

The Cult of Statistical Significance – What Economists Should and Should Not Do to Make their Data Talk

By Walter Krämer*

Abstract

This article takes issue with a rather devastating critique of statistical significance testing propagated in a recent book by Ziliak/McCloskey (2008) of the same title. Ziliak/McCloskey argue that statistical significance testing is a barrier rather than a booster for empirical research in economics and should therefore be abandoned altogether. The present article argues that this is good advice in some research areas but not in others, with the aim of making practitioners aware of various fallacies connected with the concept of statistical significance, and at the same time showing where significance testing is most fruitfully employed. Taking all issues which have appeared so far of the German Economic Review and a recent epidemiological meta-analysis as examples, the present paper shows that there has indeed been a lot of misleading work where confirmatory significance testing has played a major role, and that at the same time many promising avenues, best summarized under the heading “specification tests”, have not been used.

Zusammenfassung

Der Beitrag kritisiert die in jüngster Zeit vermehrt geäußerte Kritik an statistischen Signifikanztests aller Art, speziell in den Wirtschaftswissenschaften. Hier wird in der Tat mit konfirmatorischen Tests, wo eine Ablehnung der Nullhypothese als Bestätigung einer vorab formulierten Theorie gewertet wird, viel Unfug betrieben. Dies wird an zahlreichen Beispielen aufgezeigt. Auf der anderen Seite existiert aber ein enormes, noch nicht ausgeschöpftes Potential an Spezifikationstests, die eine etablierte Theorie einer Prüfung durch die Daten unterwerfen. Hier können statistische Signifikanztests für den Fortschritt in der Forschung durchaus nützlich sein.

JEL Classifications: B40, C12, C52

Received: April 19, 2011

Accepted: June 14, 2011

* Research supported by DFG under SFB 823. On purpose, the title is the same as that of a recent book by Ziliak/McCloskey (2008). I am grateful to Michael Bücken and Ronja Walter for excellent research assistance, to Stephen Ziliak for comments on a book review in Statistical Papers which has provided the initial impetus for this article, and to two anonymous referees for helpful suggestions and comments.

1. Introduction and Summary

A significant statistical test only means: if the null hypothesis were true – a big if – then the tail probability of the observed event would be less than a pre-chosen level of significance. This rather modest claim is even further compromised by its extreme dependence on the size and on the generation of the sample and by the common practice of disguised multiple testing, i.e. doing lots of tests and reporting only the most “significant” results, and the ensuing understatement of the true probability of an error of the first kind (“data mining”). This is what every decent statistician knows, or at least should have been taught in any introductory mathematical statistics course.

Additional, rather popular empirical improprieties are HARKing (“Hypothesizing after the Results are Known”), collective as opposed to individual data mining, and what I call an Error of the Third Kind, by which I mean mistaking a rejected null as proof that the alternative is true. All of these deficiencies figure prominently in a critical literature of long standing that is summarized in section 2 below. But the critique which is the point of departure of the present paper is much more fundamental. It dates back to the very beginning of significance testing as a scientific discipline and concludes that even if significance testing were properly done according to the rules of the game, it would still be fundamentally flawed as an approach to empirical research in many fields due to implied disregard of what really counts in many applications, the size, as opposed to the mere existence, of an effect. Ziliak/McCloskey (2008) have created the neologism “oomph” for this; they argue “that ‘oomph’, the difference a treatment makes, dominates precision” (p. xvii), and that a rather disproportionate amount of attention is devoted to the latter, taking away scarce resources from more promising avenues of research.

There are therefore three types of misleading test-based inference: (a) there is no effect, but due to technical deficiencies, “significance” nevertheless obtains, (b) there is a large effect (much “oomph”), but due to variability, it is not “significant” and therefore discarded, and (c) there is only a small effect (no “oomph”), but due to precision it is highly “significant” and therefore taken seriously. Taken together, these sources of error have led McCloskey (2002, 44) to conclude: “The progress of economic science has been seriously damaged [by the common practice of significance testing]. You can’t believe anything that comes out of [it]. Not a word. It is all nonsense, which future generations of economists are going to have to do all over again. Most of what appears in the best journals of economics is unscientific rubbish. I find this unspeakably sad. All my friends, my dear, dear friends in economics, have been wasting their time.... They are vigorous, difficult, demanding activities, like hard chess problems. But they are worthless as science.” Or even more bluntly, in her book with Ziliak (2008): “If null-hypothesis significance testing is as idiotic as we and other critics have so long believed, how on earth has it survived?” (240).

