

Advanced Statistical Methods for Observational Studies

Spring 2021 Remote Asynchronous Instruction

Instructor

David Rogosa (usually) Sequoia Hall 224, rag@stanford.edu

Note: Prof. Baiocchi, the lead instructor in the course since inception, is off-duty and will not be participating during Spring 2021.

His lectures from the prior Remote Asynchronous offering of this course in Spring 2020 will still form the core of the 10 weeks of instruction, along with the [Computing Corner](#) materials and lectures from Prof. Rogosa.

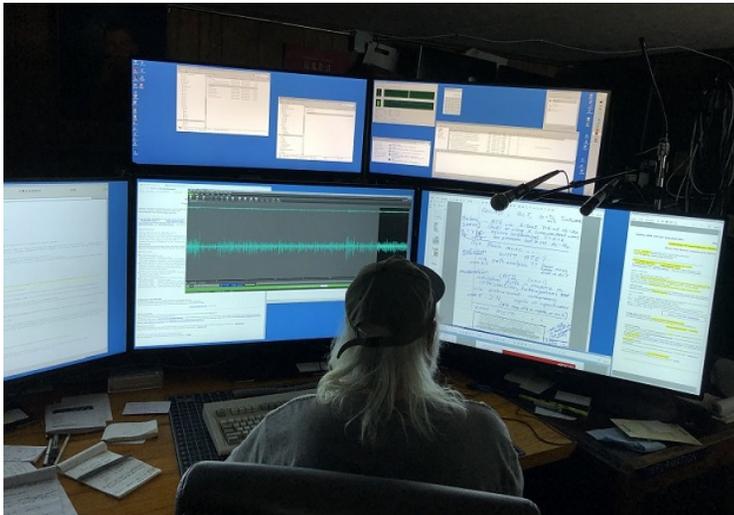
Course web page: <http://rogosateaching.com/somgen290/>

To see full course materials from Spring 2020 [go here](#)

Course Welcome and Logistics (first day stuff, posted in March, call it Week0)

[Lecture slides, week 0](#) (pdf)

[Audio companion, week 0](#)

**Registrar's Information**

EPI 292: Advanced Statistical Methods for Observational Studies (STATS 266, EDUC 260B, CHPR 266) 2-3 units

Design principles and statistical methods for observational studies.

Topics include: matching methods, sensitivity analysis, instrumental variables.

3 unit registration requires a small project and presentation. Computing is in R.

Pre-requisites: EPI 261 and 262 or STAT 209, or equivalent. See <http://rogosateaching.com/somgen290/>

Terms: Spr | Units: 2-3 | Grading: Medical Option (Med-Ltr-CR/NC)

Abbreviated Course Outline

- Week 1 - Course Introduction; Matching Methods Part 1 (intro and theory)
- Week 2 - Matching Methods Part 2 (implementation); Potential Outcomes and Study Design
- Week 3 - Full matching, Inclusion and Exclusion, and Defining Treatment Effects
- Week 4 - Models for Observational Studies and Inverse Probability Weighting
- Week 5 - Randomized Experiments and Design Sensitivity
- Week 6 - Augmenting the Primary Study: Second Outcomes, Known Nulls, Coherence, Multiple Contrast Groups, and Thick Description
- Week 7 - Alternative Designs: Discontinuity Designs and Case-Noncase Studies
- Week 8 - RCT designs with Instrumental Variables
- Week 9 - Observational Studies with Instrumental Variables
- Week 10 - Mendelian randomizations and synthetic cohorts

Note: This course was structured before the new world order of March 2020 as one two-hour class meeting per week, traditionally Monday afternoon.. Posting of weekly materials this quarter (web resources, lectures, audio) is intended for Mondays.

[Lectures, Course Files, and Readings](#)

this page is where course content resides starting 3/29

In this course students will:

- (1) Learn to identify key statistical issues in observational studies and methods and study designs to address issues of confounding.
- (2) Become proficient with advanced statistical methods for observational studies: methods for missing data, matching based inference, sensitivity analysis, propensity score methods, instrumental variables. You should know which methods are useful in different situations, and which conditions have to be checked for the method to be applicable.
- (3) Be able to perform detailed data analyses on a variety data using the statistical computation environment R. You should be able to implement all the methods presented in this course.

