Statistical mistakes and misunderstandings: More common than you might think

David Schlesinger, Ph.D.

Lars Leksell Gamma Knife Center
University of Virginia

Department of Radiation Oncology

Conflicts of Interest

None

The "truth" is trendy at the moment….

Adapted from C. Lozada, Washington Post, 7/3/2018

The truth is still hard to come by

One study's result is not necessarily the truth

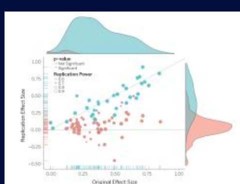

Original study effect size versus replication effect size (correlation coefficients). Diagonal line represents replication effect size equal to original effect size. Dotted line represents replication effect size of 0. Points below the dotted line were effects in the opposite direction of the original. Density plots are separated by significant (blue) and nonsignificant (red) effects.

Open Science Collaboration, Science, 49, 2015.

RESEARCH ARTICLE

PSYCHOLOGY

Estimating the reproducibility of psychological science

Open Science Collaboration*†

Attempted replication of 100 studies published in 3 psychology journals

97% of original but only 36% of replicated studies had statistically significant results

Only 47% of original effect sizes were within the 95% confidence interval of replication effect size

....including the study about non-reproducibility

On the Reproducibility of Psychological Science

Valen E. Johnson, **Richard D. Payne**, **Tianying Wang**, **Alex Asher**, and **Soutrik Mandal**
Department of Statistics, Texas A&M University, College Station, TX

Reanalyzes the previous paper's data using a model that accounts for publication bias toward significant p-values and estimates distribution of effect sizes.

Model suggests that 90% of the studies tested negligible effects
Publication bias towards p<=0.05 is a main cause of lack of reproducibility.

V.E.Johnson, et al., JASA 112, 2017

The number of retractions is sharply rising

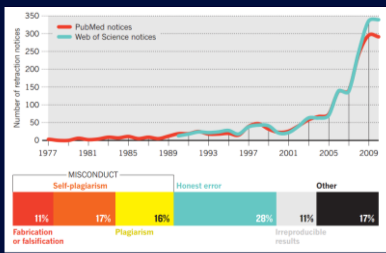

nature.com

Number of retraction notices has increased by factor of 10

Literature has only increased by 44%

30-40% of retractions are for technical reasons

R. Van Noorden, Nature, 478, 2011.

A lack of statistical fluency may be part of the problem

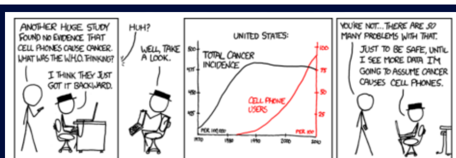

Many medical physicists receive little training in practical statistics as applied to clinical outcomes studies.

However…these studies are at the heart of our profession.

How to recognize when the statistics don't quite add up?

https://xkcd.com/925/

Session educational objectives

1. Learn about the presence of statistical problems in published studies
2. Identify common signs and symptoms of potential problems in various types of statistical tests
3. Learn methods for correctly implementing statistical analyses of the type commonly found in medical physics publications and in routine clinical activities

How we will spend the morning….

David Schlesinger University of Virginia	Study irreproducibility, philosophy of inferential statistics, how to use and misuse p-values
Jenghwa Chang Hofstra/Northwell Health	Normal and non-normal distributions: Why understanding distributions are important when designing experiments and analyzing data
William Sensakovic Florida Hospital/UCF/FSU/Adventist	Linear and logistic regressions, what they try to explain and how to interpret the results
Mike Altman Washington University	Statistical significance, statistical power, and clinical significance. How to explain your results in context.

A little quiz (not a SAMS question!)