One can hardly imagine a more devastating critique of this aspect of empirical work in economics (or in any other field).

This critique has of course been criticised itself, see for instance Hoover/Siegler (2008) or Spanos (2008). The purpose of the present article is to put the Ziliak-McCloskey view into yet another perspective, by supporting it for some types of tests but not for others. It is addressed at practitioners as both a warning of what can go wrong and as an advice as to what to do instead. The first class of tests, where much is going wrong indeed, is sometimes referred to as “presumptive” or “confirmatory” testing. It is from here that Ziliak/McCloskey (2008) and other critics draw most of their examples. Confirmatory testing means that there is a particular alternative one has in mind, with the aim or wish of establishing this as true. Section 2 summarizes various illegal ways in which this goal is often achieved in applications, plus related aberrations when interpreting confirmatory tests, and exemplifies such type (a) mistakes using a recent example from epidemiology. Section 3 considers type (b) and (c) mistakes. This is done by checking all empirical articles ever from the German Economic Review, the Journal of the Verein für Socialpolitik, which is distributed to about 4000 members four times a year. Both sections confirm Ziliak/McCloskey (2008) insofar as lots of useless and misleading inferences are unearthed. But they also show that confirmatory testing only makes sense, no matter whether one is after mere significance or “oomph”, if the underlying model is reasonably correct, and that it is the common failure to test for this which is the real threat to meaningful statistical results.

Such tests, often called “exploratory” or “specification” tests (see Krämer/Sonnberger, 1986), are the topic of the final section 4. Specification tests are not aimed at any specific alternative, so a rejection of the null only tells the investigator that he or she should look out for a better model, without establishing whichever type of “effect” there is supposed to exist. They are also more in line with the Popperian paradigm of scientific progress, where the null hypothesis corresponds to established beliefs, to be abandoned only in the presence of compelling evidence. In a sense, therefore, the Ziliak-McCloskey argument is turned on its head: in order to extract meaningful information (“oomph”) from the economy or whatever field of application via formal statistical models, one has to do a lot of significance testing first.

2. Confirmatory Testing and Errors of the Third Kind

A significance level of $\alpha = 5\%$ for a statistical test implies that, even when the null hypothesis were true, the procedure would still reject it in roughly 5 out of 100 applications. In the context of a specific alternative, usually some kind of “effect”, this means that even without any effect being present, the test

will nevertheless claim one in roughly 5 out of 100 trials. This is the well known error of the first kind.

A first objection to the routine use of statistical significance testing concerns the ease with which a significant test often leads to what I have termed above an error of the third kind: to assume that a significant test implies that the alternative is true. Take a recent German study of childhood leukemia in the vicinity of nuclear power plants – a meta analysis combining various previous investigations and some data collected independently – prepared by Greiser (2009) for the political party “Bündnis 90/Grüne”. Like many others, it purports to show that nuclear power plants induce a “statistically significantly elevated risk of leukemia for all age groups considered” (3)¹ and starts with an error of the third kind: mistaking a rejected null hypothesis as proof that the alternative is true. “AKW erhöhen das Leukämierisiko (nuclear power plants increase risk of leukemia)” was the heading of a press release distributed by Bündnis 90/Grüne in the fall of 2009.

This error of the third kind, or some variant such as “the null hypothesis is wrong with 95% probability” occurs even among professional statisticians (Haller/Krauss, 2002) or in statistics textbooks for students. Examples from the American market include Guilford (1942, and later editions), which was probably the most widely read textbook in the 1940s and 1950s, Miller/Buckhout (1973, statistical appendix by Brown, 523) or Nunally (1975, 194–196). Examples from the German market include Wyss (1991, 547) or Schuchard-Fischer et al. (1982), who on p. 83 of their best-selling textbook explicitly advise their readers that a rejection of the null at 5% implies a probability of 95% that the alternative is correct. For details, see Krämer/Gigerenzer (2005) or Krämer (2011, chapter 8).