Textbooks**Required**

Design of Observational Studies, Paul Rosenbaum, 1st Edition (Springer) Available online: [Stanford access](#)

Additional Resources

Causal Inference, Miguel Hernan & [Jamie Robins](#) Available online: <http://www.hsph.harvard.edu/miguel-herman/causal-inference-book/>

Causal Inference in Statistics, Social and Biomedical Sciences: An Introduction, [Guido Imbens and Don Rubin](#), 1st Edition (Cambridge University Press) [Stanford access](#)

Computation

We will make extensive use of the statistical computation and [programming environment R](#). It is free and open source, and it has become the de facto standard for statistical analyses in most

areas of academia and industry. For references and software: [The R Project for Statistical Computing](#)

Closest download mirror is [Berkeley](#). If Berkeley is offline, choose a mirror from the main R page (first link).

The [CRAN Task Views](#) provide an organization and overview of the many R packages.

One specialized resource this course will draw from is [Paul Rosenbaum's software page](#).

Many students may also use or be interested in an editor for R (beyond emacs or command line); a popular option is [RStudio](#), [download at](#) .

Basic R References.

Using R for Introductory Statistics, [Verzani](#), 1st Edition (Chapman & Hall/CRC)

online Verzani book resources: version of text available from [John Verzani's page](#) . [alternate version, single.pdf](#) [Using R R-package](#)

Introductory Statistics with R, [Dalgaard](#), 2nd Edition (Springer); also see `ISwR` package for data and functions

A [handbook of statistical analyses using R](#) (second edition). Brian Everitt, Torsten Hothorn CRC Press, [Index of book chapters](#) [Stanford access](#) Data sets etc [Package 'HSAUR2'](#)

Course Components, Student Work

Homeworks: Review Questions will be posted each week with solutions. These problems are for your own learning and will not be collected or graded.

Problem Sets: During this special Spring Quarter there will be only one (instead of usual two), take-home problem set(s). The problem set will be the basis for grading the two-unit enrollment.

No collaboration or external assistance is permitted.

Course Problem Set

Third unit enrollment, Project/Presentation. If you are enrolled for three units credit, the additional unit requirement to complete a small project and in the past give us a short presentation (approximately 15 minutes with handout). Form of these 'presentations' this special quarter will be arranged.

Advanced Statistical Methods for Observational Studies



LECTURE 00

website



- Weekly – before Mondays at 5pm Pacific – slide decks and accompanying audio files will be uploaded to the site to the “Course Readings, Files and Examples” section.

STAT266/CHPR266/HRP292-- Course Files, Readings, Examples

Week 1--Course Introduction; Matching Methods part 1 (intro and theory)

From 2020:

Lecture Topics Lecture 1 [slide deck \(companion audio part 1\)](#) ([companion audio part 2](#))

1. Course outline and logistics
2. A matched observational study (DOS, Chap 7)
3. Study design versus inference
4. Basic tools of multivariate matching (DOS, Secs 8.1-8.4)

Text Readings

Rosenbaum DOS: Chapters 7 and 8 (8.1-8.4)

Additional Resources

Observational Studies according to Donald B. Rubin

[For objective causal inference, design trumps analysis](#) Annals of Applied Statistics, Volume 2, Number 3 (2008), 808-840. [Rubin talk](#).
Another Rubin overview of matching: [Matching Methods for Causal Inference](#) Stuart, E.A. and Rubin, D.B. (2007). Best Practices in Quasi-Experimental Designs: Matching methods for causal inference. Chapter 11 (pp. 155-176) in Best Practices in Quantitative Social Science. J. Osborne (Ed.). Thousand Oaks, CA: Sage Publications.

