1 Risk of microcephaly associated with Zika virus infection in women infected in the first trimester.

Refer to The 2016 Lancet article "Association between Zika virus and microcephaly in French Polynesia, 2013-15: a retrospective study". It is available under the 'Zika Virus, 2016: May 21' link in this webpage: `http://www.biostat.mcgill.ca//hanley/bios601/Applications/`.

The objective is to make a simple estimate of the rate, using nothing more than the 'numerator' data in the bottom of Figure 2, the weekly consultation numbers* in the top of the Figure, the estimated 66% seroprevalence at the end of the outbreak, and the estimated numbers of women becoming pregnant each week. So, the main task is estimating (and then summing) the exposed denominators for each week. (*Hint: a Lexis diagram, or matrix, with calendar weeks as columns, and gestational weeks as rows, may help with the bookkeeping.*) To simplify matters, assume the background rate, i.e., the rate in the absence of the Zika virus, is 0%.

Once you have arrived at a numerator and an estimated denominator, calculate a point estimate and 95% CI for the rate, and compare them with the 95 cases (34-191) per 10000 estimated by the authors.

Your CI will only reflect the binomial variation. What additional sources of uncertainty might you wish to add in? Can you come up with a way to add them in?

Try to rank the various uncertainties as to how much they should widen the CI.

`* JH digitized this graph and arrived at these weekly numbers:`
`619,1163,1556,1007,1190,2041,2837,3079,3600,3089,2289,1619,1210,`
`1182,1045,895,952,583,373,286,197,137,138,152,86,36,7,17,5,11,9`

2 Bicycle weight and commuting time: randomised trial

`http://www.biostat.mcgill.ca//hanley/bios624/Resources/BikePackage.zip`.

Re-analyze the data.

3 Broad-Spectrum Sunscreen Use and the Development of New Nevi in White Children: A Randomized Controlled Trial

Water fluoridation, tooth decay in 5 year olds, and social deprivation

The relation between the availability of neonatal intensive care and neonatal mortality

Refer to Table 4. (Parameter Estimates for Variables Predicting Number of New Nevi in Vancouver Schoolchildren) in the JAMA report, available here: `http://www.biostat.mcgill.ca//hanley/c678/sunscreen.pdf`.

In the text, in reference to this table, the authors say that "A model of the effects of the independent variables on the whole-body number of new nevi is presented in TABLE 4. Total sunlight exposure, adjusted for clothing, school grade (age), the interaction term for sunscreen group, and degree of facial freckling, appears to predict nevus counts. The interaction between being randomized to the broad-spectrum sunscreen group and degree of freckling is statistically the strongest predictor of new nevi."

(Guided by the much more informative Figure 2) re-design the content of Table 4, maybe turning it into a Figure instead of a Table. Also, re-write the text to make the results more reader-friendly and readily understandable.

(Along the same lines) could the text of the article 'Water fluoridation, tooth decay in 5 year olds, and social deprivation (Jones CM et al. BMJ 1997;315:514-517, 30 August) – available on this course website `http://www.biostat.mcgill.ca/hanley/c678/` do with any improvement?

The Editor insisted the fitted regression be presented in a Table. Do so.

Convert the results in table 4 of this article `http://www.biostat.mcgill.ca/hanley/c678/neonatology.pdf` into a Figure that is accessible to lay people.

Here are JH's Notes from when he taught a (4-week) multiple regression course in the department's Summer School. `http://www.biostat.mcgill.ca/hanley/c678/class_06.pdf`.

4 Fortification of Foods with Folic Acid: Impact on Neural Tube Defects

- Neural Tube Defects and Folic Acid: https://www.marchofdimes.org/complications/neural-tube-defects.aspx, https://en.wikipedia.org/wiki/Spina_bifida and https://www.canada.ca/en/public-health/services/pregnancy/folic-acid.html
- Early study that showed the link: http://www.medicine.mcgill.ca/epidemiology/hanley/minimed/talk-MiniMed-v20141021h.pdf#page=36 taken from this 1965 'preliminary communication' http://www.biostat.mcgill.ca/hanley/c626/HibbardAndSmithells1965.pdf.
- Subsequent (non-randomized) clinical trial: http://www.biostat.mcgill.ca/hanley/c626/SmithellsTrialFolicCandNTDs1980.pdf.
- 1991 results of the MRC randomized clinical trial (RCT) 'MRC Vitamin Trial': http://www.biostat.mcgill.ca/hanley/c626/MRCvitaminTrial.pdf.
- More history: https://academic.oup.com/ije/article/40/5/1154/660590 and https://onlinelibrary.wiley.com/doi/full/10.1002/bdra.20544.
- Yongling Xiao's 2007 student presentation in bios60 http://www.medicine.mcgill.ca/epidemiology/hanley/bios601/CaseStudies/index2007.html
- R meta-analysis: https://www.ncbi.nlm.nih.gov/pubmed/28151610
- Talk by Nigel Paneth at this 2019 symposium: Start at 1.48.30 in the morning session https://www.niehs.nih.gov/news/events/pastmtg/2019/epid_2019/index.cfm

Refer to the 2001 USA and 2002 Canadian publications available on this website http://www.biostat.mcgill.ca/hanley/c626/index.html#folicAcid. The '**Exercise** on above material' is below the two sets of links, i.e., at http://www.biostat.mcgill.ca/hanley/c626/folicAcid.html

[*The 626 website now has some additional US material from 2005, as well as a news item from 2001 describing the (contrary) UK view on compulsory folic acid supplementation.*]

5 High risk of HIV-1 infection for first-born twins born to an infected mother

Refer to the article 'Twin Data that Made a Big Difference, and that Deserve to be Better-Known and Used in Teaching', available here http://www.tandfonline.com/doi/full/10.1080/10691898.2017.1381055 .

Repeat the analyses presented in the article (data are in Supplemental Material)

The Lancet 1991 and J Pediatrics 1995 articles are available here http://www.biostat.mcgill.ca/hanley/bios602/CondnlLogisticRegrn/index.html, along with some other material.

A link to the (flawed) study Induced Abortion and Secondary Infertility can also be found on the site.

If you wish to see how studies involving a sensitive topic might be better carried out, see http://www.med.mcgill.ca/epidemiology/hanley/c609/practicum2011/BrCaAbortion.pdf and http://www.med.mcgill.ca/epidemiology/hanley/c609/practicum2011/CaseControlStudiesAbortions.pdf.

6 ADHD & traffic accidents

Refer to the 'Fractions' topic in this course webpage http://www.biostat.mcgill.ca/hanley/c609/material/index.html#EF. and carry out the portion of the exercise http://www.biostat.mcgill.ca/hanley/c609/material/exerciseEFandPF2012.pdf that pertains to the Redelmeier publication on Road Trauma in Teenage Male Youth with Childhood Disruptive Behavior Disorders. http://www.medicine.mcgill.ca/epidemiology/hanley/c609/Material/redelmeierTeenagers.pdf.

7 ADHD & traffic accidents

Refer to the Supplementary Exercise 16.1 in JH's BIOS602: Notes, C&H. Ch 16 (Case Control Studies); 17 (Likelihoods for the odds ratio). Fall 2010. http://www.biostat.mcgill.ca/hanley/bios601/ch-notes-16-17.pdf.

Do questions (iii) to (vi).

There is some R code on 'Combining information on same parame-

ter' in http://www.medicine.mcgill.ca/epidemiology/hanley/bios602/MultilevelData/index.html.

8 The lacing defence: double blind study of thresholds for detecting addition of ethanol to drinks

Refer to the 'Detecting ethanol in drinks' material under 'Datasets' in http://www.medicine.mcgill.ca/epidemiology/hanley/c622/. After reading the documentation, but without consulting the (BMJ) article itself, come up with your own analysis, and an Abstract/Figure/Table that summarizes your findings.

9 On the Relation of the Direction of the Wind to the Age of the Moon

The full material on this topic is available at the bottom of this page. http://www.medicine.mcgill.ca/epidemiology/hanley/c622/. After reading the Excerpt, but without consulting the full article, write down your own '*data collection and computerization plan*' [write it as Airy might have done on his way back from his voyage – but assume it is 2018, where he wants the data entered into Excel or the like]. Then write out your '*statistical analysis plan*' [before looking at the data – i.e., as if you were applying for funding to carry out the project, and as if the plan were a set of instructions for a research assistant]. The latter should include an Abstract, and a Figure/Table that summarizes your findings [again, do this *pre* data-analysis – pretend to have analyzed them, and use made-up numbers as possible answers.]

After you have done so, refer to the summary data presented in Airy's table.

First, are you happy with the title of the table? Would you have a good sense of its contents, if (ahead of purchasing it for $45 of your own funds) this title was all you had to go on? If not, suggest a better title.