Presumably, however, in the leukemia example, not even statistical significance obtains. A first type (a) mistake is HARKing (Kerr, 1998): “Hypothesizing after the results are known”. In Germany, testing for leukemia started only after an abnormal cluster of leukemia cases was found close to the Krümmel power generation plant. But as Kruskal (1969) puts it, “Almost any set of data [...] will show anomalies of some kind when examined carefully, even if the underlying probabilistic structure is wholly random – that is, even if the observations stem from random variables that are independent and identically distributed. By looking carefully enough at random data, one can generally find some anomaly [...] that gives statistical significance at customary levels although no real effect is present” (247).

A famous example is one of the very first applications of significance testing at all, the observation made by astronomers that the orbital planes of the planets

¹ English translation. The German original says: “Die Ergebnisse zeigen ein statistisch signifikant erhöhtes Erkrankungsrisiko an Leukämie für alle untersuchten Altersgruppen.”

are quite close together. In 1734, Daniel Bernoulli and his son John computed the probability that this is due to chance (given that orbital planes are determined randomly; in modern language, they computed the prob-value of a test). This probability however is only correct if the particular anomaly had been established beforehand, and is larger otherwise.

Then there is the publication bias, which is bound to particularly affect any meta-analyses which collect together previous work. “There is some evidence that in fields where statistical tests of significance are commonly used, research which yields nonsignificant results is not published” (Sterling, 1959, 30). “Such research being unknown to other investigators may be repeated independently until eventually by chance a significant result occurs.” Taken to the limit, this argument implies that a “significant” effect will be found eventually almost surely, no matter what.

A second, even more popular mistake is to claim some nominal significance level α when in reality the reported test statistic is the most significant one among several trials, each conducted at the level α . The true significance level is then simply the probability that the maximum of several test statistics is larger than some critical value; it increases rapidly with the number of tests that are performed. This multiple testing problem has of course long been recognized in statistical research (see Krämer and Sonnberger (1986, chapter 6) for an overview of the early literature in econometrics), but it seems that the enormous theoretical work that has been done here has not yet made it into routine empirical applications. And even if it had, the many restrictions that are attached to most multiple testing approaches would severely limit their impact on the problem we are discussing here.

In economics, this habit of reporting only the most “significant” results is sometimes referred to as “data mining” (Lovell, 1983).² It is of course strictly illegal and rightly frowned upon, but has nevertheless been common practice in empirical economics ever since statistical tests of significance have been introduced.

In the leukemia-example, the alleged significance appears to be mostly due to data mining and publication bias; it vanishes completely – in fact, the sign of the observed effect is reversed –, once an obvious failure of the underlying model, the total disregard of important confounding factors, is accounted for. According to Ries et. al (1999, figure 6 and table 1.5), and confirmed by many others, important risk factors for childhood leukemia are race and sex. For instance, childhood cancer incidence in the U.S. is 30% higher for boys as compared to girls and almost double for whites as compared to blacks. For leukemia only, the highest incidence rates are observed among hispanics (48.5 per

² Not to be confused with the serious business of the same name that is a subject of modern computer science.

million as compared to 41.6 per million for whites and 25.8 per million for blacks). By far the lowest rates for any type of childhood cancer are observed for American Indians.

Also, leukemia incidence correlates strongly with income – the higher the income of the parents, the larger the risk of leukemia for kids. In Scotland, for instance, the incidence of childhood leukemia between the richest and the poorest subpopulations differs by as much as 50%. Other risk factors which have been identified so far are population density and population mixing, which both might likewise lead to an increased exposure of susceptible individuals to infections and local epidemics which in turn could later promote the onset of cancers of many types.

Now, the plant that contributes most to the surplus of 158 leukemia cases in the Greiser study is San Onofre Nuclear Generating Station in Southern California, in the northwestern corner of San Diego County, south of the city of San Clemente. According to Greiser (2009, 21, table 4) there were 281 cases of childhood leukemia close to San Onofre (which in this case means: in San Diego County) in the 2001–2006 time period, compared to only 177 expected cases, an excess of 104. Therefore, this single data point contributes almost all of the 158 excess cases on which the “significant” increase of childhood leukemia in the vicinity of nuclear power plants is based.

Looking closer at the San Onofre site, however, it appears that virtually all confounding factors which have so far been established in the literature are higher there than elsewhere in the U.S.. For instance, San Diego County is rather wealthy, with average household income 20% above the national average. In addition, San Diego county has an above-average population of Hispanics and very few blacks. Also, both population density and population mixing are more pronounced in San Diego county than elsewhere in the U.S.. San Diego is the largest concentration of naval facilities in the world, with a constant moving in and out of families, which is even further accentuated by a large University and many more military facilities such as training camps, air-bases, Marine Corps Recruit Depots and coast guard stations. All of these variables correlate strongly with childhood leukemia.