You have a treatment you expect might alter performance on certain task:

Test with control and experimental groups of 20 subjects each
Compare the means of each group's performance
Result is significant per independent means t-test
($t=2.7$, $df=18$, $p=0.01$)

G. Gigerenzer et al., pub in The Sage handbook of quantitative methodology for social sciences, 2004.

Which statements are true? (several or none may be correct)

Remember: (t=2.7, df=18, p=0.01)

1) You have absolutely disproved the null hypothesis
2) You have found the probability of the null hypothesis being true
3) You have absolutely proved your experimental hypothesis
4) You can deduce the probability of the experimental hypothesis being true
5) If you decide to reject the null hypothesis, you know the probability that you are making the wrong decision
6) You have a reliable experimental finding – if you repeat the experiment a large number of times you would obtain a significant result 99% of the time.

G. Gigerenzer et al., pub in The Sage handbook of quantitative methodology for social sciences, 2004.

Cellphone use causes cancer

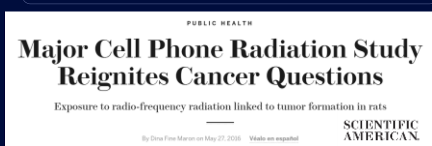
PUBLIC HEALTH

Major Cell Phone Radiation Study Reignites Cancer Questions

Exposure to radio-frequency radiation linked to tumor formation in rats

By Dina Fine Maron on May 27, 2016 Véalo en español

SCIENTIFIC AMERICAN

- Led by the National Toxicology Program (NTP) under the NIH
- Rodents exposed to calibrated RF (GSM and CDMA) radiation for 9 hours/day over 2 years
- Division into groups by SAR exposure
- Association between exposure and cardiac schwannoma in male rodents (no association in female rodents)

http://www.cancer.gov/about-cancer/causes-prevention/risk/radiation/cell-phones-fact-sheet

Cellphone use causes cancer – maybe?

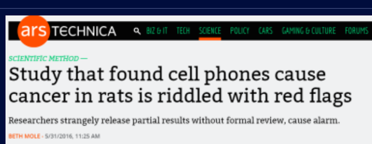
ars TECHNICA BIZ & IT TECH SCIENCE POLICY CARS GAMING & CULTURE FORUMS

SCIENTIFIC METHOD —

Study that found cell phones cause cancer in rats is riddled with red flags

Researchers strangely release partial results without formal review, cause alarm.

BETH MOLE - 5/31/2016, 11:25 AM

- Study was released before complete peer-review on a pre-publication website
- Control rats showed less than expected natural rate of tumor incidence and died early
- Incidence of tumor development correlates with age, so the early control death may have magnified the statistical findings

http://arstechnica.com/science/2016/05/study-that-found-cell-phones-cause-cancer-in-rats-is-riddled-with-red-flags/

Human studies are mostly one-sided

Publication Year	Study	Type	# participants	Outcome
2010	Interphone Study Group	Case-control study	~5000 cases; ~5000 matched controls; 13 countries	No overall risk*
2001 (updated 2007, 2011)	Danish cohort study	Cohort study	358,000	No association
2013 (updated 2014)	Million Women Study	Prospetive cohort study	791,710	Yes (acoustic neuroma), then no association
2014	CERENAT	Multicenter case control	447 cases, 892 matched controls	No association with regular use; yes association with heaviest use
2011	Swedish pooled analysis	Pooled analysis of 2 case control studies	1251 cases, 2438 controls	Increased risk of glioma

Cellphone use causes cancer?

We still don't know!

What question are we trying to ask?

Does cellphone use cause cancer?

Or more specifically since we are scientists running a study:

What is the probability that cellphone use causes cancer given the data from our study?

Two ways to consider the problem

Bayesian methods	Frequentist methods - Null-hypothesis significance testing (NHST)
Directly calculates conditional probability of a hypothesis	Determines how extreme the observed data is
Requires an estimate of the prior probability of a hypothesis	Never gives the probability of a hypothesis
Data can be used as it comes in	Does not require an estimate of prior probability
Assumes data is fixed and hypotheses vary	Requires the exact specification of the experiment in advance
	Assumes hypothesis is fixed and data varies

http://cfss.uchicago.edu/persp017_inference.html

Bayes' Theorem

$$P(H|D) = \frac{P(D|H) \cdot P(H)}{P(D)}$$

Bayes' Theorem

$$P(hypothesis|data) = \frac{P(data|hypothesis) \cdot P(hypothesis)}{P(data)}$$

Bayes' Theorem

Probability of our hypothesis given the data (posterior probability)

Probability we would see our data if our hypothesis is correct (likelihood)

Our prior estimate of the probability (prior probability)

$$P(hypothesis|data) = \frac{P(data|hypothesis) \cdot P(hypothesis)}{P(data)}$$

Looks a lot like the question we want to answer!