If possible, apply your method to these summary data, so that you can give some numbers to go along with Airy's inference from the 'general result' contained in the Table

> And, while it shows that there is *great uncertainty in the verification of an empirical law,* even from nearly ninety lunations, *it seems very distinctly to negative the asserted law which gave rise to the inquiry.*[Italics by JH]

10 Effects of neonicotinoid pesticides on honey bees and wild bees

Refer to the article and to the data-tables in the supplementary material http://www.medicine.mcgill.ca/epidemiology/hanley/bios601/Applications/BeesNeonicotinoids.pdf

To make it easier, JH has made a single datafile http://www.medicine.mcgill.ca/epidemiology/hanley/bios601/Applications/dsABC.txt with the columns of data from Tables S2A, S2B and S2C.

Re-do the data-analysis 'your way.' Explain the models/approaches the authors used, and how yours differ from (or are the same as) them.

11 Quantification of *Mycobacterium bovis* transmission in a badger vaccine field trial

Refer to the scientific report and newspaper article near the bottom of this site http://www.medicine.mcgill.ca/epidemiology/hanley/c622/.

Make your own lay summary. Also prepare a more technical paragraph that explains why they used that generalized linear model.

12 (Left)Handedness and Longevity

See publications by Halpern and Coren in Nature (1998) and NEJM (1991), as well as other authors, under 'Longevity Comparisons' in http://www.biostat.mcgill.ca/hanley/bios601/CandHchapter06/

13 Sequelae of first trimester rubella infections

http://www.medicine.mcgill.ca/epidemiology/hanley/minimed/talk-MiniMed-v20141021h.pdf

http://www.medicine.mcgill.ca/epidemiology/hanley/minimed/Gregg1941.pdf

http://www.medicine.mcgill.ca/epidemiology/hanley/minimed/DeafnessRubellaAustralia.pdf

14 Statistical Sudoku – Penrose's data on Downs Syndrome and Maternal Age

see http://www.biostat.mcgill.ca/hanley/StatisticalSudoku/

15 Effects of experience and commercialisation on survival in Himalayan mountaineering: retrospective cohort study

Contact JH for data.

16 Does month-of-birth influence the probability of reaching the professional sports leagues? (NHL, NBL, NBA, ...)

Refer to Supplementary Exercise 16.5 in the Notes regarding Chapter16 of Clayton & Hills. http://www.epi.mcgill.ca/hanley/bios602/ch-chapter16PlusNotes2015.pdf .

The R code and data can be found in bullets 1 and 4 under Software and Data near the bottom of this webpage http://www.epi.mcgill.ca/hanley/bios601/CaseControlStudies/index.html .

17 The full moon and motorcycle related mortality

See http://www.biostat.mcgill.ca/hanley/c622/RedelmeierFullMoon.pdf for the article and http://www.biostat.mcgill.ca/hanley/c622/Persons.txt for the dataset JH extracted from the FARS database. This User's Manual https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812827 explains the variables in upper-case. The lower-case variable names were derived by JH. The variable lune uses the numbering system in the website https://www.timeanddate.com from which JH scraped this file http://www.biostat.mcgill.ca/hanley/c622/MoonData.csv.

18 How much per square foot? Pattern in condominium prices

Here is a page from a 2005 brochure http://www.biostat.mcgill.ca/hanley/c634/rates/leloft2.pdf.

Assume that back in 2005 the vendor has a formula for the price per square foot, that may included the location within the floor, as well as the floor. [Quite apart from inflation, this function would probably be different today, as tall buildings now change the views (see http://lelowney.prevel.ca/en/condos-plans/.)]

Under 2005 'Condo Prices' on this site http://www.medicine.mcgill.ca/epidemiology/hanley/c634/rates/index.html#CondoPrices you will find the raw data and some SAS and R code. [Ignore the 2020 data]

The data-analysis objective is to figure out what formula the realtor might have used for the \$/pi^2 'rate' as an (increasing) function of the floor. Other unit-specific factors, such as location on floor, probably entered into it too, but ignore them for now; likewise, ignore the first (ground) floor units. Here are two 'rate' possibilities:

(i) a rate that was linear for every extra floor above the second
..... e.g.rate = \$/pi^2 = \$200 + \$5 for every extra floor ('additive')

(ii) a 'multiplicative' one, where rate was more progressive,
..... e.g. (a) rate = \$/pi^2 = \$ 200 $\times (1 + 0.025)^{no.extra.floors}$
..... (b) rate = \$/pi^2 = \$200 $\times e^{0.025 \times no.extra.floors}$.

These are analogs of 'simple-' and 'compound-interest' models

(i) Using model (i), write out the formula for the (average) price of a unit as a function of floor and square footage, and use the dataset to estimate the "base", i.e. "floor-2", rate (what Clayton and Hills http://www.medicine.mcgill.ca/epidemiology/hanley/ch/Statistical%20Models%20II%20-%2022.pdf call the "corner") and "per extra floor" amount in 'rate model' (i).

(ii) Using model (ii b), write out the formula for the (average) price of a unit as a function of floor and square footage, and use the dataset to estimate the corresponding coefficients in 'rate model' (ii b). Compare the parameter estimates, and fits, of the 2 models. Work in the \$/pi^2 scale, not the log[\$/pi^2] one!

19 The lifespan of male fruitflies: the effect of sexual activity

`http://www.biostat.mcgill.ca/hanley/bios602/LifetableRegression/FruitfliesPartridgeFarquhar1981.pdf` and (for data and documentation) `http://www.biostat.mcgill.ca/hanley/c622/index.html`

See also the JH article `http://www.biostat.mcgill.ca/hanley/Reprints/Multivariate_analysis_1983.pdf` that uses these data to illustrate the role of multiple regression in making contrasts FAIRER and SHARPER.

20 Survival/Longevity, Mortality Rates

Refer to the material on the website for course 634. `http://www.medicine.mcgill.ca/epidemiology/hanley/c634/`, and in particular the more complete set of topics covered when it was first given in 2006-2007 `http://www.medicine.mcgill.ca/epidemiology/hanley/c634/index2006-2007.html` .

See in particular the exercises, beginning March 2, 2007, related to

• the longevity of Oscar actors/actresses `http://www.epi.mcgill.ca/hanley/Reprints/Sylvestre_et_al_OSCARS.pdf`. The data/documentation are 1/2 way down this page `http://www.epi.mcgill.ca/hanley/c634/rates/`.

• Weekend versus Weekday Admission and Mortality from Myocardial Infarction `http://www.biostat.mcgill.ca//hanley/bios601/BiasReduction/weekend-weekday-mi.pdf`. The data and R code: the fourth bullet here: `http://www.epi.mcgill.ca//hanley/c681/cox/index.html` .

21 Fitting a Prognostic Probability Model to Clinical-trial Data

This activity is sometimes (quite inaccurately) called 'risk prediction.'

When considering treatment options, a physician needs to know the prognosis corresponding to the risk profile of the patient seeking treatment. However, as we describe in our 2008 article `http://www.medicine.mcgill.ca/epidemiology/hanley/Reprints/Julien-Hanley-Clin-Trials-2008.pdf` ('Profile-specific survival estimates: Making reports of clinical trials more patient-relevant'), reports of clinical trials generally address treatment-specific survival probabilities only in the aggregate, i.e., for the typical patient, and often express the difference in survival as a hazard ratio.

Such summaries do not provide treatment-specific survival probabilities (and thus the absolute difference in these probabilities) for patient profiles that are not near the typical ('average') of those in the trial. Despite the fact that Cox intended his hazard regression method to be used to produce such profile-specific survival estimates, and even showed how to calculate them, many investigators are either unaware that this is possible, or else choose not to report them.

In a presentation JH made on this at the time
http://www.medicine.mcgill.ca/epidemiology/hanley/Reprints/StMarysPresentation.pdf
he imagined that he had been contacted by a male relative of a certain age who had recently been diagnosed with localized prostate cancer of a certain aggressiveness [a Gleason Score, measured on a scale from least (2) to most (10) aggressive]. Knowing that JH was a medical statistician, the relative asked JH to combine his profile with the results in the Albertsen et al. (JAMA) paper and the subsequent Bill-Axelson et. al the (NEJM) paper on the RCT to tell him what would be the prognosis (probability of dying of prostate cancer) if he were to (a) undergo or (b) decline a radical prostatectomy. Unless it is an 'average risk' case, JH would have to say that the reported trial results were not sufficient.

In this exercise, you are asked to use your favourite software package, along with the relevant code given in Table A1 of the Clin-Trials-2008 article, to calculate profile-specific survival estimates from a Cox model fitted to data from The Systolic Hypertension in the Elderly Program (SHEP) trial.[1]

[1] This type of probability-fitting is not limited to data from clinical trials; its counterpart in non-experimental studies (minus the treatment option) is nicely illustrated by these two 'risk calculators' for cardiovascular risk assessment, derived using data from the Framingham Heart Study. `https://www-uptodate-com.proxy3.library.mcgill.ca/contents/calculator-cardiovascular-risk-assessment-10-year-general-cardiovascular-disease-men-framingham-2008-paper` for men, and for women `https://www-uptodate-com.proxy3.library.mcgill.ca/contents/calculator-cardiovascular-risk-assessment-10-year-general-cardiovascular-disease-women-framingham-2008-paper` for cardiovascular risk assessment, derived using data from the Framingham Heart Study.

