

Discussion Paper

No. 2019-65 | November 18, 2019 | <http://www.economics-ejournal.org/economics/discussionpapers/2019-65>

Inference in economic experiments

*Norbert Hirschauer, Sven Grüner, Oliver Mußhoff,
and Claudia Becker*

Abstract

Replication crisis and debates about p-values have raised doubts about what we can statistically infer from research findings, both in experimental and observational studies. With a view to the present debate on inferential errors, this paper systematizes and discusses experimental designs with regard to the inferences that can and – perhaps more important – that cannot be made from particular designs.

JEL B41 C18 C90

Keywords Economic experiments; ceteris paribus; confounders; control, inference; Internal/external validity; randomization; random sampling; superpopulation

Authors

Norbert Hirschauer, Martin Luther University Halle-Wittenberg, Halle (Saale), Germany, norbert.hirschauer@landw.uni-halle.de

Sven Grüner, Martin Luther University Halle-Wittenberg, Halle (Saale), Germany

Oliver Mußhoff, Georg August University Göttingen, Göttingen, Germany

Claudia Becker, Martin Luther University Halle-Wittenberg, Halle (Saale), Germany

Citation Norbert Hirschauer, Sven Grüner, Oliver Mußhoff, and Claudia Becker (2019). Inference in economic experiments. Economics Discussion Papers, No 2019-65, Kiel Institute for the World Economy.

<http://www.economics-ejournal.org/economics/discussionpapers/2019-65>

Received October 15, 2019 Accepted as Economics Discussion Paper November 5, 2019

Published November 18, 2019

© Author(s) 2019. Licensed under the [Creative Commons License - Attribution 4.0 International \(CC BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)

1 Introduction

Starting with CHAMBERLIN (1948), SAUERMANN and SELTEN (1959), HOGGATT (1959), SIEGEL and FOURAKER (1960), and SMITH (1962), economists have increasingly adopted *experimental* designs over the last decades. Their motivation to do so was to obtain – compared to *observational* studies – more trustworthy information about the causalities that govern human behavior. Unfortunately, it seems that in the process of adopting the experimental method, no tightly inference-focused systematization of economic experiments has emerged. Some scholars use randomization as the defining quality and equate “experiments” with “randomized controlled trials” (ATHEY and IMBEN 2017). Despite ensuing changes in the nature of feasible inferences, other researchers include non-randomized designs into the definition as long as behavioral data are generated through a treatment manipulation (HARRISON and LIST 2004). One might speculate that economists tend to conceptually stretch the term “experiment” because the seemingly attractive label suggests that they have adopted “trustworthy” research methods that are comparable to those in the natural sciences. Whatever the reason, confusion regarding the different types of research designs that are labeled as experiments entails the risk of inferential errors.

The inferences that can be made from controlled experiments based on the *ceteris paribus* approach where “everything else but the item under investigation is held constant” (SAMUELSON and NORDHAUS 1985: 8) are different from those that can be made from observational studies. The former rely on the research design to ex ante ensure *ceteris paribus* conditions that facilitate the identification of causal treatment effects. Observational studies, in contrast, rely on an ex post control of confounders through statistical modeling that, despite attempts to move from correlation to causation, does not provide a way of ascertaining causal relationships that is as reliable as a strong ex ante research design (ATHEY and IMBEN 2017). But even within experimental approaches, different designs facilitate different inferences.

In this paper, we address the question of statistical and scientific induction and, more particularly, the role of the p -value for making inferences beyond the confines of a particular experimental study. We aim at an adequate differentiation of experimental designs that contributes to a better understanding of the inferences that can and – perhaps more important – that cannot be made from particular designs. For the sake of simplicity, we limit the discussion of treatment comparison to binary treatments.