Normalization term (marginal probability)

Collect data that you then use to update your previous belief

When there are many possible hypotheses

$$P(hypothesis|data) = \frac{P(data|hypothesis) \cdot P(hypothesis)}{P(data|hypothesis) + \sum P(data|alt\ hypotheses)}$$

Denominator is the total probability of the data (mostly treated as a scaling constant)

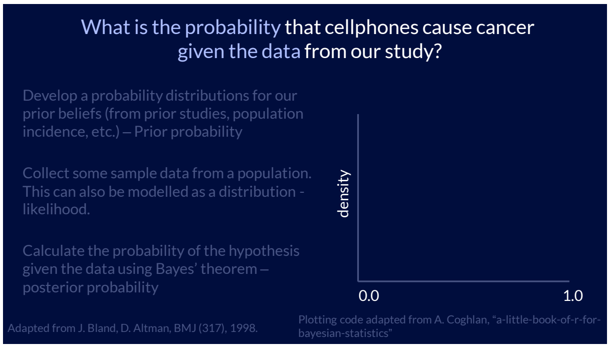
What is the probability that cellphones cause cancer given the data from our study?
Develop a probability distributions for our prior beliefs (from prior studies, population incidence, etc.) – Prior probability
Collect some sample data from a population. This can also be modelled as a distribution - likelihood.
Calculate the probability of the hypothesis given the data using Bayes' theorem – posterior probability
density
0.0
1.0
Adapted from J. Bland, D. Altman, BMJ (317), 1998.
Plotting code adapted from A. Coghlan, "a-little-book-of-r-for-bayesian-statistics"

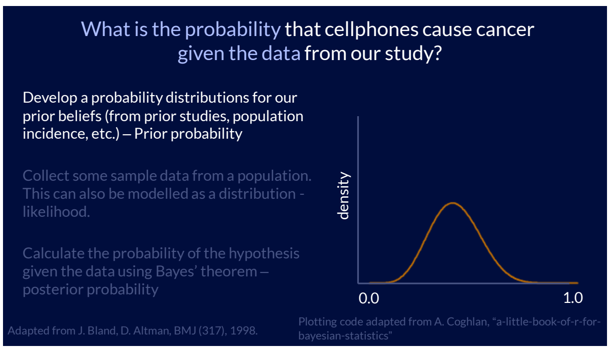
What is the probability that cellphones cause cancer given the data from our study?
Develop a probability distributions for our prior beliefs (from prior studies, population incidence, etc.) – Prior probability
Collect some sample data from a population. This can also be modelled as a distribution - likelihood.
Calculate the probability of the hypothesis given the data using Bayes' theorem – posterior probability
density
0.0
1.0
Adapted from J. Bland, D. Altman, BMJ (317), 1998.
Plotting code adapted from A. Coghlan, "a-little-book-of-r-for-bayesian-statistics"

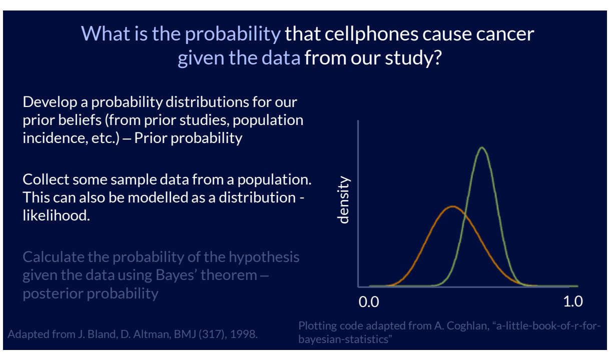
What is the probability that cellphones cause cancer given the data from our study?
Develop a probability distributions for our prior beliefs (from prior studies, population incidence, etc.) – Prior probability
Collect some sample data from a population. This can also be modelled as a distribution - likelihood.
Calculate the probability of the hypothesis given the data using Bayes' theorem – posterior probability
density
0.0
1.0
Adapted from J. Bland, D. Altman, BMJ (317), 1998.
Plotting code adapted from A. Coghlan, "a-little-book-of-r-for-bayesian-statistics"