The SHEP dataset `dataSHEPtrial4701.txt` (used in the Clin-Trials-2008 article) is available in the `zip` file accessible from the top of the page at this link `http://www.medicine.mcgill.ca/epidemiology/hanley/software/`.

if you are up for it, you can also (or instead) try to fit the 'smooth-in-time' hazard function models that Hanley and Miettinen used to come up with an alternative to the 'step-function' risk functions[2] one is limited to by the Cox model. Now that two of our biostatistics students (SB and MT) and a former younger colleague (OS) have produced the `casebase` package for `R`, it is much easier than it was when they published the 'Fitting of Smooth-in-Time Prognostic Risk Functions via Logistic Regression' in The International Journal of Biostatistics in 2009.

22 The Brinks Case

The 'documentation' on this site `http://www.medicine.mcgill.ca/epidemiology/hanley/c678/#brinks` contains a summary of a legal case involving 'shortfalls' in the amounts of money collected by Brinks employees from New York City parking meters some 40 years ago (when the average revenue per meter-day was close to 1 dollar). It also tells you the meaning of the columns of data in the 'Brinks data' file. **In this exercise**, using just these data, you are asked to come up with your estimate of the amount of the shortfall – the amount Brinks should compensate the city for.

(A few years ago, after JH contacted him, William Fairley kindly supplied the additional Table and Graphs and his and Bruce Levin's testimony at the trial).

JH was reminded of this Brinks case when he read the following story in the Montreal Gazette: `http://www.medicine.mcgill.ca/epidemiology/hanley/c678/EmployeesWhoCollectedPparkingMeterCoinsConvictedOfTheft.pdf`.

He is attempting to obtain the 'analysis done by a forensic accountant' so as to provide a more local, and contemporary data-challenge.

Incidentally, why do the equations shown in the green box above the input section have the form they have?

10-Year Risk $= 100 * (1 - 0.88936^{\exp[Risk factors]})$[men] and $100 * (1 - 0.95012^{\exp[Risk factors]})$[women]

McGill subscribes to `https://www.uptodate.com/home`, widely used by clinicians. But JH believes that one could teach an entire course on regression models using the medical calculators: `https://www.uptodate.com/home/medical-calculators`.

[2] This short note `http://www.medicine.mcgill.ca/epidemiology/hanley/Reprints/BreslowEstimatorEpidemiology2008.pdf` tries to explain how one can come up with a 'baseline' hazard function (i.e., one for the profile where all the X values are zero) from a (Cox) model that purports to *not* model it.

23 Report on certain enteric fever inoculation statistics. (Meta-analysis by K. Pearson, 1904)

See `http://www.biostat.mcgill.ca/hanley/bios602/MultilevelData/PearsonMetaAnalysis.pdf`

24 Effect of Sunscreen Application Under Maximal Use Conditions on Plasma Concentration of Sunscreen Active Ingredients

The article and the raw data are available here `http://www.biostat.mcgill.ca//hanley/c622/SunscreenIngredientsPlasma.zip`. The zip file also contains a newer second study, with a larger sample size, that addresses 'the systemic absorption and pharmacokinetics of the 6 active ingredients in 4 sunscreen products under single- and maximal-use conditions.'

25 Chlamydia pecorum prevalence in south Australian koala populations: Identification of a population [in Kangaroo Island] free from infection?

The article and extensive supplement are available here:
`http://www.biostat.mcgill.ca//hanley/c622/KoalasChlamydiaKangaroosIsland.zip`.
The article contains links to the model to assess freedom from C. pecorum (implemented in MS Excel (2013) and to the dataset.

26 Training for a First-Time Marathon Reverses Age-Related Aortic Stiffening

The article is available here `http://www.biostat.mcgill.ca//hanley/c622/MarathonsArteries.pdf`.

27 Effect of Folic Acid and Zinc Supplementation in Men on Semen Quality and Live Birth Among Couples Undergoing Infertility Treatment

The article, and a media story, are available here http://www.biostat.mcgill.ca//hanley/c622/FolicAcidZincFertility.zip.

28 Multi-modal survey of Adélie penguin mega-colonies

The article and supplement are available here http://www.biostat.mcgill.ca//hanley/c622/CountingPenguins.zip.

29 Hockey Games and the Incidence of ST-Elevation Myocardial Infarction

The article and associated materials are available here http://www.biostat.mcgill.ca//hanley/c622/HabsGames.zip.

30 Over the Hill at 24: Persistent Age-Related Cognitive- Motor Decline in Reaction Times in an Ecologically Valid Video Game Task Begins in Early Adulthood

The article is here http://www.biostat.mcgill.ca/hanley/bios601/Surveys/OverTheHillAt24.PDF. Other material on reaction times (and their measurement) can be found 7/8ths the way down this page, http://www.biostat.mcgill.ca/hanley/bios601/Surveys/index.html under the heading 'Measuring reaction times'.

31 Patterns of performance degradation and restoration during sleep restriction and subsequent recovery: a sleep dose-response study

Article is here: http://www.biostat.mcgill.ca/hanley/bios601/Surveys/Belenky2003.pdf. Related material on reaction times (and their measurement): 7/8ths the way down this page, http://www.biostat.mcgill.ca/hanley/bios601/Surveys/index.html under the heading 'Measuring reaction times'. Dataset used here.

32 Galton's Anthropometric Laboratory

2/3 the way down the same (Surveys) link. The article is here: http://www.biostat.mcgill.ca/hanley/bios601/Surveys/GaltonsData100yearsLater.pdf

33 Does Exposure to Scientific Theories Affect Women's Math Performance?

Article, supplement, data provided by authors, and some notes by JH on 'Confounding: reducing it by regression': http://www.medicine.mcgill.ca/epidemiology/hanley/tmp/Applications/WomenMath.pdf

34 Short term adverse reactions to measles-mumps-rubella vaccine

Short version of published article 'Day-to-Day Reactogenicity of Measles-Mumps- Rubella Vaccination' (abstracted by JH): http://www.medicine.mcgill.ca/epidemiology/hanley/Reprints/RCH/03MMRvaccinesVirtanen.pdf

Thesis version http://www.biostat.mcgill.ca/hanley/c622/VirtanenThesis.pdf

35 Should you drink if you're pregnant?

Message, `http://www.biostat.mcgill.ca/hanley/c626/lbw.html`, from `https://en.wikipedia.org/wiki/Seagram` in The Montreal Gazette, Tuesday, September 27, 1983.

Article: `http://www.biostat.mcgill.ca/hanley/c626/LBW-alcoholLancet1983.pdf`

See also: Wacholder S. Binomial regression in GLIM: estimating risk ratios and risk differences. American Journal of Epidemiology. 123(1):174-84, 1986. `http://www.biostat.mcgill.ca/hanley/c626/WacholderBinomialGLIM.pdf`

and `http://www.biostat.mcgill.ca/hanley/c626/Wacholder.pdf`

36 Lidkoping Accident Prevention Program

`http://www.biostat.mcgill.ca/hanley/c626/`
`LidkopingAccidentPreventionProgram.zip`

37 Pregnancy rates in cohorts of Ontario girls eligible and ineligible for the HPV vaccination program

`http://www.medicine.mcgill.ca/epidemiology/hanley/bios601/`
`StatisticalModelsCh06-2018.pdf#page=11`

38 How the switch to/from daylight savings time affects rates of motor vehicle accidents

In (Ferguson1995's) references to US and UK studies in the 1970s and 1980s in Ferguson1995, the focus was the sudden shift in *light* conditions. In an short letter in the New England Journal of Medicine in 1998, the psychologist Staley Coren was concerned with focused on *insufficient sleep and disrupted circadian rhythms.* In it he addressed suggestions that "as a society we are chronically sleep-deprived and that small additional losses of sleep may have consequences for public and individual safety," and went on to "use noninvasive techniques to examine the effects of minor disruptions of circadian rhythms on normal activities" by taking "advantage of annual shifts in time keeping."'

> More than 25 countries shift to daylight savings time each spring and return to standard time in the fall. The spring shift results in the loss of one hour of sleep time (the equivalent in terms of jet lag of traveling one time zone to the east), whereas the fall shift permits an additional hour of sleep (the equivalent of traveling one time zone to the west). Although one hour's change may seem like a minor disruption in the cycle of sleep and wakefulness, measurable changes in sleep pattern persist for up to five days after each time shift. This leads to the prediction that the spring shift, involving a loss of an hour's sleep, might lead to an increased number of "micro- sleeps," or lapses of attention, during daily activities and thus might cause an increase in the probability of accidents, especially in traffic. The additional hour of sleep gained in the fall might then lead conversely to a reduction in accident rates.