2 Experiments aimed at identifying causal treatment effects

The label “experiment” is first of all used for studies that, instead of using survey data or pre-existing observational data, are based on a deliberate intervention (treatment) *and* a design-based control over confounders. Identifying the effects of the treatment on the units (subjects) under study requires a comparison; often no-treatment observations are compared to with-treatment observations. Two different designs are used to ensure control and thus *ceteris paribus* conditions: (1) Randomized controlled trials rely on a between-subject design and randomization to generate equivalence between compared groups; i.e. we randomly assign subjects to treatments to ensure that known *and* unknown confounders are balanced across treatment groups (statistical independence). (2) Non-randomized controlled trials, in contrast, rely on a within-subject design and before-and-after comparisons; i.e. we try to hold everything but the treatment constant over time and compare the before-and-after-treatment outcomes for all subjects who participate in the experiment.¹

The persuasiveness of causal claims depends on the credibility of the alleged control. Comparing randomized treatment groups is generally held to be a more convincing device to identify causal relationships than before-and-after treatment comparisons (CHARNESS et al. 2012). This is due to the fact that

¹ The term “non-randomized controlled trial” is also used for between-subject designs when the technique of assigning subjects to treatments, e.g. alternate assignment, is not truly random but claimed to be as good as a random.

randomization balances known *and* unknown confounders across treatment groups and thus ensures statistical independence.² In contrast, efforts to hold everything else but the treatment constant over time in before-and-after comparisons are limited by the researcher’s capacity to identify and fix confounders. A particular threat to causal inference arises when subjects’ properties change through treatment exposure. That is, holding “everything” but the treatment constant over time can be difficult because sequentially exposing subjects to multiple treatments may cause order effects that violate the *ceteris paribus* condition (CHARNESS et al. 2012). However, as CZIBOR et al. (2019) emphasize, within-subject designs also have their advantages: besides the fact that they can more effectively make use of small experimental groups, they facilitate the identification of higher moments of the distribution. Whereas between-subject designs are limited to estimating average treatment effects, within-subject designs enable researchers to look at quantiles and assess heterogeneous treatment effects among subjects.

Due to the particular credibility of randomization as a means to establish control over confounders, the use of the term “experiment” – accompanied by the label “natural” – has even been extended to observational settings where, instead of a deliberate treatment manipulation by a researcher, the socio-economic or natural environment has randomly “assigned treatments” among some set of units. Regarding this terminology, DUNNING (2013: 16) notes “that the label ‘natural experiment’ is perhaps unfortunate. [...], the social and political forces that give rise to as-if random assignment of interventions are not generally ‘natural’ in any ordinary sense of that term. [...] and], natural experiments are observational studies, not true experiments, again, because they lack an experimental manipulation. In sum, natural experiments are neither natural nor experiments” but may be structurally close to randomization.³

3 Inferences in experiments based on treatment comparisons

Sharing the essential approach of providing for an *ex ante*, design-based control over confounders through the introduction of a well-defined treatment into an otherwise controlled environment, randomized-treatment-group comparisons and before-and-after-treatment comparisons facilitate causal inferences. The meaning of statistical inference and the *p*-value, however, are different in the two cases. In *randomized-treatment-group comparisons*, the *p*-value linked to the treatment difference is usually based on the approximation of the *randomization distribution* (cf. RAMSEY and SCHAFER 2013), i.e. the distribution of the *difference between group averages* and the standard error used in a two-independent-sample *t*-test. Regardless of how participating subjects were recruited, the resulting *p*-value targets the following question: when there is no treatment-group difference, how likely is it that we would find a difference as large as (or larger than) the one observed when we repeatedly assigned the experimental subjects at random to the treatments under investigation (VOGT et al. 2014: 242). In randomized controlled experiments, the evaluation of internal validity and causal inference can be aided by statistical inference based on the *p*-value, which represents a continuous measure of the strength of evidence against the null hypothesis of there being no treatment effect in the group of experimental subjects. While scientific inferences beyond the confines of the experimental group under study are often desired, it must be recognized that randomization-based inference is no help for generalizing from experimental subjects to a broader population from which they have been recruited. Using statistical inference to help make such generalizations would require that, besides being randomized, the recruited experimental

² DUFLO et al. (2007) note that randomization only achieves that confounders are balanced across treatments “in expectation;” i.e. balance is only ensured if we have a “sufficiently” large experimental group to start with (law of large numbers). In the case of a very small experimental group, an observed difference between two randomized treatment groups could easily be contaminated by unbalanced confounders.