What is the probability that cellphones cause cancer given the data from our study?

Develop a probability distributions for our prior beliefs (from prior studies, population incidence, etc.) – Prior probability

Collect some sample data from a population. This can also be modelled as a distribution - likelihood.

Calculate the probability of the hypothesis given the data using Bayes' theorem – posterior probability

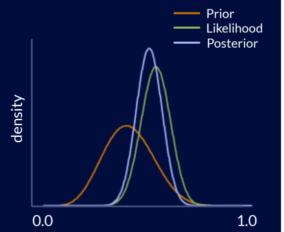

Adapted from J. Bland, D. Altman, BMJ (317), 1998.

Plotting code adapted from A. Coghlan, "a-little-book-of-r-for-bayesian-statistics"

Another (implicit) Bayesian example – linac QA

TG-142 monthly QA-mechanical

Table II. Monthly.

Procedure	Machine-type tolerance Non-IMRT	IMRT	SRS/SBRT
Mechanical			
Light/radiation field coincidence[b]		2 mm or 1% on a side	
Light/radiation field coincidence[b] (asymmetric)		1 mm or 1% on a side	
Distance check device for lasers compared with front pointer		1 mm	
Gantry/collimator angle indicators (@ cardinal angles) (digital only)		1.0°	
Accessory trays (i.e., port film graticle tray)		2 mm	
Jaw position indicators (symmetric)[c]		2 mm	
Jaw position indicators (asymmetric)[d]		1 mm	
Cross-hair centering (walkout)		1 mm	
Treatment couch position indicators[e]	2 mm/1°	2 mm/1°	1 mm/0.5°
Wedge placement accuracy		2 mm	
Compensator placement accuracy[f]		1 mm	
Latching of wedges, blocking tray[g]		Functional	
Localizing lasers	±2 mm	±1 mm	< ±1 mm

Excerpted from E. Klein, et al., Med Phys 36(9), 2009.

QA tests are a way to update prior probability of machine function to be beyond a decision threshold.

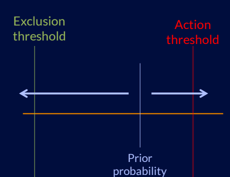

P. Chang, AJR Am J Roentgenol (152), 1989

But...we mostly don't do this....

Bayes is sometimes impractical

Prior can be subjective or unknown

$$P(hypothesis|data) = \frac{P(data|hypothesis) \cdot P(hypothesis)}{P(data)}$$

Real-world problems involve distributions. Often must be solved numerically. (Markov-chain Monte-Carlo)

Sometimes requires generating all of the possible probability combinations (although we generally can avoid this)

Instead, here is what we do…

We get rid of what we don't know…

$$P(hypothesis|data) \cancel{=} \frac{P(data|hypothesis) \cdot \cancel{P(hypothesis)}}{\cancel{P(data)}}$$

In real life, we often don't have data that can help us describe a prior probability. So, Bayes' theorem won't work directly…..

…and we turn the problem around

All we are left with is → $P(data|hypothesis)$

We create a reference hypothesis called the Null Hypothesis (H_0) → $P(data|H_0)$

We calculate how likely our data is assuming that this reference hypothesis is true

Philosophically, we measure how unexpected our experimental data is.
The more unexpected, the less likely the null hypothesis is true.