To study this, he

> used data from a tabulation of all traffic accidents in Canada as they were reported to the Canadian Ministry of Transport for the years 1991 and 1992 by all 10 provinces. [...] Data for analysis were restricted to the Monday preceding the week of the change due to daylight savings time, the Monday immediately after, and the Monday one week after the change, for both spring and fall time shifts. Data from the province of Saskatchewan were excluded because it does not observe daylight savings time.

His concluded the letter by stating that

> these data show that small changes in the amount of sleep that people get can have major consequences in everyday activities. The loss of merely one hour of sleep can increase the risk of traffic accidents. It is likely that the effects are due to sleep loss rather than a nonspecific disruption in circadian rhythm, since gaining an additional hour of sleep at the fall time shift seems to decrease the risk of accidents.

For the full article, and the subsequent correspondence with the journal (including one from a researcher at Transport Canada that used data from the years 1984 to 1993), see the file 1996CorenDaylightSavings.pdf in this **zipped folder**.

The folder also contains several reports since then, including the one `2014MyocardialIfarctions.pdf` referred to by 'sleep-expert' MW in

minute 8.45 of this TED talk https://www.ted.com/talks/matt_walker_sleep_is_your_superpower/transcript?language=en#t-567051 Andrew Gelman's blog has more on some of MW's claims and some further follow-up of his work.

The US Fatal Accident Reporting System (FARS*) contains data from 1975 onwards, on every motor vehicle traffic crash involving at least 1 fatality. From it, JH was able to extract the 4-way frequency array contained in the file `fr.Rdata`. It is a $31 \times 12 \times 42 \times 4$ array of day.of.month $\times$ month $\times$ year $\times$ unit.counted counts extracted from it, based on the 42 years from 1975 to 2016, where the last dimension is 1: crashes, 2: all fatalities; 3: persons involved (fatal+non-fatal) and 4: driver fatalities.

* More at https://www.nhtsa.gov/research-data/fatality-analysis-reporting-system-fars.

Also included in the `.zip` file are 2 `.Rdata` files, each 42 years long, with the dates when Daylight Savings Time began and ended each year.

Exercise: use these data to measure the magnitudes of the shifts in crash rates following the shift to/from Daylight Savings Time.

39 The Mode of Delivery and the Risk of Vertical Transmission of Human Immunodeficiency Virus Type 1

See the article http://www.biostat.mcgill.ca//hanley/c678/MaternalChildHIVNEJMmetaAnalysis.pdf The data (from Table 4) are here http://www.biostat.mcgill.ca/hanley/c626/hiv.dat.txt. These SAS-INSIGHT- (now SAS JMP-) based notes on logistic regression http://www.biostat.mcgill.ca/hanley/c678/outlin11.pdf use this dataset as a worked example.

40 Distribution of discovered Downs' syndrome cases and of total live births by maternal age and birth order, Michigan, 1950-1964.

The data (and some SAS code, and some notes on the origin of the data) are here: http://www.medicine.mcgill.ca/epidemiology/hanley/c626/downs_sas.txt

The fuller dataset can be found in Stark CR and Mantel N (1966) Effects of maternal age and birth order on the risk of mongolism and leukemia J Natl Cancer Inst 37 687-698, available here http://www.biostat.mcgill.ca/hanley/c626/StarkMantelArticle.pdf

41 Mystery Dataset: What happened to these people?

Data: http://www.biostat.mcgill.ca/hanley/c626/mystery.txt

Graph: http://www.biostat.mcgill.ca/hanley/minimed/MysteryData.pdf

SPOILER ALERT: Other graphs of these data can be found here and here.

42 What events were these?

See graph at http://www.biostat.mcgill.ca/hanley/MysteryData/.

43 The Heritability of Otitis Media

Article: http://www.biostat.mcgill.ca/hanley/bios602/MultilevelData/otitisJAMAarticle.pdf;

Data: in .csv file, tall format;

Notes: http://www.biostat.mcgill.ca/hanley/bios602/MultilevelData/otitis-data-correspondence.txt;

R code for Anova - Pairs Only http://www.biostat.mcgill.ca/hanley/bios602/MultilevelData/otitis-analysis-anovaPairsOnly.R
input file for this: http://www.medicine.mcgill.ca/epidemiology/hanley/bios602/MultilevelData/otitis-data-R.txt

R JAGS and Winbugs code/data for MCMC estimation of variance components: http://www.biostat.mcgill.ca/hanley/bios602/MultilevelData/otitis-data-winbugs-rjags.R

44 Framingham Heart Study

Overview: https://www.nhlbi.nih.gov/science/framingham-heart-study-fhs

5209 Records from Framingham Data File:

Code sheet and some SAS code: http://www.biostat.mcgill.ca/hanley/c681/cox/fram_sas_pgm.txt;

Data: http://www.biostat.mcgill.ca/hanley/c681/cox/fram_data.txt;

More on these records, and the textbook's introduction to the study, can be found here: http://www.biostat.mcgill.ca/hanley/c626/framnghm.txt

45 Completing the Results of the 2013 Boston Marathon

Article: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3984103/

Data: here.

46 Effects of city-wide 20 mph (30km/hour) speed limits on road injuries in Bristol, UK

Article: http://www.biostat.mcgill.ca/hanley/c622/SpeedLimitsBristol.pdf

Data and R code: http://www.biostat.mcgill.ca/hanley/c622/BristolAnalysis.R.txt

47 Effect of altitude on physiological performance: a statistical analysis using results of international football games

Article: http://www.biostat.mcgill.ca/hanley/c622/altitude-football.pdf

Data (scraped by JH): http://www.biostat.mcgill.ca/hanley/c622/citydata.csv and http://www.biostat.mcgill.ca/hanley/c622/results.csv

SAS **code**: http://www.biostat.mcgill.ca/hanley/c622/fifa-sas.txt

48 Are football (soccer) referees more likely to give red cards to players with dark skin than to players with light skin?

Article: http://www.biostat.mcgill.ca/hanley/c622/crowdsourcedDataAnalysesSkinColourRedCards.pdf

Data: http://www.biostat.mcgill.ca/hanley/c622/CrowdstormingDataJuly1st.csv

R **code**: http://www.biostat.mcgill.ca/hanley/c622/SkinColourRedCardsAnalysis.R.txt

.

49 Battle for the thermostat: Gender and the effect of temperature on cognitive performance

Article: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0216362

Data: http://www.biostat.mcgill.ca/hanley/c622/

TemperaturesGenderData.txt

50 Sex-Age-CalendarTime Patterns in population mortality rates in Denmark

Notes by JH: http://www.medicine.mcgill.ca/epidemiology/hanley/c634/rates/epi634-2010-RegressionOverview_hw_02.pdf and http://www.medicine.mcgill.ca/epidemiology/hanley/c634/rates/epi634-2010-hw02%20-%20Ans.pdf

R code for plots and informal fit of a regression model : http://www.medicine.mcgill.ca/epidemiology/hanley/c634/rates/DanishPopMortalityRates.R

51 How long did their hearts go on? A Titanic study

Several other articles from the Christmas Edition of the British Medical Journal: http://www.biostat.mcgill.ca/hanley/c626/index.html#XmasBMJ

The article on the longevity of the Titanic passengers who survived the sinking: http://www.biostat.mcgill.ca/hanley/Reprints/article_bmj_xmas_2003.pdf

The data, along with various bits of R code are here: http://www.biostat.mcgill.ca/hanley/bios602/b-d-II-ch-1-2-3/index.html#dataCode

These have been the subject of various homework exercises in courses bios601 and bios602 over the last 15 years, such as thus one from 2017 http://www.biostat.mcgill.ca/hanley/bios601/StatisticalModelsCh06.pdf

The source of our data, the Encyclopedia Titanica website https://www.encyclopedia-titanica.org has become much fancier in the years since we extracted the data in 2003. [*Warning: the website can become somewhat of a time-sink!*] But let JH know if you are interested in tracing the 65 passengers who remained 'untraced' in 2003.

52 Mystery Canadian Data: what was going on this Sunday afternoon of 2010.02.28? And what critical element is missing from the label on the vertical axis?

See this graphic from the Edmonton Edmonton's water utility: http://www.biostat.mcgill.ca//hanley/minimed/EdmontonFigBlinded.png

- What was going on this particular day in Canada in 2010?
- On the vertical axis, ML stands for Mega Litres, or a million litres. But what word(s) is(are) missing from the label?

  By the way: the annual report from this source, http://www.biostat.mcgill.ca/hanley/minimed/EdmontonWaterUtilitiesStatistics.html and a little back-calculation, should help you figure out the *specific* missing word.

  [***Spoiler Alert***:] The List of Abbreviations in this City of Winnipeg Water and Waste Department report https://www.winnipeg.ca/waterandwaste/pdfs/water/2010WaterConsumptionSummaryReport.pdf#page=6: contains the correct label.