Computing Corner: [see 2019 course website](#) for 2019 Computing Corner Week 1 pdf slides [2020 audio companion](#)

Week 1 Review Questions

From Computing Corner

1. In Week 1 Computing Corner with the Lalonde data (effect of job training on earnings), we started out (see R-session) by showing the ubiquitous [epidemiology to economics] analysis for observational data of an analysis of covariance, aka tossing the treatment variable and all the confounders into a regression equation predicting outcome and hoping for the best (c.f 2016 Week 1 *in the news* analyses: mom fish consumption on child cognition). The statement made in class (technical details week 1 stat209) is that regression does not “control” for confounders; instead the coefficient of treatment (putative causal effect) is obtained from a straight-line regression of outcome on the residuals from a prediction of treatment by all the other predictors in the regression. Demonstrate that equivalence using the ancova in CC1.

[Solution for Review Question 1](#)

2. RQ1 uses the Week 1 Computing Corner Lalonde data (effect of job training on earnings) analysis of covariance: tossing the treatment variable and all the confounders into a regression equation predicting outcome and hoping for the best. Compare that ancova with an ancova that uses just the significant predictors of re78. Also compare with an ancova which uses the single available covariate/confounder having the highest correlation with outcome. Are these analyses consistent?

[Solution for Review Question 2](#)

EPI292/STAT266/CHPR266/EDUC260B-- Lectures, Course Files, and Readings

Week 1--Course Introduction; Matching Methods part 1 (intro and theory)

Lecture Topics Lecture 1 [slide deck \(companion audio part 1\)](#) ([companion audio part 2](#))

1. Course outline and logistics
2. A matched observational study (DOS, Chap 7)
3. Study design versus inference
4. Basic tools of multivariate matching (DOS, Secs 8.1-8.4)

Text Readings

Rosenbaum DOS: Chapters 7 and 8 (8.1-8.4)

Additional Resources

Observational Studies according to Donald B. Rubin

[For objective causal inference, design trumps analysis](#) Annals of Applied Statistics, Volume 2, Number 3 (2008), 808-840. [Rubin talk](#). Another Rubin overview of matching: [Matching Methods for Causal Inference](#) Stuart, E.A. and Rubin, D.B. (2007). Best Practices in Quasi-Experimental Designs: Matching methods for causal inference. Chapter 11 (pp. 155-176) in Best Practices in Quantitative Social Science. J. Osborne (Ed.). Thousand Oaks, CA: Sage Publications.

Computing Corner: Extended Data Analysis Examples

Lalonde NSW data (DOS sec 2.1). Subclassification/Stratification and Full matching. [pdf slides for CCI](#) [2021 audio companion](#)

Week 1 [handout](#) Lalonde NSW data

[Rogosa R-session \(using R 3.3.3\)](#) [4/1/18 redo in R 3.4.4](#) (sparse)

[2019 lalonde Matchit: full matching, balance with cobalt, love, plot and bal.tab](#)

[2019 lalonde optmatch: fullmatch with outcome analysis](#)

Resources:

MatchIt provides a wrapper that can call optmatch or Sekhon's genetic matching

MatchIt: [Nonparametric Preprocessing for Parametric Casual Inference](#) Daniel Ho, Kosuke Imai, Gary King, Elizabeth Stuart

[MatchIt vignette](#)

JSS May 2011 exposition: [MatchIt: Nonparametric Preprocessing for Parametric Causal Inference](#)

Cobalt: [Using cobalt with Other Preprocessing Packages](#) [Covariate Balance Tables and Plots: A Guide to the cobalt Package](#)

Optmatch: Ben Hansen (local hero) [optmatch manual](#) [R News Oct 2007](#)

[optmatch/fullmatch vignette](#) [optmatch another version](#) another good tutorial [optmatch Functions for Optimal Matching](#)

Week 1 Review Questions

From Computing Corner

1. In Week 1 Computing Corner with the Lalonde data (effect of job training on earnings), we started out (see R-session) by showing the ubiquitous [epidemiology to economics] analysis for observational data of an analysis of covariance, aka tossing the treatment variable and all the confounders into a regression equation predicting outcome and hoping for the best (c.f 2016 Week 1 in the news analyses: mom fish consumption on child cognition). The statement made in class (technical details week 1 stat209) is that regression does not "control" for confounders; instead the coefficient of treatment (putative causal effect) is obtained from a straight-line regression of outcome on the residuals from a prediction of treatment by all the other predictors in the regression. Demonstrate that equivalence using the ancova in CCI.