Removing San Onofre from the Greiser (2009) data set, and adding some studies he has overlooked, the initial surplus of leukemia cases turns into a deficit, see Krämer/Arminger (2010). This section therefore shows that one complaint against significance testing raised by Ziliak/McCloskey (2008) – spurious significance due to bad practice – is certainly warranted by current practice in many fields. But it also shows that any claims as to significance of any sort require that the underlying model be reasonably correct.

3. Eleven Years of Significance Testing in the German Economic Review

This section turns to type (b) and (c) mistakes, i.e. neglecting large effects which are not “significant” and celebrating trivial effects which are significant only due to sample size. To that purpose, I scrutinized all issues which have appeared so far of the German Economic Review, the official Journal of the Verein für Socialpolitik, an association of about 4000 German speaking economists from all over the world. It was inaugurated in 2000 as the English language successor to the venerable Zeitschrift für Wirtschafts- und Sozialwissenschaften, also known as “Schmollers Jahrbuch”, with a history dating back to 1871. At the time of this writing it is in its 12th year of existence, so there are 11 complete volumes which in this section will be scrutinized for misleading applications of statistical tests of significance.

Table 1 provides a summary. It shows that about 40% of all articles published so far rely for their results both on data and on formal tests of significance of the confirmatory type (a vast majority being *t*-tests of the significance of some effect). This percentage has been increasing recently, but is still somewhat less than Ziliak/McCloskey (2008, chapters 6 and 7) find for the American Economic Review 1980–1999.

The rather astonishing number of more than 10.000 tests of significance, i.e. about 1000 tests per volume, is of course due to the routine production of such tests by commercial software packages that are used by the authors to fit their models. And more often than not, they result in comments of the type “*X* has a significantly positive impact on *Y*”, without any reference to its magnitude. In 486 cases, the exact magnitudes of the estimated coefficients are not even reported, the only information given being that they were “significant”.

Even this attribute is doubtful given the prevalence of comments like “table 2 presents the results of our final model estimation”. Obviously, this means that various estimates and tests were computed beforehand, with only the most “significant” results remaining to be shown, so the scientific value of such tests is close to zero: “Cheap *t*-tests, becoming steadily cheaper with falling computational costs, have in equilibrium a marginal scientific product equal to their cost” (Ziliak/McCloskey, 2008, 112). And when the costs of tests tend to zero, their informational value seems to follow straight in line.

This is true even if the tests as such were properly done. Table 2 confirms that not even this is true in many cases. It provides information on various additional features of the 110 papers which report tests of significance.

Among papers that rely on some form of regression, most choose a linear functional form without much discussion. Some also exclude or include variables solely on the basis of statistical significance, paying little attention to pertinent economic theory. And when only final versions of regression models

were presented, only very rarely was there awareness of the multiple testing problem. Table 3 gives the details.

Table 1

Confirmatory significance tests in the German Economic Review

Volume	Number of articles	Articles with confirmatory significance tests	Number of confirmatory significance tests	Only sign, no effect
1	21	8	421	4
2	24	9	527	0
3	21	7	725	332
4	22	5	176	22
5	20	10	994	40
6	25	10	1359	87
7	22	9	653	0
8	24	11	1375	0
9	24	11	1171	0
10	28	12	1809	0
11	27	18	1365	1
Together	258	110	10575	486

Table 2

Three types of mistakes

Type (c) error: Confusion of economic and statistical significance of estimated coefficients or effects ("significant" used for both? Much ado about statistically significant but economically small coefficients or effects?)	62 / 110 = 56.4%
Type (b) error: Economically significant and plausible effects or coefficients discarded due to lack of statistical significance?	31 / 110 = 28.2%
No or only passing discussion of the dependence of "significance" on the correct specification of the model? (Independent and/or identically distributed observations etc.)	78 / 110 = 70.1%

Table 3

Selected deficiencies of papers that use some sort of regression model

No detailed discussion of the appropriateness of the chosen model (no theoretical justification for the particular functional form, no or only cursory diagnostic testing etc)	56 / 98 = 57.1%
Explanatory variables included or excluded exclusively or almost exclusively on the basis of statistical significance?	20 / 98 = 20.4%
Several models tried, only the final one presented, but no awareness of the multiple testing problem	14 / 98 = 14.2%