53 Do Oscar Winners Live Longer than Less Successful Peers?

- Here is the original 2001 article, and the 2006 re-analysis http://www.epi.mcgill.ca/hanley/Reprints/Sylvestre_et_al_OSCARS.pdf.
- Here is a sampling of news items on this question:

  2001 - 2006 - 2006 - 2007 - 2007 - 2010 - 2011 - 2012 - 2014 - 2014 - 2015
- So that you get a better sense of what is involved in a fully computerized (and thus more black-box) analysis of the full dataset, first complete Q 3 from this exercise in course 634 in 2007: http://www.medicine.mcgill.ca/epidemiology/hanley/c634/hw_4.pdf#page=2
- Several versions of the full data are available 1/2 way down this page http://www.epi.mcgill.ca/hanley/c634/rates/.

- Using the '1 record per performer' dataset, try to reproduce the longevity difference reported in the abstract of the 2001 article (Note: the dataset sent to JH included new nominees/winners and some additional followup gathered after the 2001 publication).

- To compare the mortality rates of performers while in the 'still hoping to win' state vs. in the 'already a winner' state, make sure you use a dataset that divides the performer-years into those lived as Oscar nominees and those lived as Oscar winners. Use it to compare the death rates in the performer-years spent as nominees versus those spent as winners. The different data formats allow you to use classical (non-regression) techniques such as standardized rates and Mantel-Haenszel summary rate ratios, as well as fully parametric (Poisson) and semi-parametric (Cox) regression models. Whichever one(s) you choose, apply it (them) to both version (c) – correctly allocated time – and version (d) – some time incorrectly allocated. [3] . [4]

54 Age- and calendar-year-specific rates of pregnancies in 4 Ontario birth-cohorts

Parts 1, 3 and 4 in this exercise: `http://www.biostat.mcgill.ca/hanley/bios601/StatisticalModelsCh06-2020.pdf#page=11`

55 Media interpretations of an epidemiologic research article

Here are five media reports of this article 'Dairy, soy, and risk of breast cancer: those confounded milks'
`http://www.biostat.mcgill.ca/hanley/bios691/MilkBrCa.pdf`.

- Science Daily
- CTV News
- The Montreal Gazette
- Global News
- Yahoo.com

i. ***Compare the five as for how faithfully/accurately they reported the findings described in the article***. It is more difficult to be 'detached' / 'neutral' when discussing the strengths/limitations of the studies: they are seen very differently by different constituencies, e.g., 1, 2, 3, and 4. Thus, just ***focus on the numbers shown in, and reported on, in the Int. J. Epi. article***.

ii. Use the 'linear fit' model (along with the fact that there were 1057 cases in the 7.9 years of the 52,795 women) to compute absolute 8-year risks for women at the 0, 1/3, and 1-cup consumption levels shown in Figure 1 of the IJE article. State any assumptions you made.

56 Vaccine Efficacy in BioNTech/Pfizer, Moderna, and Oxford/AstraZeneca Trials

The protocols, announcements, journal reports, editorial, and links to the video of and material presented to/discussed at the FDA meetings can be found here.

i. Comment on the null hypothesis (first sentence section 9.2, page 83 of the Moderna protocol

ii. What does a 'proportional hazards VE' (top of p. 84) mean? Is proportional hazards a reasonable assumption?

iii. What is a 'stratified Cox proportional hazard model'? How it is fitted?

iv. Why use of a *1-sided* 0.025 significance level? (c.f. FDA guidelines).

v. "The primary analysis population for efficacy will be the Per-Protocol (PP) Set, defined in Section 9.4. In the primary analysis of efficacy, cases will be counted starting 14 days after the second dose of IP."
Explain what is meant by the "Per-Protocol (PP) Set."

vi. Rework the statistical calculations shown in Table 6. *Do so 'from scratch,' using first principles, i.e. don't use the referenced* `R` *package. Begin with the simplest case, where there are no interim analyses.*

[3] It would be nice if a team that is expert in Stata would show us how, starting with the one-record per performer dataset, one deals with time-varying covariates.

[4] For more on this topic, including a recent and far more serious instance of incorrectly-allocated person time in a Danish study, see the slides and lyrics from this 2016 presentation by JH: `http://www.biostat.mcgill.ca/hanley/Reprints/3FRIAS2016_Session1_Invited1_Hanley.pdf`.

vii. Rework the statistical calculations in Table 18 of the Moderna brief.

viii. Likewise, rework the statistical calculations shown in Table 19.

ix. Some authorities are planning to use just a single dose. Can you use the data in the Kaplan Meier plots to obtain a point and and interval estimate of the efficacy of a single dose, and if you can, what are the limitations? (*You might find the information in Fig. 3 of the NEJM report of the BioNTech/Pfizer trial easier to work with*).

(See also the paragraph beginning with "Delaying or deferring the second dose of the Pfizer or Moderna vaccine has the same problem" in Risks of altering vaccination protocols are unknown by Montreal Gazette columnist Dr. Christopher Labos, the comments by McGill's Dr. Donald Vinh, as well as this BBC piece).

x. Comment on Senn's Serendipity dividend? on the AstraZeneca/Oxford announcement. "As the press release put it, *two different dosing regimens demonstrated efficacy with one showing a better profile*."

xi. Paraphrase the message in this item from the StatsChat website at the University of Auckland.

xii. Repeat the statistical calculations in the 'SputnikV' report. Also comment on any major design/analysis differences vis-a-vis the other trials you have looked at.

57 Bidirectional associations between COVID-19 and psychiatric disorder: retrospective cohort studies of 62 354 COVID-19 cases in the USA

Article, and editorial Supplement.

"In patients with no previous psychiatric history, a diagnosis of COVID-19 was associated with increased incidence of a first psychiatric diagnosis in the following 14 to 90 days compared with six other health events (HR 2.1, 95% CI 1.8–2.5 vs influenza; 1.7, 1.5–1.9 vs other respiratory tract infections; 1.6, 1.4–1.9 vs skin infection; 1.6, 1.3–1.9 vs cholelithiasis; 2.2, 1.9–2.6 vs urolithiasis, and 2.1, 1.9–2.5 vs fracture of a large bone; all $p < 0.0001$).

The data presented in this paper and the appendix can be freely accessed at* `https://osf.io/fjnw8` . "

`OSF_COVIDPsychiatricSequelae.Rdata`

The article and supplement have several 'going-up' Kaplan-Meier curves for onset of first psychiatric diagnoses after COVID-19 diagnosis compared with influenza and other cohorts. While the Number at Risk is shown for 30, 45, 60, 75, and 90 days post COVID-19 diagnosis, the daily / total numbers of new first psychiatric diagnoses up to 90 days are not shown, or mentioned in the text.

In order to better understand these curves, it would help to be able to recreate the daily numbers of patients receiving a first psychiatric diagnoses as well as the daily numbers at risk. Focus on the two curves (COVID-19 patients vs. influenza patients) in the top left [F20-F48] panel of Figure 1 in the main text, and on 'Table 1 – Characteristics of the COVID-19 and influenza cohorts before and after matching' on page 9 of the supplement.

i. Using just the numbers at risk at days 30, 45, 60, 75, and 90 shown at the bottom of the relevant Figure 1 panel, and the plotted curves, make a rough estimate of the daily and total number of patients receiving a first psychiatric diagnosis.
Hint: First: use rough interpolations to arrive at approx. daily numbers at risk, and multiply these by the daily (conditional) probabilities) of receiving a first psychiatric diagnosis. The almost linear patterns of the plotted outcome probabilities make it easy to arrive at approx. daily conditional probabilities.

ii. (*de luxe version*) Using the *data*[5] in the `OSF_COVIDPsychiatricSequelae.Rdata` file, use reverse-engineering to arrive at the daily numbers at risk and the daily and total number of patients receiving a first psychiatric diagnosis.
Hint: Among others, this involves knowing and using the (Greenwood) formula for the variance of the $\widehat{S(t)}$ *at each day* t*, and figuring out which* `conf.type` option in `survfit(Surv(time,dx) 1,conf.type=?)` *was used to produce the lower and upper limits of the confidence interval.*

iii. Comment on the follow-up patterns you found, on why they might have arisen, and on what might be their implications.

[5]* The `Rdata` file does not contain raw (individual-level) data. It contains the information needed to reproduce the K-M plots and the shaded areas representing 95% CIs.

58 'Irish babies born on January 1st expected to live to 105 years of age – UN'

On reading this article on the website of this newspaper, **JH emailed the journalist**

Sent: 01 January 2021 03:58
To: Sorcha Pollak <spollak@irishtimes.com>
Subject: Life expectancy story (Irish Times)

Hi Sorcha

Interesting story for 2021

I went to this site to see where you found the numbers for Ireland and the other countries you mention, https://population.un.org/wpp/DataQuery/
Might you send me a screenshot / description of what setting you used (as I cannot get above about 80 y., or maybe 85 if I chose girls).
Your 105 seems very optimistic

I wonder if the UN did these calculations before COVID came over the horizon.