³ HARRISON and LIST (2004) speak of “*doing* natural field experiments” to tag field experiments with subjects from the social group of interest and a *covert* manipulation of subjects’ real-life environment. For terminological clarity it should be noted that such field experiments are not natural experiments but deliberate interventions.

subjects had been randomly drawn from a defined parent population. If they are not, extending inference from the experimental subjects to any broader group must be based on scientific reasoning beyond statistical measures such as p -values. This implies accounting for contextual factors and the entirety of available knowledge including external evidence for the phenomenon under study.⁴

When we not only randomize a given group of experimental subjects but also recruit them from a defined parent population through random sampling, the question arises of how to link randomization-based inference, which is concerned with internal validity and causality, to sampling-based inference, which is concerned with external validity and generalization towards the broader parent population. The “true” standard error of the randomization distribution would reflect the idea of frequently re-randomizing a *given* group of, let’s say, $n=100$ subjects in hypothetical experimental replications. The standard error in a two-independent-*sample* t -test, in contrast, presumes that we repeatedly draw random samples of $n=100$ subjects from a population before carrying out the randomized experiment. As stated above, two-sample t -tests are often also used for causal inferences from randomized-treatment-group comparisons even though they are conceptually based on random sampling from populations. If we accept the sampling-based standard error as an approximation of the randomization-based standard error (ATHEY and IMBEN 2017) – it is an upwardly-biased approximation because it considers sampling error in addition to randomization error – the resulting p -value can be used as an aid for simultaneously assessing internal and external validity. One should always be explicit about the fact, however, that the interpretation of the p -value must be strictly limited to causal inferences *within* the *given* group of experimental subjects when the group of experimental subjects was not recruited through random sampling.

Contrary to randomization, a p -value associated with the treatment difference in *before-and-after-treatment comparisons* is conceptually per se based on random sampling and the *sampling distribution*, i.e. the distribution of the *average individual before-and-after difference* and thus the standard error in a paired t -test. This is just another label for a one-sample t -test on the variable “individual before-and-after differences.” Statistical inference based on the one-sample p -value implies that we concern ourselves with the question of what we can learn about the population mean from a random sample. In other words, we are asking the following question: assuming there is no difference in the population, how likely is it that we would find an average before-and-after difference as large as (or larger than) the one observed if we carried out very many statistical replications and subjected repeatedly drawn random samples to the same treatment procedure. Therefore, our p -value is a continuous measure of the strength of evidence against the null of there being no treatment effect in the parent population. While being an inferential tool to help make generalizations from the sample of experimental subjects to a broader population (external validity), it must be recognized that a p -value in before-and-after comparisons is no help whatsoever for assessing causality. Instead, causality claims hinge on the credibility of the *ceteris paribus* claim and must be based on transparent experimental protocols that show what exactly researchers did to hold everything but the treatment constant over time. A p -value in a one-sample t -test informs us about the random sampling error, *irrespective* of whether our experimental procedure was successful in holding everything but the treatment constant over time or not. The only important assumption is that the treatment that leads to the observation of individual before-and-after differences presumably remains *unchanged* over all statistical replications. One should be clear that there is no role for a p -value when subjects in before-and-after-treatment comparisons are not randomly recruited.

⁴ Even though the p -value is the cornerstone of the statistical methodology that is currently in dominant use, it is but a summary statistic of the data at hand from which inductive inferences do not flow automatically. Elaborating on the per se limited inferential content of the p -value is not within this paper’s scope, however. For a discussion see for example, AMRHEIN et al. (2019), HIRSCHAUER et al. (2018, 2019), MCSHANE et al. (2017), TRAFIMOV et al. (2018), WASSERSTEIN and LAZAR (2016), WASSERSTEIN et al. (2016, 2019), or ZILIAK and McCloskey (2008).

Being a probabilistic concept based on a chance model (i.e. a hypothetical replication of a chance mechanism), p -values are not applicable if there is no random process of data generation (either randomization or random sampling). When there is no randomization, maintaining the p -value's probabilistic foundation therefore poses serious conceptual challenges when we already have the data of the whole target population (DENTON 1988: 166f.). An example is an experimental within-subject design where experimental subjects are clearly a non-random convenience sample, or where we do not want to generalize beyond the confines of the particular sample to start with. In such cases, the sample already constitutes the finite population to which we are limited. Due to the lack of a chance mechanism that could hypothetically be replicated, there is no role for the frequentist p -value and statistical significance testing. The fact that there is no room for statistical inference when we already have data of the entire inferential target population is formally reflected in the finite population correction factor. Rather than assuming that a sample was drawn from an infinite population – or at least that a small sample of size n was drawn from a very large population of size N – the finite population correction factor $(1-n/N)^{0.5}$ accounts for the fact that, besides absolute sample size, sampling error decreases when the sample size becomes large relative to the whole population. The correction reduces the standard error and is commonly used when sample share is more than 5% of the population (KNAUB 2008). Having the entire population corresponds to a correction factor of zero and thus a corrected standard error of zero.