The critique summarized in tables 1 to 3 is not meant to denigrate a particular journal or empirical work in the German Economic Review as such. In fact, the particular approach which is criticised here appears to be common to most economic journals in the world and is even more prevalent, if Ziliak/McCloskey (2008) are to be believed, in the American Economic Review, which is the leading journal in the field. Also, there are many fine empirical papers in the German Economic Review which, to convey their message, do not rely on confirmatory significance tests at all.³ And even if tests are reported, it is sometimes with some sort of tongue in cheek, as in Wagner (2010), to pay respect to some tradition which not even the authors do take seriously any more. And it is exactly this what the present paper wants to emphasize: that the endless tables of t -values which adorn most empirical papers nowadays are indeed what Ziliak/McCloskey call them – a needless waste of space and time.

4. Specification Testing vs. Searching for Effects

An important point often overlooked in the significance debate is that any such claim – no matter whether it only concerns the existence or also the size of an effect – is only valid if the underlying statistical model describes the data reasonably well. In particular, as was shown in section 3, one has to ensure that all relevant explanatory variables have been properly accounted for. And it is here that statistical test of significance can help researchers along their way a lot. This is best exemplified with the help of the standard linear regression model

$$y_i = \beta_0 + \beta_1 x_{i1} + \dots + \beta_K x_{iK} + u_i (i = 1, \dots, n),$$

where the y 's are to be explained, the x_{ik} are observations on K explanatory variables (regressors, exogenous variables, design variables), and the u_i 's are unobservable disturbance terms, presumably uncorrelated, with equal variance and expectation zero. Confirmatory testing in this context means establishing the “significance” of individual regressors, i.e. testing $H_0 : \beta_k = 0$ for some $k = 1, \dots, K$. As was shown in section 4, and is confirmed by independent investigations by Ziliak/McCloskey (2008), well above 99% of all statistical tests reported in a typical economics journal are of this type, with all the ensuing complications discussed in sections 2 and 3.

What is much more rarely done, but should be standard practise, is testing whether the model that is entertained provides a proper approximation to the

³ Recent examples are Bachmann/Burda (2010), who convincingly summarise and explain labour market dynamics in Germany, using lots of tables and figures, but no t -tests whatsoever, or, to take another journal, Becker/Übelmesser (2010), who provide a complete t -test-free empirical analysis of the effects of advertising in the German health care business.

data in the first place. Only in that case do tests of the confirmatory sort apply. And as shown in Krämer et al. (1985), most empirical papers, even in decent journals, fail such specification tests, often by wide margins. Among things that can go wrong here are omitted regressors (see section 3), non-linearity, in particular interaction effects, measurement errors, endogeneity or structural changes in the β -coefficients. Only if such deficiencies can be ruled out with some confidence does it make sense to talk about “oomph”, i.e. the size of the β s, (and, if one so chooses, to test for their significance). Or to put this differently, one should first test whether the assumptions concerning the (conditional) first moments of the y 's are indeed correct before proceeding to establish any kind of effect. So if one takes seriously what most critics of standard statistical significance testing maintain, that it is the size and not the significance of effects which really counts, then one has to do some significance testing first.

As compared to first moments, the correct specification of second moments or the distributional form of the dependent variable appears to be less important: Whether or not the disturbances in a regression model are normal, homoscedastic or uncorrelated among themselves will certainly affect the efficiency or the size of standard estimators and tests, but not the possible estimation of “effects”. Also, deviations from assumptions concerning second moments are quite easily accounted for by robust procedures such as heteroskedasticity and autocorrelation resistant t-tests and confidence intervals.

Krämer/Sonnberger (1986) provide additional arguments in favour of focusing on first moments and give an overview of the early literature of statistical specification testing. An extremely simple procedure known as the RESET (Regression Specification Error Test) for instance only involves adding artificial regressors like squares, cubes or cross products of the initial regressors and testing whether they are significant. If so, there is evidence that the initial linear functional form is not correct. Or, for time series data, one could simply compare parameter estimates obtained from the initial model to estimates obtained from the same data after first differencing. If the model were correct, both estimators estimate the same things and should be close to each other. If not, there is evidence again that something is wrong with the model (an idea which has been generalized by Hausman (1978) to various other pairs of estimators). Krämer/Sonnberger (1986) collect together a generous toolbox of such techniques for checking model adequacy.