The null hypothesis (H_0)

Serves as a reference hypothesis

Is usually the opposite of the hypothesis you hope to be true

Is frequently stated as “no difference”, but doesn’t have to be

Any statistically significant result assumes the null hypothesis is true

A. Field, et al., Discovering Statistics Using R, 2012

What is the probability I would get data this (or more) extreme, assuming cellphones do not cause cancer?

1. Identify null and experimental hypotheses.
2. Determine the appropriate test statistic and its distribution
3. Specify the significance level you are going to use and get critical value
4. Calculate value of test statistic from your data
5. See if this is more extreme than critical value (i.e. calculate a p-value)

A. Field, et al., Discovering Statistics Using R, 2012

What is the probability I would get data this (or more) extreme, assuming cellphones do not cause cancer?

1. Identify null and experimental hypotheses. — H_0,
2. Determine the appropriate test statistic and its distribution
3. Specify the significance level you are going to use and get critical value
4. Calculate value of test statistic from your data
5. See if this is more extreme than critical value (i.e. calculate a p-value)

A. Field, et al., Discovering Statistics Using R, 2012

What is the probability I would get data this (or more) extreme, assuming cellphones do not cause cancer?

1. Identify null and experimental hypotheses. — H_0,
2. Determine the appropriate test statistic and its distribution — X^2, t, F, Z, etc.
3. Specify the significance level you are going to use and get critical value
4. Calculate value of test statistic from your data
5. See if this is more extreme than critical value (i.e. calculate a p-value)

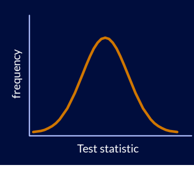

A. Field, et al., Discovering Statistics Using R, 2012

What is the probability I would get data this (or more) extreme, assuming cellphones do not cause cancer?

1. Identify null and experimental hypotheses. — H_0,
2. Determine the appropriate test statistic and its distribution — X^2, t, F, Z, etc.
3. Specify the significance level you are going to use and get critical value — *p=0.05 (convention)*
4. Calculate value of test statistic from your data
5. See if this is more extreme than critical value (i.e. calculate a p-value)

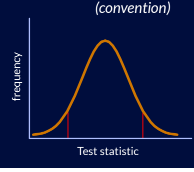

A. Field, et al., Discovering Statistics Using R, 2012

What is the probability I would get data this (or more) extreme, assuming cellphones do not cause cancer?

1. Identify null and experimental hypotheses. — H_0,
2. Determine the appropriate test statistic and its distribution — X^2, t, F, Z, etc.
3. Specify the significance level you are going to use and get critical value — p=0.05 (convention)
4. Calculate value of test statistic from your data
5. See if this is more extreme than critical value (i.e. calculate a p-value)

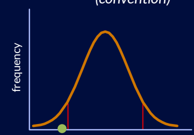

A. Field, et al., Discovering Statistics Using R, 2012

What is the probability I would get data this (or more) extreme, assuming cellphones do not cause cancer?

1. Identify null and experimental hypotheses. — H_0,
2. Determine the appropriate test statistic and its distribution — X^2, t, F, Z, etc.
3. Specify the significance level you are going to use and get critical value — p=0.05 (convention)
4. Calculate value of test statistic from your data
5. See if this is more extreme than critical value (i.e. calculate a p-value)

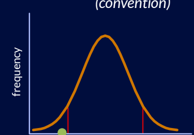

A. Field, et al., Discovering Statistics Using R, 2012

What should I do with the p-value?

Option 1 (Fisherian method)	Option 2 (Neyman-Pearson method)
Report the p-value without any statement of "rejecting the null hypothesis"	"Reject" the null hypothesis if p-value is below significance threshold Assumes you specified alternate hypothesis (Balance type I and type II errors)

In both cases, report the statistical test, test statistic, degrees of freedom, etc.
Not just the p-value!

H. Motulsky, 2018; G. Gigerenzer, 2004.

Ok…but what is a p-value?

Ok…but what is a p-value?

“a *p*-value is the probability under a specified statistical model that a statistical summary of the data (e.g., the sample mean difference between two compared groups) would be equal to or more extreme than its observed value.”