If p -values are nonetheless calculated for entire populations (or non-random samples for that matter), one would have to imagine an infinite “unseen parent population” (or “superpopulation”), i.e. an underlying stochastic mechanism that is hypothesized to have generated the observations in the observed sample. DENTON (1988) critically notes that this rhetorical device, which is also known as “great urn of nature,” does not evoke wild enthusiasm from everybody. “However, some notion of an underlying [random] process – as distinct from merely a record of empirical observations – has to be accepted for the testing of hypotheses in econometrics to make any sense” (DENTON 1988: 167). We would add that researchers who resort to the p -value in such circumstances should explicitly explain why and how they base their inferential reasoning on the notion of a superpopulation. When doing so, they should be clear that this notion does not facilitate statistical inference in the conventional sense of generalizing towards a numerically larger parent population. Instead, inferences would be limited to the unseen superpopulation in terms of a random process that is supposed to “apply” to only and exclusively the subjects who happen to be in the sample.

4 Inferences in experiments without treatment comparisons

In experimental treatment comparisons, the term “control” means first of all generating *ceteris paribus* conditions (*ex-ante* control over confounders) with the objective of identifying causal treatment effects. We know that this *ex-ante* control comes in two forms: in randomized-treatment-group comparisons, control over confounders is achieved without exercising control over the environment; i.e. randomization, which balances confounders (including unknown ones) across treatment groups, replaces environmental control. In before-and-after-treatment comparisons, in contrast, control over confounders requires that we exercise control over the environment and fix and maintain all factors that could influence subjects' behaviors besides the treatment under investigation.

Often, economic experiments do not settle for identifying causal treatment effects among experimental subjects in more or less artificial experimental environments. Instead, experimenters want to learn what governs the behaviors of certain social groups in relevant real-world contexts and, eventually, how policy interventions would work in these contexts. This requires not only going beyond internal validity and causality. It also requires moving external validity beyond statistical inference, which is solely concerned with random error in repeated random sampling from the same population and thus the *sample-*

population relationship. That is, we cannot limit ourselves to the question of how we can generalize from the behavior of experimental subjects in a particular but potentially uninformative experiment to the would-be behavior of the parent population in this very experiment. Instead, we need to address the *experiment-real-world relationship*. Or using a well-known expression coined by SMITH (1982), we should exercise “control over subjects’ preferences” and search for experimental designs which ensure that subjects’ choices in the experiment reveal their “true” preferences. In the terminology of measurement theory we would say that, besides the uncertainty of the measurement due to sampling error (measurement precision/reliability; signal-to-noise ratio), we are now concerned with the accuracy of the measurement (measurement validity) and the question of whether the measurement instrument “experiment” yields a manifest variable (observed experimental behavior) that is informative regarding the latent variable of interest, i.e. people’s true preferences. It should be noted that an experiment’s measurement accuracy cannot be evaluated by statistical tools. It can only be evaluated based on the logical consistency and plausibility of the argument that is put forward in justification of the particular experimental design and/or in relation to a presumed standard of knowledge.

Control over subjects’ preferences is crucial for the external validity of economic experiments irrespective of whether they are based on treatment comparisons or not. However, this aspect of external validity is often more salient in economic experiments that study only one treatment and do not aim for causal inferences through *ceteris paribus* treatment comparisons. While still relying on an experimenter’s intervention, such experiments are focused on measuring latent preferences such as individual risk or social preferences. Prominent examples are experimental games such as prisoner’s dilemmas, trust games, or public goods games that are implemented to find out, for instance, whether the choices made by individuals are in line with conventional rational choice predictions.⁵ For example, one might deliberate how large the real payments (incentives) that are linked to subjects’ abstract earnings in a dictator “game” would have to be to achieve a valid measurement in that these incentives make subjects reveal their true prosocial preferences. Another example is the attempt to avoid “experimenter demand effects” that often threaten external validity because subjects are usually aware of participating in an experiment and often inclined to please experimenters (DE QUIDT et al. 2018). When assessing the quality of the experimental control over subjects’ preferences, one should be clear that this aspect of external validity has nothing to do with *p*-values. In other words, we may jointly have randomization *and* random sampling *and* control over subjects’ preferences in an experiment. However, we may also have an experiment without randomized treatment comparison and without random recruitment, but with an attempted control over subjects’ preferences. Imagine an incentivized dictator game carried out with a non-random convenience sample of students who happen to be in an experimenter’s classroom on a particular Friday. In this case, all inductive inferences – be they towards the experimental behavior of a broader population of students or other demographic groups, or towards the real-life behavior of the classroom students or broader populations – must be based on scientific arguments beyond *p*-values. It would therefore be a gross abuse to use the term “statistical significance” for a purported corroboration of such inferences.