A related class of specification tests do not challenge a given model, because the underlying model is in most cases rather obvious and simple, but test whether or not certain parameters in this model are compatible with established economic theory. An example is the test for weekday anomalies for stock returns, see Krämer/Runde (1992). Financial theory requires that expected excess returns are positive and equal to each other for all days of the week, or that successive returns have autocorrelation zero. Again, one is not interested here in the size of the effect, but rather in whether one exists in the

first place, the only problem being the distinction between statistical and practical significance (small deviations from theory cannot be exploited due to trading costs). One might call such procedures “theory-attacking-tests” (as opposed to “theory-confirming-tests”, which are the prime target of the Ziliak-McCloskey critique.)

Unfortunately, specification tests and theory attacking tests are a distinct minority among statistical significance tests reported in economics journals. Table 4 gives the respective figures for the German Economic Review.

Table 4

**Papers with exploratory or “specification” test
(i.e. test where an acceptance of H_0 is fine)**

Number of papers where such tests are done at all	26
Number of papers which discuss the power of such tests	4/26 = 15.4%

Examples of papers from recent volumes of the German Economic Review that rely at least partially on specification tests are Zarzoso et al. (2009, 327): “Specification tests also rejected the inclusion of a quadratic aid-term in the estimated equation”) or Feld/Reulier (2009), who investigate the effect on a Swiss canton’s personal income tax rate of various regressors, including the corresponding rates of neighbouring cantons. In addition to lots of confirmatory testing, they also test whether their regressors are truly exogenous: “Equation (1) cannot consistently be estimated by OLS because there is an obvious endogeneity problem. Hausman tests indicate that the neighbouring tax rates at the local level or at the regional levels are endogenous” (98). More recently, Feld/Schneider (2010, 130) test for the adequacy of an indicator model of the shadow economy. However, tests of these types are still dwarfed in number by mindless batteries of t -tests attached to parameter estimates. In the German Economic Review, the terms “specification test”, “specification testing”, RESET or “Hausman test” appear less than ten times each in eleven years, that is less than once a year.

One possible reason is that specification tests, other than confirmatory t -tests, are not routinely supplied by most statistical software packages. In fact, when Krämer et al. (1985) first applied a battery of specification tests to selected empirical papers from the literature, they had to program most procedures themselves. Although this is not difficult at all in many cases, this might still deter practitioners not familiar with the software which they use.

Another reason why this toolbox is not as popular as it deserves is the ease with which econometric models can be shown to be incorrect. For instance, there was not a single empirical paper among the ones surveyed by Krämer

et al. (1985) which survived all specification tests that were applied. And as there is no such thing as a truly “correct” model in all of econometrics (all econometric models are only approximations to reality), practitioners might be afraid to admit failures of this type, the more so as rejections occur ever more often given the large samples which are nowadays common in empirical work in economics and elsewhere, which allow even miniscule deviations form a model to be detected by a test. This calls for a distinction between “incorrectness” in a statistical and in an economic sense (when is a model “reasonably correct”) , similar to the distinction between economic and statistical significance in the context of confirmatory testing or theory attacking tests, which however is an issue far beyond the scope of the present paper.

5. Conclusion

The admonition often heard recently to stop testing in empirical economics is partially mistaken. While it is true that confirmatory testing, where the null hypothesis is only entertained as a dummy to help establishing a prearranged alternative, has a huge potential to mislead, and does indeed mislead in many applications, specification testing is more important than ever. Therefore, the advice should be, not to abandon the concept of significance, but to shift the focus to other types of null hypotheses.

References

- Bachmann, R./Burda, M.* (2010): Sectoral transformation, turbulence and labor market dynamics in Germany, *German Economic Review* 11, 37–59.
- Becker, B./Uebelmesser, S.* (2010): Health insurance competition in Germany – the role of advertising, *Schmollers Jahrbuch* 130, 23–39.
- Bündnis 90/Grüne* (2009): AKWs erhöhen das Leukämierisiko, Press Release, Sep. 7.
- Feld, L. P./Reulier, E.* (2009): Strategic tax competition in Switzerland: Evidence from a panel of the Swiss cantons, *German Economic Review* 10, 91–114.
- Feld, L. P./Schneider, F.* (2010): Survey on the shadow economy and undeclared earnings in OECD countries, *German Economic Review* 11, 109–149.
- Greiser, E.* (2009): Leukämie-Erkrankungen bei Kindern und Jugendlichen in der Umgebung von Kernkraftwerken in fünf Ländern, Report prepared by the political party Bündnis 90/Grüne, see http://www.gruene-bundestag.de/cms/archiv/dokbin/302/302113.studie_leukaemierisiko.pdf.
- Guilford, J. P.* (1942): *Fundamental Statistics in Psychology and Education*, McGraw-Hill, New York.
- Haller, H./Kraus, S.* (2002): Misinterpretation of significance: A problem students share with their teachers? *Methods of Psychological Research Online* 7, 1–20.