R. Wasserstein, et al., ASA Statement on Statistical Significance and P-Values, American Statistician (70), 2016.

Ok…but what is a p-value?

“a *p*-value is the probability under a specified statistical model that a statistical summary of the data (e.g., the sample mean difference between two compared groups) would be equal to or more extreme than its observed value.”

Not clear?

R. Wasserstein, et al., ASA Statement on Statistical Significance and P-Values, American Statistician (70), 2016.

Some facts about p-values and NHST

A p-value is the probability of obtaining data equal to or more extreme than the data you actually collected if you assume the null hypothesis is true.

A p-value is always between 0.0 and 1.0

Scientific conclusions should not be solely based on whether a p-value passes a threshold

Using a threshold of $p<0.05$ is completely arbitrary

p-values depend on exact experimental setup (some of which can be implicit)

p-values depend on the sample size and spread of the data

H. Motulsky, Intuitive Biostatistics: A nonmathematical Guide to Statistical Thinking, 4th ed., Oxford University Press, 2018.

What a p-value is NOT

A p-value is not the probability that the null hypothesis is true

1.0 minus the p-value is not the probability the alternative hypothesis is true

1.0 minus the p-value is not the probability the results will hold up under repeated experiments

A high p-value does not mean the null hypothesis is true

a p-value is not the probability of your results being “a random coincidence”

p-values are not a measure of the effect size or importance of a result

H. Motulsky, Intuitive Biostatistics: A nonmathematical Guide to Statistical Thinking, 4th ed., Oxford University Press, 2018.

The ASA's statement on p-values: context, process, and purpose

Ronald L. Wasserstein & Nicole A. Lazar

1. *P-values can indicate how incompatible the data are with a specified statistical model.*
2. *P-values do not measure the probability that the studied hypothesis is true, or the probability that the data were produced by random chance alone.*
3. *Scientific conclusions and business or policy decisions should not be based only on whether a p-value passes a specific threshold.*
4. *Proper inference requires full reporting and transparency*
5. *A p-value, or statistical significance, does not measure the size of an effect or the importance of a result.*
6. *By itself, a p-value does not provide a good measure of evidence regarding a model or hypothesis.*

R. Wasserstein et al., The American Statistician, 2016.

So why do we use NHST?

It is objective in that everyone will agree on the p-value for given data/statistical test/experimental design

Doesn't require prior probabilities

Requires careful description of the experiment and choice of p-value thresholds ahead of time

Computationally simple (and widely available)

Used for a long time (over 100 years)

http://cfss.uchicago.edu/persp017_inference.html

The key is to understand the limits of each method

Example - Two approaches, two different answers

a-priori probability of machine fault → 0.01 P(fault)
probability QA test (T) is positive fault exists → 0.99 P(test+ | fault)
probability QA test (T) is positive if no fault → 0.01 P(test+ | no fault)

What is the probability P(fault | test+) that a positive test means machine really has the fault?

Using Bayes' theorem	Using NHST (H_0 = machine does NOT have a fault)
P(fault | test+) = 0.50	P(test+ | no fault) = 0.01 so p=0.01 (in this case)

Some of our QA tests show this difference in action!

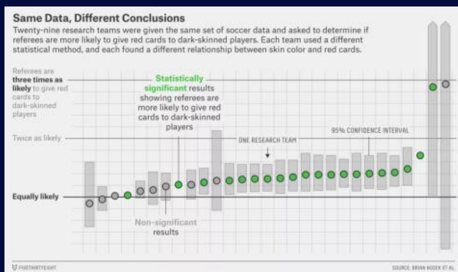

http://fivethirtyeight.com/features/science-isnt-broken/#part2

So now…back to our quiz

Which statements are true? (several or none may be correct)

Remember: (t=2.7, df=18, p=0.01)

1) You have absolutely disproved the null hypothesis

2) You have found the probability of the null hypothesis being true

3) You have absolutely proved your experimental hypothesis

4) You can deduce the probability of the experimental hypothesis being true

5) If you decide to reject the null hypothesis, you know the probability that you are making the wrong decision

6) You have a reliable experimental finding – if you repeat the experiment a large number of times you would obtain a significant result 99% of the time.