(it would be interesting to think about what if you had written this story 100 years ago, what no. you would have given, and it might have been a bit different again if you had selected 1918 or 1919, although the Spanish flu was not that tough on infants.. it hit your age group more)

Best for 2021
James Hanley

http://www.medicine.mcgill.ca/epidemiology/hanley/
http://www.medicine.mcgill.ca/epidemiology/hanley/BridgeOfLife/
http://www.medicine.mcgill.ca/epidemiology/hanley/Pandemics/

I use life expectancy a lot in my teaching...

incl. this one told by my sec. school teacher at our 50th re-union 5 years ago !! http://www.medicine.mcgill.ca/epidemiology/hanley/bios601/ch05-2020-orig-notes-exercises.pdf#page=14 (Supp. Exercise 5.6).
Am always on lookout for journalism stories to make the material more engaging.
(no shortage of material this past year!)
http://www.medicine.mcgill.ca/epidemiology/hanley/bios601/COVIDvaccineTrials/

and received this reply...

Sorcha Pollak <spollak@irishtimes.com>
Sun 2021-01-03 5:22 AM
To: James Hanley, Dr.

Dear James,

Many thanks for your email. I agree that 105 seems like a very optimistic figure but it was the number calculated by Unicef, here's

the link to the data I used

Unicef focused on the World Population Prospects data but also combined info from civil registration stats, national household survey data and data from the UN Department of Economic and Social Affairs Population Division.

I'd say even if I'd written this story 20 years ago the number would have been significantly lower. Again though, I agree that the Unicef data does seem a bit overly optimistic.

I hope this helps.

All the very best,

Sorcha

Exercise

The footnote to the table 'ESTIMATES OF BIRTHS AND COHORT LIFE EXPECTANCY BY COUNTRY' that she used leads us to this document. Section 11 (page 28-) deals with life expectancy at birth, but doesn't mention the country-specific projections shown in the Table.

i. If you are able to locate the methods by which these country-specific projections were made, summarize them, and highlight the key assumptions.

ii. If you are not, provide your own commentary on how plausible the projections are.

iii. For a historical perspective, you might wish to visit this website, and replay (gaze down on, like Mirza did![6] Addison's *Bridge of Life* for the various countries/periods/cohorts. If your computer doesn't allow you to run the (much faster) java animations, use the `R` code instead. Comment, in particular, on the 3 'bridges' based on the data from France.

iv. A number of authors, e.g., in Spain and the USA have begun to use drops in life-expectancy to measured the impact of COVID-19. Do you think this metric in this context? What metric(s) would you recommend?

59 How large is the 'nocebo' effect?

N-of-1 Trial of a Statin, Placebo, or No Treatment to Assess Side Effects

Article, and Supplementary Appendix data scraped from Figure 1

"In patients who had discontinued statin therapy because of side effects, 90% of the symptom burden elicited by a statin challenge was also elicited by placebo."

Coverage by ScienceDaily Commentary, JAMA Medical News & Perspectives

"The original primary end-point analysis showed a *nocebo ratio* of 2.2 (95% confidence interval [CI], -62.3 to 66.7). This value was high and had a wide confidence interval because ... "

"An independent statistician therefore recommended a different analysis win which[...] This analysis showed a *nocebo ratio* of 0.90."

i. Explain in plain words what they mean by a *nocebo ratio*. (If you like, use numbers to illustrate)

ii. Explain in your own words why, despite the reasonable sample size, the first approach produced unsatisfactory results.

[6] See the presentation.

iii. The second approach produced a sensible point estimate, but no interval estimate. Using the data from Figure 1, supply one. (*deluxe version not required*)

iv. What is the broader statistical methods message from this example?

v. Comment on the last paragraph of page 6 of the Supplementary Appendix ["As there were no prior data... "]

60 Effect of Cannabidiol and Δ9-Tetrahydrocannabinol on Driving Performance

A Randomized Clinical Trial

Article, Editorial, and Supplementary Online Content

data scraped from Figure 2

IMPORTANCE Cannabis use has been associated with increased crash risk, but the effect of cannabidiol (CBD) on driving is unclear.

OBJECTIVE To determine the driving impairment caused by vaporized cannabis containing Δ9-tetrahydrocannabinol (THC) and CBD.

DESIGN, SETTING, AND PARTICIPANTS A double-blind, within-participants, randomized clinical trial was conducted at the Faculty of Psychology and Neuroscience at Maastricht University in the Netherlands between May 20, 2019, and March 27, 2020. Participants (N = 26) were healthy occasional users of cannabis.

INTERVENTIONS Participants vaporized THC-dominant, CBD-dominant, THC/CBD-equivalent, and placebo cannabis. THC and CBD doses were 13.75 mg. Order of conditions was randomized and balanced.

MAIN OUTCOMES AND MEASURES The primary end point was standard deviation of lateral position (SDLP; a measure of lane weaving) during 100 km, on-road driving tests that commenced at 40 minutes and 240 minutes after cannabis consumption. At a calibrated blood alcohol concentration (BAC) of 0.02%, SDLP was increased relative to placebo by 1.12 cm, and at a calibrated BAC of 0.05%, SDLP was increased relative to placebo by 2.4 cm.

RESULTS Among 26 randomized participants (mean [SD] age, 23.2 [2.6] years; 16 women), 22 (85%) completed all 8 driving tests. At 40

to 100 minutes following consumption, the SDLP was 18.21 cm with CBD-dominant cannabis, 20.59 cm with THC-dominant cannabis, 21.09 cm with THC/CBD-equivalent cannabis, and 18.28 cm with placebo cannabis. SDLP was significantly increased by THC-dominant cannabis (+2.33 cm [95% CI, 0.80 to 3.86]; P ¡ .001) and THC/CBD-equivalent cannabis (+2.83 cm [95% CI, 1.28 to 4.39]; P ¡ .001) but not CBD-dominant cannabis (-0.05 cm [95% CI, -1.49 to 1.39]; P ¿ .99), relative to placebo. At 240 to 300 minutes following consumption, the SDLP was 19.03 cm with CBD-dominant cannabis, 19.88 cm with THC-dominant cannabis, 20.59 cm with THC/CBD-equivalent cannabis, and 19.37 cm with placebo cannabis. The SDLP did not differ significantly in the CBD (-0.34 cm [95% CI, -1.77 to 1.10]; P ¿ .99), THC (0.51 cm [95% CI, -1.01 to 2.02]; P ¿ .99) or THC/CBD (1.22 cm [95% CI, -0.29 to 2.72]; P = .20) conditions, relative to placebo. Out of 188 test drives, 16 (8.5%) were terminated due to safety concerns.

CONCLUSIONS AND RELEVANCE In a crossover clinical trial that assessed driving performance during on-road driving tests, the SDLP following vaporized THC-dominant and THC/CBD-equivalent cannabis compared with placebo was significantly greater at 40 to 100 minutes but not 240 to 300 minutes after vaporization; there were no significant differences between CBD-dominant cannabis and placebo. However, the effect size for CBD-dominant cannabis may not have excluded clinically important impairment, and the doses tested may not represent common usage.

See also: Montreal Gazette article by Christopher Labos

From the Outcomes portion of the Methods section

"The prespecified primary end point was mean SDLP during the on-road driving test. Lateral position, which is the distance between the vehicle and the lane boundary to the left of the vehicle, was recorded by a camera mounted onto the roof of the vehicle and sampled continuously at 4 Hz. Measurements of lateral position over the time of the driving test were aver- aged to yield the mean lateral position, and standard deviation was calculated to determine the mean SDLP. Larger numbers indicate greater variability (ie, reduced stability) in lane positioning. A 2.4-cm drug vs placebo increase in SDLP is typical of a driver with a blood alcohol concentration (BAC) of 0.05% and is thought to indicate the lower limit of clinically relevant driving impairment"

Exercises

i. 4 Hz means 4 samples a second, so, for each person, driving an hour in each cannabis condition, that is n=4 x 60 x 60 = 14400 measurements that go into the mean distance from the white line, and also into the SD.

   • If someone drives at a mean lateral position of 84 cm from the white line (95% CI, 80.01 to 88.82) in the placebo condition, and has a SD of 18 cm , what is the probability of going over the white line? What if the SD is increased to 25 cm?

   • State the assumption(s) you would need to calculate this.

ii. Does the large no. of measurements make the distribution of the 14400 individuals measurements Gaussian? Why/why not?

iii. What DOES the CLT say? Is it relevant here?

iv. SD's are typically harder to estimate, and take bigger n's than means. How precise is a SD based on 14400 measurements?

v. "Statistical Analysis
Sample size was determined by power calculation using the effect size obtained in a previous study of dronabinol (10-20 mg THC) on SDLP during on-road driving. This indicated that 20 participants were needed to detect an equivalent effect (Cohen f = 0.62; ΔSDLP = 1.94 cm; approximately 0.04% BAC) with 95% power."