Control over the environment, in terms of shaping, knowing, and describing all behaviorally relevant factors besides the treatment of interest, generally decreases from lab experiments to field experiments, irrespective of whether they are based on treatment comparisons or not. Any taxonomic proposal that takes account of the diminishing control over the environment from the lab to the field is open to debate – at least for non-randomized experiments. Attaching the label “experiment” to studies that rely on proper randomization to control for confounders is likely to cause little controversy even when they are carried out in the field where it is difficult to know, let alone fix all relevant factors besides the treatment.

⁵ Of course, all these games could also be used within a randomized design. Simply imagine an experiment in which subjects are randomly assigned to two dictator experiments with differing initial endowments.

In non-randomized designs, in contrast, the classification is likely to become controversial at some point; i.e. an arguable minimum level of control over the relevant environment would seem to be a prerequisite for calling a non-randomized approach an experiment. Irrespective of the label, we must take account of the specific research design when making inferences: (1) Causal inference must be based on scientific arguments but cannot be supported by a p -value when an experiment is not based on randomization. An important example are experimental within-subject designs. When causal inferences are based on doubtful claims of control over confounders, one should consider alternative experimental designs (e.g. randomized instead of non-randomized designs) or even a regression-based statistical control of observable confounders.⁶ (2) Inference dealing with the *sample-population relationship* (generalization) must be based on scientific reasoning but cannot be supported by a p -value when there was no random sampling from a broader (numerically larger) population. This is the case, for example, when randomized experiments are carried out with subjects from non-random convenience sample. (3) Inference dealing with the *experiment-real-world relationship* and thus the question of whether experimental subjects reveal their “true” preferences in a particular experiment cannot be supported by a p -value at all. When the control over subjects’ preferences is in question, one should avoid overhasty conclusions and check the robustness of results in replication studies with more valid experimental designs – preferably in field experiments carried out with subjects from the relevant parent population and a manipulation of subjects’ real-life environments.

5 Inferences in quasi-experiments

Often, non-randomized study designs focus on the behavioral outcomes induced by an intervention in one social group as opposed to another. Such designs are examples of “quasi-experiments” (CAMPBELL and STANLEY 1966) in which the *ceteris paribus* condition is in question. For illustration, imagine a dictator “game” in which a mixed-sex group of experimental subjects are used as first players who can decide which share of their initial endowment they give to a second player (one person acts as second player for the whole group). Additionally, assume that the experimental subjects are a convenience sample but not a random sample of a well-defined broader population. What kind of statistical inferences are possible? Neither one of the two chance mechanisms – random sampling or randomization – applies. Consequently, there is no role for the p -value: (i) Statistical inference towards a wider population beyond our experimental subjects is not possible because we are limited to a non-random sample. (ii) Statistical inference regarding causal relationships is not possible because there was no random assignment of subjects to treatments. Instead, *one* treatment was used to obtain a behavioral measurement in *two* predefined social groups. We should therefore simply describe, without reference to a p -value, the observed difference and the experimental conditions – or carry out a regression analysis to control for

⁶ There is no need to resort to regression when proper randomization ensures *ex ante* that confounders are statistically independent of treatments. In some cases, for instance when only a small experimental group is available (cf. footnote 2), switching to an *ex-post* control of confounders in a statistical model may be appropriate, however. It may therefore be useful to realize how, in the simplest case without confounders, a treatment-group comparison relates to a linear model where we regress the response to a binary treatment dummy and a constant. Generally speaking, the sampling distributions of estimated regression coefficients $\hat{\beta}_j$ that link predictors x_j to response y are the distributions of the point estimates derived from a hypothetically repeated random sampling of the response variable at the fixed values of the predictors (RAMSEY and SCHAFER 2013: 184). Using a dummy regression (and a p -value based on the *sampling distribution*) instead of comparing two group averages (and a p -value based on the *randomization distribution*) can therefore be questioned on the grounds that it implies switching to a chance model that is at odds with the actually applied chance mechanism. There are specific constellations (equal variance in both groups or, alternatively, heteroscedasticity-robust regression standard errors) that lead to identical standard errors. However, group comparison and dummy regression only coincide as long as the former is based on the sampling-based approximation of the standard error of the randomization distribution. If the group comparison were based on the “true” standard error of the randomization distribution, we would obtain a lower standard error compared to which the standard error in the regression would be upwardly biased (ATHEY and IMBEN 2017).