G. Gigerenzer et al., pub in The Sage handbook of quantitative methodology for social sciences, 2004.

Table 1
Percentages of False Answers (i.e., Statements Marked as True) in the Three Groups of Figure 1

Statement (abbreviated)	Germany 2000: Psychology students	Germany 2000: Professors and lecturers: not teaching statistics	Germany 2000: Professors and lecturers: teaching statistics	United Kingdom 1986: Professors and lecturers
1. H_0 is absolutely disproved	34	15	10	1
2. Probability of H_0 is found	32	26	17	36
3. H_1 is absolutely proved	20	13	10	6
4. Probability of H_1 is found	59	33	33	66
5. Probability of wrong decision	68	67	73	86
6. Probability of replication	41	49	37	60

Note. For comparison, the results of Oakes' (1986) study with academic psychologists in the United Kingdom are shown in the right column.

G. Gigerenzer et al., pub in The Sage handbook of quantitative methodology for social sciences, 2004.

Reporting statistical tests with p-values has become a de-facto requirement

Medicine increasingly relies on p-values

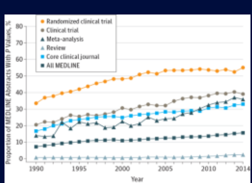

Article Category	Abstracts, No.	Abstracts Reporting P Values, No. (range per year)
Randomized clinical trial	332 262	165 968 (2056-12 101)
Clinical trial	197 921	60 317 (736-3700)
Meta-analysis	52 639	16 018 (30-2960)
Review	1 396 965	17 125 (176-2245)
Core clinical trial	796 103	210 045 (2907-10 090)
All MEDLINE	12 821 790	1 608 736 (20 709-138 218)

Original Investigation
Evolution of Reporting *P* Values in the Biomedical Literature, 1990-2015

JAMA The Journal of the American Medical Association

Chavalarias et al., JAMA, 3/15/2016.

But there are problems…..

p-values just below p=0.05 are over-represented

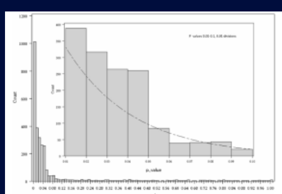

The distribution of probability values in medical abstracts: an observational study

Bastiaan Ginsel[1*†], Abhinav Aggarwal[1†], Wei Xuan[2] and Ian Harris[1†]

Conclusions:

p-values immediately below 0.05 appear to be over-represented in the literature relative to their expected frequency

Shows evidence of systematic error including publication bias, selective reporting, methodological errors, or fraud.

B. Ginsel, et al., BMC Res Notes 8, 2015.

Authors love to invent ways of getting around the "rules"

"a trend towards significance (p=0.06)"

"not absolutely significant but very probably so (p>0.05)"

"not significant in the narrow sense of the word (p=0.29)"

"teetering on the brink of significance (P=0.06)"

"tantalizingly close to significance (p=0.104)"

"possibly statistically significant (p=0.10)"

"a nonsignificant trend toward significance (p=0.1)"

https://mchankins.wordpress.com/2013/04/21/still-not-significant-2/

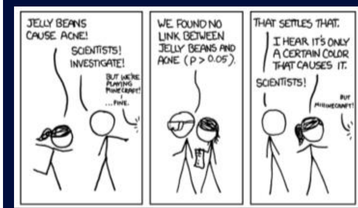

The search for a p-value you like

P-hacking
Data dredging
Cherry-picking
etc.…

https://xkcd.com/882/

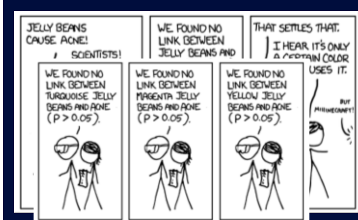

The search for a p-value you like

P-hacking
Data dredging
Cherry-picking
etc.…

https://xkcd.com/882/

The search for a p-value you like

P-hacking
Data dredging
Cherry-picking
etc.…

https://xkcd.com/882/

The search for a p-value you like

- P-hacking
- Data dredging
- Cherry-picking
- etc….

