5(2), 25152459221080396. <https://doi.org/10.1177/25152459221080396>

- Manski, C. F. (2019). Treatment Choice With Trial Data: Statistical Decision Theory Should Supplant Hypothesis Testing. *The American Statistician*, 73(sup1), 296–304.
<https://doi.org/10.1080/00031305.2018.1513377>
- Matthews, R. A. J. (2019). Moving Towards the Post $p < 0.05$ Era via the Analysis of Credibility. *The American Statistician*, 73(sup1), 202–212.
<https://doi.org/10.1080/00031305.2018.1543136>
- Mayo, D. G. (1996). *Error and the Growth of Experimental Knowledge* (1st ed.). University of Chicago Press.
- (2018). *Statistical inference as severe testing: How to get beyond the statistics wars*. Cambridge University Press.
- Mayo, D. G., and Spanos, A. (2011). Error statistics. *Philosophy of Statistics*, 7, 152–198. DOI: [10.1016/B978-0-444-51862-0.50005-8](https://doi.org/10.1016/B978-0-444-51862-0.50005-8)
- Meehl, P. E. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology*, 46, 806–834, DOI: [10.1037/0022-006X.46.4.806](https://doi.org/10.1037/0022-006X.46.4.806)
- (1990). Appraising and amending theories: The strategy of Lakatosian defense and two principles that warrant it. *Psychological Inquiry*, 1, 108–141. DOI: [10.1207/s15327965pli0102_1](https://doi.org/10.1207/s15327965pli0102_1)
- Mudge, J. F., Baker, L. F., Edge, C. B., and Houlahan, J. E. (2012). Setting an optimal α that minimizes errors in null hypothesis significance tests. *PLOS ONE*, 7, e32734, DOI: [10.1371/journal.pone.0032734](https://doi.org/10.1371/journal.pone.0032734)
- Neyman, J. (1957). ‘Inductive behavior’ as a basic concept of philosophy of science. *Revue de l’Institut International de Statistique / Review of the International Statistical Institute*, 25, 7-22. DOI: [10.2307/1401671](https://doi.org/10.2307/1401671)

- Neyman, J., and Pearson, E. S. (1933). On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 231, 289–337. DOI: [10.1098/rsta.1933.0009](https://doi.org/10.1098/rsta.1933.0009)
- Nickerson, R. S. (2000). Null hypothesis significance testing: A review of an old and continuing controversy. *Psychological Methods*, 5, 241–301. DOI: [10.1037//1082-989X.5.2.241](https://doi.org/10.1037//1082-989X.5.2.241)
- Niiniluoto, I. (1987). *Truthlikeness*. Springer.
- Pogrow, S. (2019). How Effect Size (Practical Significance) Misleads Clinical Practice: The Case for Switching to Practical Benefit to Assess Applied Research Findings. *The American Statistician*, 73(sup1), 223–234. <https://doi.org/10.1080/00031305.2018.1549101>
- Popper, K. R. (1962). *Conjectures and Refutations: The growth of scientific knowledge*. Basic Books.
- (2002). *The Logic of Scientific Discovery*. Routledge. First published in 1959.
- Rozeboom, W. W. (1960). The fallacy of the null-hypothesis significance test. *Psychological Bulletin*, 57, 416–428. DOI: [10.1037/h0042040](https://doi.org/10.1037/h0042040)
- Scheel, A. M., Tiokhin, L., Isager, P. M., and Lakens, D. (2020). Why hypothesis testers should spend less time testing hypotheses. *Perspectives on Psychological Science*. DOI: [10.1177/1745691620966795](https://doi.org/10.1177/1745691620966795)
- Spanos, A. (1999). *Probability theory and statistical inference: Econometric modeling with observational data*. Cambridge University Press.
- Spanos, A., and Mayo, D. G. (2015). Error statistical modeling and inference: Where methodology meets ontology. *Synthese*, 192, 3533–3555. DOI: [10.1007/s11229-015-0744-y](https://doi.org/10.1007/s11229-015-0744-y)
- Suppes, P. (1962). Models of data. In E. Nagel, P. Suppes and A. Tarski eds., *Logic, Methodology, and Philosophy of Science: Proceedings of the 1960 International Congress*. Stanford University Press, pp. 252–261.
- (1974). The structure of theories and the analysis of data. In F. Suppe ed., *The structure of scientific theories*. Urbana: University of Illinois Press, pp. 266–283.

- Tukey, J. W. (1960). Conclusions vs decisions. *Technometrics*, 2, 423–433. DOI: [10.1080/00401706.1960.10489909](https://doi.org/10.1080/00401706.1960.10489909)
- Trafimow, D. (2019). Five Nonobvious Changes in Editorial Practice for Editors and Reviewers to Consider When Evaluating Submissions in a Post $p < 0.05$ Universe. *The American Statistician*, 73(sup1), 340–345. <https://doi.org/10.1080/00031305.2018.1537888>
- Uygun Tunç, D., and Tunç, M. N. (*in press*). A Falsificationist Treatment of Auxiliary Hypotheses in Social and Behavioral Sciences: Systematic Replications Framework. *Meta-Psychology*. Preprint DOI: [10.31234/osf.io/pdm7y](https://doi.org/10.31234/osf.io/pdm7y)
- Wald, A. (1950). *Statistical Decision Functions* (1st ed.). Wiley.
- Wasserstein, R. L., Schirm, A. L., and Lazar, N. A. (2019). Moving to a world beyond ‘ $p < 0.05$ ’,” *The American Statistician*, 73, 1–19, DOI: [10.1080/00031305.2019.1583913](https://doi.org/10.1080/00031305.2019.1583913)
- Woodward, J. (1989). Data and phenomena. *Synthese*, 79, 393–472.
- (2000). Data, phenomena, and reliability. *Philosophy of Science*, 67, S163–S179. DOI: [10.1086/392817](https://doi.org/10.1086/392817)
- Zabell, S. L. (1992). R. A. Fisher and Fiducial Argument. *Statistical Science*, 7(3), 369–387. <https://doi.org/10.1214/ss/1177011233>