confounders if necessary; for example, the male subjects may be more or less wealthy than the female subjects which could be another explanation for the differences between the two groups.

Due to engrained disciplinary habits, researchers might be tempted to implement “statistical significance testing” routines in our dictator game example even though there is no chance model upon which to base statistical inference. While there is no random process, implementing a two-sample t -test might be the spontaneous reflex to find out whether there is a “statistically significant” difference between the two sexes. One should recognize, however, that doing so would require that some notion of a random mechanism is accepted. In our case, this would require imagining a randomization distribution that would have resulted if money amounts had been randomly assigned to sexes (“treatments”). Our question would be whether the money amounts transferred to the second player differed more between the sexes than what would be expected in the case of such a random assignment. We must realize, however, that there was no random assignment of subjects (with all their potentially confounding characteristics) to treatments, i.e. the sexes might not be independent of covariates. Therefore, the p -value based on a two-sample t -test for a difference in mean does not address the question of whether the difference in the average transferred money amount is *caused* by the subjects’ being male or female. That could be the case, but the difference could also be due to other reasons such as female subjects being less or more wealthy than male subjects. As stated above, it would therefore make sense to control for known confounders in a regression analysis *ex post* – again, without reference to a p -value as long as the experimental subjects have not been recruited through random sampling.

6 Conclusion

Systematizations of economic experiments have not predominantly addressed the inferences that can be made in different types of experimental designs. Usages of the term “experiment” range from a narrow view of “applying randomization” to identify causal effects, to a broad perspective of “trying something out” or measuring something. Our paper has shown that an adequate differentiation of experimental designs advances the understanding of what we can infer from different types of experimental studies. Several points should be kept in mind: *first*, a random process of data generation – either random assignment or random sampling – is required for frequentist tools such as p -values to make any sense, however little it may be. *Second*, the informational content of p -values are different in randomization-based inference as opposed to sampling-based inference. Randomization-based inference is concerned with internal validity and causality, whereas sampling-based inference is concerned with external validity in terms of generalizing from a sample to its parent population. *Third*, while being conceptually different, the sampling-based standard error used in a two-sample t -test can be used as an approximation in randomization-based inference. If one accepts the approximation, and if experimental subjects are recruited through random sampling, the resulting p -value can be used as an aid both for assessing internal validity and for generalizing to the parent population. However, if experimental subjects are not randomly recruited, statistical inferences must be limited to assess the causalities within the *given* study population. *Forth*, in the context of economic experiments, there are two essentially different meanings of the term “control” that must not be confused. In experiments aimed at identifying causal treatment effects, control means first of all ensuring *ceteris paribus* conditions (statistical independence of treatments). Besides that, the term “control” is concerned with external validity beyond the sample-population relationship. The expression “control over preferences” is used to indicate experimental designs in which a valid measurement is achieved in that experimental subjects can be believed to reveal their true real-world preferences. This design quality, which is crucial for making valid inferences, is part of scientific reasoning but cannot be aided by p -values.

Acknowledgment

The authors would like to thank the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) for funding.