   • Comment on how the effect size was chosen, and whether it is an appropriate way to specify an effect size.

   • This use of Cohen's method to arrive at a sample size is common in the social sciences. Use a method that is more mainstream in medical and epidemiological research.

vi. (See dataset) Why would it be inappropriate to use a 2-independent-samples t-test to compare the CBD v.s Placebo data-points in panel A of Figure 2?

vii. Why is it appropriate to use a 1-sample t-test to analyze the ΔSDLP data-points in the first column of panel B of Figure 2?

viii. We don't have sufficient information to figure out, in Figure 2 A, which person in the placebo column is which person in the other columns. But, we can get a rough idea, by comparing e.g., the variances of (a) the SDLP in Placebo and (b) the SDLP in CBD in panel A, with the variances of (c) the ΔSDLP in 'CBD' in panel B. Thus, we should be able to say that in panel A we would expect the person at the top of the placebo column to be a certain distance down from the top in the CBD column, and likewise that we would expect the person at the bottom of the placebo column in A to be a certain distance up from the bottom in the CBD column.

   • Using these 3 variances, derive the correlation between the SDLP in Placebo and the SDLP in CBD. From this, calculate the predicted (fitted) SDLP in CBD for each SDLP in Placebo, and vice versa.

ix. Each subject was tested in all (or almost all) 8 experimental conditions. In the psychology field where these authors work, they would call this a 'repeated measures design'.

A very good reference textbook is Winer's *Statistical principles in experimental design*. McGill library has the print version and limited access to the 1991 edition by by Winer, Brown, and Michels.

We in epidemiology and statistics tend to use the repeated measures term for any (even non-experimental) situation where measurements are taken over time, but with nothing else 'changed'. They refer to the 8 sets of measurements as 8 (within-subject) "conditions" – a wise choice. And they, unlike us, tend to avoid the word 'groups' which suggests different people in different groups.

To JH, the randomization that Labos highlights is less of an added benefit than the within-person aspect – which can yield considerable narrowing of the noise. We suspect that the driver near the top of the first column in Fig2A is also near the top in the other columns. Another reader in JH's household didn't catch from his article the fact that the design was a full crossover one, i.e, that everything was within-subject. They thought that 26 divided by 4 was a very small number.

• Quantify the gain from using the crossover design. In other words, if you want a specified precision/power, how many subjects would you need in the within-person (serial) design relative to the equivalent number in the 'independent groups' (parallel) design? Begin by limiting your calculations to the placebo vs. CBD contrast.

• What implications does the difference in designs have in terms of budget, if the variable costs are mainly (a) recruitment (b) on-road sessions?

x. Comment on Labos' suggestion that their doses were small and that they should test higher doses.

Figure 2. The Standard Deviation of Lateral Position During On-Road Driving Tests

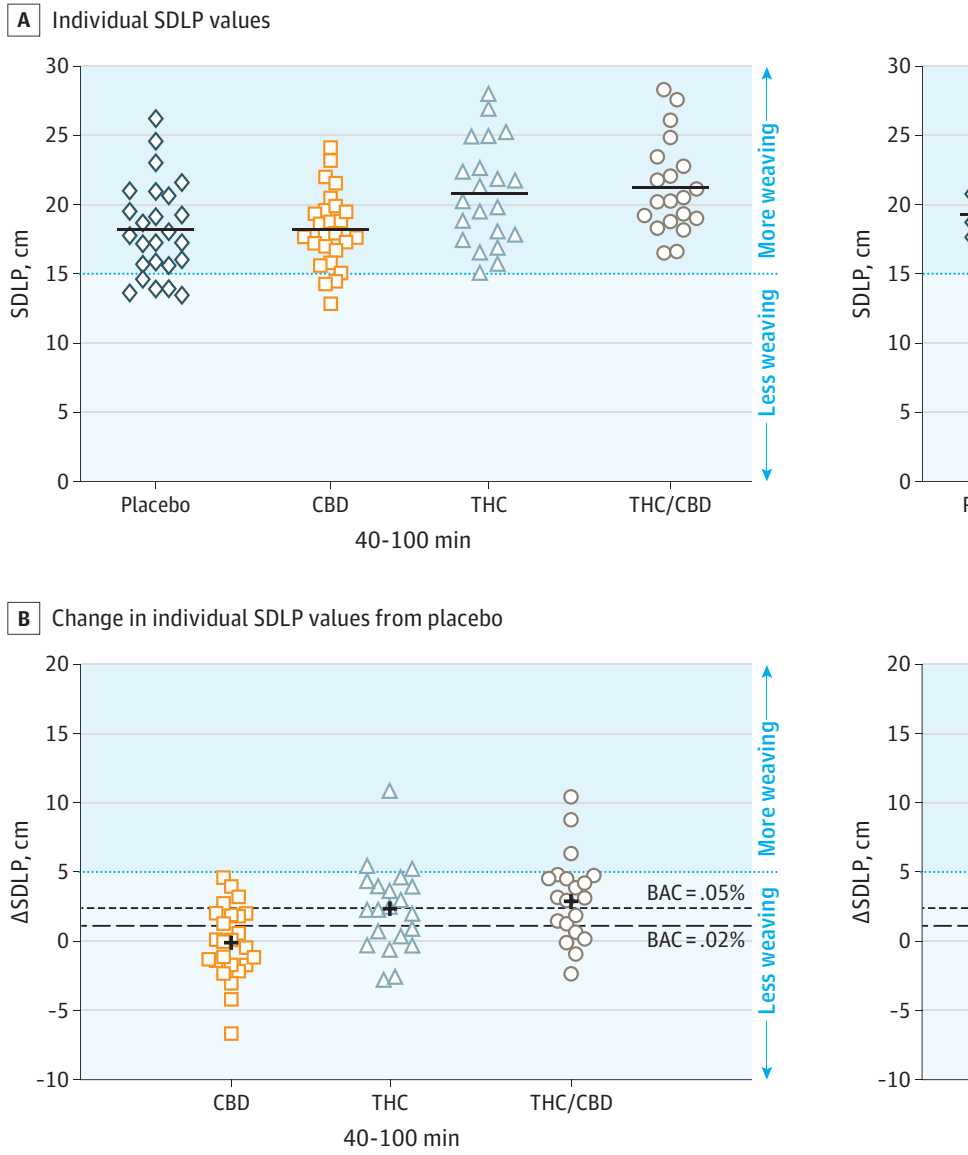

The x-axes indicate minutes postvaporization. Higher values on the y-axes indicate more weaving vs less weaving for lower values. A, The horizontal black bars indicate the mean standard deviation of lateral position (SDLP) in each condition. B, The dashed lines indicate the mean SDLP increase associated with a blood alcoh mean change THC, Δ^9-tetra

.

61 Estimating Efficacy of Moderna Vaccine against COVID-19

The protocols, announcements, journal reports, editorial, and links to the video of and material presented to/discussed at the FDA meetings can be found here.

See panel B ('Modified Intention-to-Treat Analysis') in Figure 3 of the NEJM report of the trial.

> Vaccine efficacy was defined as the percentage reduction in the hazard ratio for the primary end point (mRNA-1273 vs. placebo). A stratified Cox proportional hazards model was used to assess the vaccine efficacy of mRNA-1273 as compared with placebo in terms of the percentage hazard reduction. (Details regarding the analysis of vaccine efficacy are provided in the Methods section of the Supplementary Appendix.)

JH has been able to use the numbers at risk at selected timepoints shown at the bottom of the panel, along with the ('de-cumulated') cumulative event rate function to (almost perfectly) reconstruct the (arm-specific) daily numbers at risk, and daily numbers of cases. The daily data are available here, where 'P' and 'V' denote the Placebo and Vaccine arms respectively.

i. Check that the daily numbers of cases add up to the total numbers reported at the bottom of the panel.

ii. Because of the interpolation JH had to employ, there are a few small anomalies in the numbers at risk on the days in between the selected days (check that they match on the selected days). Identify a few of these small anomalies (they are more common at the front end).

iii. > After day 1 and through November 25, 2020, a total of 269 Covid-19 cases were identified, with an incidence of 79.8 cases per 1000 person-years (95% confidence interval [CI], 70.5 to 89.9) among participants in the placebo group with no evidence of previous SARS-CoV-2 infection. [from first paragraph of Results, EFFICACY; these numbers are repeated in panel B.]