https://xkcd.com/882/

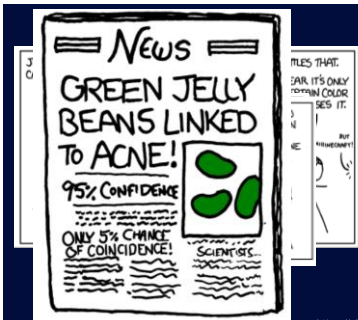

The search for a p-value you like

- P-hacking
- Data dredging
- Cherry-picking
- etc….

https://xkcd.com/882/

Try it yourself - there are many ways to achieve a desired story

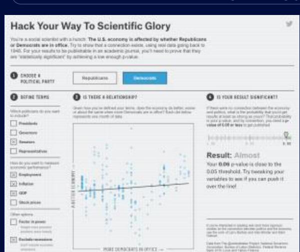

FiveThirtyEight

http://fivethirtyeight.com/features/science-isnt-broken/#part2

Many papers are published with statistical shortcomings

Category	New England J Med (n=31)		Nature Med (n=22)	
	N	%	N	%
Use of a wrong or suboptimal statistical test	5	16.1	6	27.3
No sample size / power calculation	13	41.9	22	100.0
Failure to prove test assumptions are not violated	16	51.6	13	59.1
Failure to define all statistical tests clearly and correctly	20	64.5	20	90.9
Failure to state which values of p indicate statistical significance	14	45.2	15	68.2

Excerpted from A. Strasak et al., American Statistician (61), 2007

(Several) Concluding thoughts....

Statistical literacy is only becoming more important

Characteristic	Article Year			p-value (X^2 test)
	1990 (n=133)	2000 (n=122)	2010 (n=106)	
SAS	5.3%	25.4%	49.1%	<0.001
STATA	0.0%	5.7%	32.1%	<0.001
SPSS	2.3%	4.9%	13.2%	0.002
t-test	21.1%	25.4%	26.4%	0.577
Chi-square	40.6%	41.8%	41.8%	0.471
Multiple regression	24.1%	42.6%	48.1%	<0.001
Survival analysis	14.3%	22.1%	43.4%	<0.001

Excerpts from tables in Arnold LD, et al., PLOS ONE 8(10), 2013.

There is no substitute for scientific reasoning

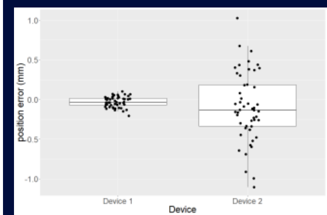

Let's say we have 2 positioning devices and we are comparing position error

Device 1: -0.029 (0.067)mm
Device 2: -0.097(0.436)mm

Welch t-test: t=1.097, df=51.325, p=0.2778

Levene's Test: F value=46.374, p<0.001

Some (I think) great resources

G. Gigerenzer et al., "The Null Ritual: What You Always Wanted to Know about Significance Testing but Were Afraid to Ask", in The Sage handbook of quantitative methodology for the social sciences, Sage, 2004.

H. Motulsky, Intuitive Biostatistics: A nonmathematical Guide to Statistical Thinking, 4th ed., Oxford University Press, 2018.

A. Field, et al., Discovering Statistics Using R, Sage, 2012. (There is a new SPSS edition available as well).

Special thanks to:

Jeff Sloan, Ph.D., Mayo Clinic

Michael Altman, Ph.D., Washington University

Questions?

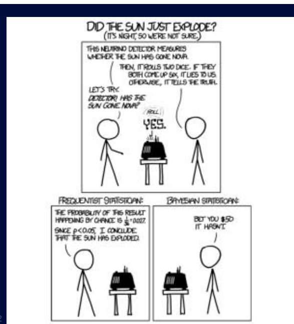

HTTP://XKCD.COM/1132