References

- Athey, S., Imben, G.W. (2017): The Econometrics of Randomized Experiments. In: Banerjee, A.V., Duflo, E. (eds.): Handbook of Field Experiments. Volume 1. Amsterdam: Elsevier: 74–140.
- Amrhein, V., Greenland, S., McShane, B. (2019): Retire statistical significance. *Nature* 567: 305–307.
- Campbell, D.T., Stanley, J.C. (1966): *Experimental and Quasi-experimental Designs for Research*. Chicago: Rand McNally.
- Chamberlin, E.H. (1948): An experimental imperfect market. *Journal of Political Economy* 56(2): 95–108.
- Charness, G., Gneezy, U., Kuhn, M.A. (2012): Experimental methods: Between-subject and within-subject design. *Journal of Economic Behavior & Organization* 81(1): 1–8.
- Czibor, E., Jimenez-Gomez, D., List, J.A. (2019): The Dozen Things Experimental Economists Should Do (More of). NBER Working Paper No. 25451.
- Denton, F.T. (1988): The significance of significance: Rhetorical aspects of statistical hypothesis testing in economics. In: Klamer, A., McCloskey, D.N., Solow, R.M. (eds.): *The consequences of economic rhetoric*. Cambridge: Cambridge University Press: 163–193.
- de Quidt, J., Haushofer, J., Roth, C. (2018): Measuring and Bounding Experimenter Demand. *American Economic Review* 108(11): 3266–3302.
- Duflo, E., Glennerster, R., Kremer, M. (2007): Using randomization in development economics research: A toolkit. In: Schultz, T. Strauss, J. (eds.): *Handbook of Development Economics*, Volume 4. Amsterdam: Elsevier: 3895–3962.
- Dunning, T. (2012): *Natural Experiments in the Social Sciences: A Design-based Approach*. Cambridge: Cambridge University Press.
- Harrison, G.W., List, J.A. (2004): Field Experiments. *Journal of Economic Literature* 42(4): 1009–1055.
- Hirschauer, N., Grüner, S., Mußhoff, O., Becker, C. (2019): Twenty steps towards an adequate inferential interpretation of p-values in econometrics. *Journal of Economics and Statistics* 239(4): 703–721.
- Hirschauer, N., Grüner, S., Mußhoff, O., Becker, C. (2018): Pitfalls of significance testing and p-value variability: An econometrics perspective. *Statistics Surveys* 12: 136–172.
- Hoggatt, A.C. (1959): An experimental business game. *Behavioral Science* 4: 192–203.
- Knaub, J. (2008): Finite Population Correction (fcp) Factor. In: Lavrakas, P. (ed.): *Encyclopedia of Survey Research Methods*. Thousand Oaks: Sage Publications: 284–286.
- McShane, B., Gal, D., Gelman, A., Robert, C., Tackett, J.L. (2017): Abandon Statistical Significance. <http://www.stat.columbia.edu/~gelman/research/unpublished/abandon.pdf>
- Ramsey, F.L., Schafer, D.W. (2013): *The statistical sleuth: a course in the methods of data analysis*. Belmont: Brooks/Cole.
- Samuelson, P.A., Nordhaus, W.D. (1985): *Economics*. 12th ed. New York: McGraw-Hill.
- Sauermann, H., Selten, R. (1959): Ein Oligopolexperiment. *Zeitschrift für die gesamte Staatswissenschaft* 115: 427–471.
- Siegel, S., Fouraker, L.E. (1960): *Bargaining and group decision making*. New York: McGraw-Hill.
- Smith, V.L. (1962): An Experimental Study of Market Behavior. *Journal of Political Economy* 70:111–137.
- Smith, V.L. (1982): Microeconomic systems as an experimental science. *The American Economic Review* 72(5): 923–955.
- Trafimov et al. (2018): Manipulating the Alpha Level Cannot Cure Significance Testing. *Frontiers in Psychology*. <https://doi.org/10.3389/fpsyg.2018.00699>
- Vogt, W.P., Vogt, E.R., Gardner, D.C., Haeffele, L.M. (2014): *Selecting the right analyses for your data: quantitative, qualitative, and mixed methods*. New York: The Guilford Publishing.
- Wasserstein, R.L., Lazar N.A. (2016): The ASA’s statement on p-values: context, process, and purpose, *The American Statistician* 70(2): 129–133.
- Wasserstein, R.L., Schirm, A.L., Lazar, N.A. (2019): Moving to a World Beyond “ $p < 0.05$ ”. *The American Statistician* 73(sup1): 1–19.
- Ziliak, S.T., McCloskey, D.N. (2008): *The Cult of Statistical Significance. How the Standard Error Costs Us Jobs, Justice, and Lives*. Ann Arbor: The University of Michigan Press.

Please note:

You are most sincerely encouraged to participate in the open assessment of this discussion paper. You can do so by either recommending the paper or by posting your comments.

Please go to:

<http://www.economics-ejournal.org/economics/discussionpapers/2019-65>

The Editor