   Use the data set to try to reproduce this incidence and confidence interval.

iv. 'Expand' this 'daily' dataset into a data frame with just over 29,000 'rows', one per participant, and with columns representing (i) a unique participant ID number, (ii) which arm the participant was in (iii) when (which day) follow-up of that participant ended and (iv) whether the follow-up was terminated by the diagnosis of COVID-19 (event=1); or ('administratively') when the data-file was locked, or by death of an unrelated cause, or a (genuine) 'loss' to follow-up (event=0). (Note: the few anomalies in part i. will pose some 'book-keeping' problems, so you can either try to remedy them, or just ignore them, since the moving of a few censoring times by a few days makes so little difference.

v. Then, try to reproduce the Vaccine Efficacy and 95% CI shown within panel B, by applying the `coxph` function in the `R survfit` package (or the corresponding function in you favourite statistical package) to this individual-level dataset. [The authors mentioned a stratified Cox model, but we don't have access to the information on the strata.]

vi. Since on some days there are 'ties' (two or more events on same day), repeat the calculations with the non-default methods of handling ties.

vii. Since on some days there are 'ties' (two or more events on same day), repeat the calculations with the non-default methods of handling ties.

viii. We spoke in class about the factors that influenced the hazard function in the placebo arm, and why Cox avoided the challenge of fitting it by (instead) matching on time, and assuming a constant-over-time efficacy. In this particular application, there is very little interest in this very 'particularistic' hazard function, but on other applications (e.g., in say calculating the 5- or 10-year Framingham risk), one needs to estimate it – indeed it is the starting point (the 'corner' in Clayton & Hills terminology) in such prognostic-probability models.

Among the many 'calculators' on this important website are the CARDIOVASCULAR MEDICINE CALCULATORS. Just look at the 'Risk' equation shown in the coloured box above the 'Input' heading.

As we say in this piece, few end-users are aware that is possibility exists, and how 'compact' and 'app-friendly' it is. Some heuristics behind how this (non-parametric) function is 'fitted' (estimated) are given in this expository article.

Exercise: By applying the appropriate 'post-fit' function from your favourite statistical package (see Table A1 in this piece), derive the $\widehat{S_0(t)}$ survival function for the placebo arm, and the $\widehat{S_1(t)}$ one for the Vaccine arm, and plot their complements (ie. the 'Risk' or 'cumulative Event Rate' functions) as overlays on the graph in Figure 3b.

ix. Fit a 'matched by day' PH model to the reduced (day by day summary) data file with 118 rows. Do so by conditioning on (treating as fixed) the total number (C) of new diagnoses (Cases) on each day, and treating the split of the C_j cases on day j into C_{j0} cases in the P arm and C_{j1} in the V arm as the realization of a Binomial random variable with 'denominator' C_j and with the probability [that a case is a *vaccinated* case] :

$$\pi_j = \frac{n_{j1} \times (100 - VE)}{n_{j1} \times (100 - VE) + n_{j0} \times 100},$$

or odds

$$\Omega_j = \frac{n_{j1} \times (100 - VE)}{n_{j0} \times 100} = \frac{100 - VE}{100} \times \frac{n_{j1}}{n_{j0}},$$

or log odds

$$\log \Omega_j = LOGIT_j = \log \left[\frac{100 - VE}{100} \right] + \log \left[\frac{n_{j1}}{n_{j0}} \right].$$

[*Hint*: The second term in this equation is an 'offset'. The code for fitting a simpler (1-row) example of this models can be found under 'Binomial regression model for the Binomial split' using the data from the Salk Polio VaccineTrial in this item. In the present example, we have as many binomials as we have riskets (days with at least one COVID-19 case).]

x. Repeat, but by fitting the '2 independent Poisson RV's' shown earlier in this item (here you will need to treat 'day' as a factor.)

xi. Now, *pretend* that you had the daily numbers of cases in each arm, but that Moderna had not disclosed the randomization ratio (or the numbers at risk that day). They might even have used a different randomization ratio each day (that could happen, for example, if the vaccine supply was limited, or they were using adaptive randomization).

Suppose that, for each day j, Moderna would sample the list of *participants at risk on day j* and let you know which selected ones were in the V and P arms respectively. For each day you specify, Moderna will charge you \$10 (per selected participant, per day!) for telling you which arm each randomly selected person was in. Your total budget is \$10,000.

Exercise: Let JH (Moderna's agent) know how many you wish to have randomly selected each day (it doesn't have to be the same number each day). For each day, JH will tell you how many of those selected were in the V and P arms. Moderna will not tell him/you if any of those selected that day were selected on an earlier day, and will not tell if these participants became a case on the day in question. All you get for (say) day 29 is that of the (say) 3 persons you asked to have selected from those at risk on day 29, say 2 were in the V arm and 1 was in the P arm – and they will charge you \$30 just for these 3 pieces of information.

Then, using your (free of charge) numerator series, and your (it cost you your entire PhD budget!) denominator series, estimate the VE, and supply a 95% CI for it.

xii. Suppose that (before you went ahead with your plan) a rich relative of yours took pity on you and donated an additional \$20,000, so you have \$30,000 in all. Redo the order to JH/Moderna, and re-do the analysis using the denominator-series that this bought you.

xiii. Suppose that (before you went ahead with your plan) CIHR reduced you budget from \$10,000 to \$2,500, and that you don't have any rich relative/supervisor. Redo the order, and re-do the analysis with the reduced denominator-series.

62 Hockey Epidemiology

Refer to the 'slides' from this 2014 presentation in the department's lecture series for the public, to commemorate its 50th anniversary, and in particular to the last 5 slides, beginning with how to beat Canadien's goalie Patrick Roy. Here is a separate hi-resolution version of the full 5-page coverage that JH has since found.

hi-resolution version.

i. What is wrong with the inference 'Pour battre Roy, mieux vaut langer bas...'?

ii. What could you do to rectify the study?

iii. Suppose time is limited, and that your reserach assistant only got to (randomly) sample just 60 of the 600 shots, and to classify them as shown in the table. *i.e. ignore the data on the full 600.* (Note that you did not ask your RA to record whether any of the 60 shots resulted in goals; nor is there any need to!). Limiting yourself to the Low (index category) versus High (reference category)

   (a) calculate the (empirical) *difference* in the success probabilities, along with a 95% CI. If, as is often the situation, the sampling fraction were unknown, could one calculate the absolute difference?

   (b) calculate the (empirical) *ratio* of the success probabilities, along with a 95% CI. If, as is often the situation, the sampling fraction were unknown, could one still calculate the ratio?

iv. Repeat the calculations with the full number of shots (i.e., the 297 and the 202).

v. What are some of the messages this example holds for

   (a) terminology, i.e., use of the term '*denominator'* series rather than '*control*' series, and which is more natural and self-expalanatory ?

   (b) whether it complicates or simplifies matters if the two different teams working on the numerator series and the denominators series respectively need to co-ordinate their efforts. i.e., should they adopt a 'don't ask, don't tell' approach? i.e., the 'numerator team' doesn't tell the denominator team which shots resulted in goals? i.e. the 'denominator' team can see enough about the shot to classify its target, but not whether it resulted in a goal. In other words, shots that result in goals are eligible to be selected, just as shots that do not are (this is different from the traditional 'case-control' approach).

   (c) whether we could completely eliminate the odds ratio as an estimand (a natural parameter, to be estimated), and just call the estimator (ad/bc) a *cross-product-ratio* a ratio that '*looks like* an empirical odds ratio, but is *nothing more (in this case) that a ratio of success probabilities using estimated denominators*'? In other words, if the estimand is a ratio of probabilities, there is no need to explain to the ardent hockey aficianado what an *odds ratio* is. AND, there is no need to get into the 'rare disease' or 'rare goal' assumption. This (partial) denominator series would also hold if at issue were goals on penalty shots (where the success probability is much higher)!

   (d) moving on from the old-fashioned but still-prevailing practice of teaching that in 'case-control' studies we 'compare cases with controls' using 'exposure odds'? [instead of comparing rates in the index and reference categories of the 'exposure'?

   For more on these concepts/principles, see this ignored classic and this update, and reference 16 in it. **The modern (case-base) version is discussed** starting at minute 50 (and again at 1hr 15min) in the full interview here.

63 Birth defects – and Other Reproductive – Epidemiology

Refer to these 'slides' from this 2014 presentation in the department's lecture series for the public, to commemorate its 50th anniversary.

i. What is the name usually given to this type of 'study' reported under the title 'Congenital cataract following German measles in the mother'?

   For more on the author, see here. The congenital deafness produced by the in-utero rubella exposure was discovered by this multi-talented author, using routinely collected census data.

ii. What type of 'study' is reported under the title 'Thalidomide and Congenital Abnormalities'?

   Would you be able to calculate CI's for the possible effect measures?

   For more on the author, see here.

iii. What type of 'study' is reported under the title 'Folic Acid Metabolism and Human Embryopathy'?

   Would you be able to calculate CI's for the possible effect measures?

   For more on the the fallout from this study, see here.

iv. What type of 'study' is reported under the title 'Adenocarcinoma of the vagina: association of maternal stilbesterol therapy with tumor appearance in young women'?

   Are you surprised at the exposure histories in the 32 'controls'?

   Would you be able to calculate point and interval estimates for the effect measures for each of the eight (separate) exposures reported on?

   For more on the story, see here or Google 'Herbst stilbesterol' or visit this site.