- Hausman, J. A. (1978), Specification Tests in Econometrics, *Econometrica* 46, 1251–1271.
- Hoover, K. D./Siegler, M. V. (2008): The rhetoric of ‘signifying nothing’: a rejoinder to Ziliak and McCloskey, *Journal of Economic Methodology* 15, 57–68.
- Kerr, N. L. (1998): HARKing (Hypothesizing After the Results are Known), *Personality and Social Psychology Review* 2, 196–217.
- Krämer, W. (2011): *Denkste – Trugschlüsse aus der Welt des Zufalls und der Zahlen*, 9th paperback edition, München.
- Krämer, W./Armingier, G. (2010): True believers. Or: numerical terrorism at the nuclear power plant, *Jahrbücher für Nationalökonomie und Statistik*, forthcoming.
- Krämer, W./Gigerenzer, G. (2005): How to confuse with statistics. The use and misuse of conditional probabilities, *Statistical Science* 20, 223–230.
- Krämer, W./Runde, R. (1992): The holiday effect: yet another capital market anomaly? in: S. Schach/G. Trenkler (eds.): *Data analysis and statistical inference: Festschrift in honour of Friedhelm Eicker*, Bergisch-Gladbach, 453–462.
- Krämer, W./Sonnberger, H. (1986): *The Linear Regression Model under Test*, Heidelberg.
- Krämer, W./Sonnberger, H./Maurer, J./Havlik, P. (1985): Diagnostic checking in practice, *Review of Economics and Statistics* 68, 118–123.
- Kruskal, W. (1968): Tests of statistical significance, in: D. Sills et al. (ed.): *International Encyclopedia of the Social Sciences*, New York, 238–250.
- Lovell, M. C. (1983): Data Mining, *Review of Economics and Statistics* 65, 1–12.
- McCloskey, D. (1985): The loss function has been mislaid: The rhetoric of significance tests, *American Economic Review* 75, 201–205.
- McCloskey, D. (2002): *The Secret Sins of Economics*, New York.
- Miller, G. A./Buckhout, R. (1973): *Psychology: The Science of Mental Life*, 2nd ed., New York.
- Nunnally, J. C. (1975): *Introduction to Statistics for Psychology and Education*, New York.
- Ries, L. A. G./Smith, M. A./Gurney, J. G./Linnet, M./Tamra, T./Young, J. L./Bunin, G. R. (1999): *Cancer Incidence and Survival among Children and Adolescents: United States SEER Program 1975–1995*, National Cancer Institute, Bethesda, MD.
- Schuchard-Fischer, C./Backhaus, K./Hummel, H./Lohrberg, W./Plinke, W./Schreiner, W. (1982): *Multivariate Analysemethoden – Eine anwendungsorientierte Einführung*, 2nd edition, Berlin.
- Spanos, A. (2008): Review of S. T. Ziliak and D. N. McCloskey’s *The Cult of Statistical Significance*, *Erasmus Journal for Philosophy and Economics*, 1 (1): 154–164.
- Sterling, T. R. (1959): Publication decisions and their possible effects on inferences drawn from tests of significance – or vice versa, *Journal of the American Statistical Association* 54, 30–34.

- Wagner, J.* (2010): Entry, exit and productivity: empirical results for German manufacturing, *German Economic Review* 11, 78–85.
- Wyss, W.* (1991): *Marktforschung von A–Z*. Luzern.
- Zarzoso, I./Nowak-Lehmann, F./Klasen, S./Larch, M.* (2009): Does German development aid promote German exports? *German Economic Review* 10, 317–338.
- Ziliak, S./McCloskey, D.* (2008): *The Cult of Statistical Significance: How the Standard Error Costs Us Jobs, Justice and Lives*, University of Michigan.