[Solution for Review Question 1](#)

2. RQ1 uses the Week 1 Computing Corner Lalonde data (effect of job training on earnings) analysis of covariance: tossing the treatment variable and all the confounders into a regression equation predicting outcome and hoping for the best. Compare that ancova with an ancova the uses just the significant predictors of re78. Also compare with an ancova which uses the single available covariate/confounder having the highest correlation with outcome. Are these analyses consistent?

[Solution for Review Question 2](#)

From Lecture

3. We will be working a lot with matching based techniques. One of the best thinkers/writers on the topic of matching is Elizabeth Stuart from Johns Hopkins. For this problem, take a look at her paper: ["Matching Methods for Causal Inference: A Review and a Look Forward."](#) In lecture 01 you were introduced to "balance tables" (a.k.a. "Table 1") which summarizes the covariate distribution of the observations. A handful of questions: (a) as concisely as possible, state why we focus on balance assessments as part of our argumentation when attempting to perform causal inference, (b) in addition to a balance table, name other tools used to report balance, (c) why do we use standardized mean differences instead of p-values to assess balance when assessing the quality of a match design?, and (d) why is it kinda weird to use a p-value of the covariates in a randomized trial to assess balance?

[Solution for Review Question 3](#)

4. In lecture 1 we quickly outlined some of the big challenges to causal inference when using observational data (see slide 41, "There should be strong effort to show the two groups are similar..."). These challenges include: inclusion/exclusion of observations, observational units that may be completely missing (censored, survival bias), missing data, imbalances in observed data, and imbalances in unobserved data. We'll address each of these at different points in the course. But let's focus on the decision to include/exclude observations. What we're doing when matching -- i.e., removing observations that do not have adequate counterparts in the contrast group -- may seem a bit subversive. The intuition is: why "throw away" data? I think there are two reasons people worry about "throwing away data." First, it seems like limiting the kinds of observations in our study we may be losing the ability to generalize our conclusions to a wider swath of the population. The counter to that is: yes, we are trading off the ability to generalize (i.e., external validity) for the ability to make stronger claims about a candidate causal effect (i.e., internal validity). The second concern is that it seems like more data is better. Formulate a response to this concern. (Note: OMG, this question seems so nebulous. Yup. That's how this works; you're playing Big Kid academics now. We made sure to mention this argument during lecture 01, so you know it. It's a common statistical argument nowadays. If you want to read your way out of this one... [here's a good paper.](#))

[Solution for Review Question 4](#)

From Computing Corner

5. Exercise in pair matching. In DOS Sec 2.1, Rosenbaum works with the randomized experiment data from NSW. In Week 1,2 Computing Corner we used the constructed observational study version of these data. Use the observational study data to do a version of the 1:1 matching in DOS section 2.1. Compare the balance improvement achieved from nearest neighbor matching with the full matching results in Computing Corner Week 1,2.

[Solution for Review Question 5](#)

6. For the fullmatch analysis done in the Lalonde class presentation weeks 1 and 2, the outcome comparison was carried out using lmer to average the treatment effects over the 104 subclasses. A hand-wavy analogy to the paired t-test here would be to use the mean difference within each subclass. Show that (because some of the subclasses are large) this simplified analysis doesn't well replicate the lmer results.

[Solution for Review Question 6](#)

7. optmatch package, fullmatch, lalonde.

MatchIt uses the optmatch package fullmatch command for its "full" option, as used in the class example. Using the raw optmatch (without the matchit wrapper) allows additional specifications and controls for the full or optimal matching.

For lalonde data try out optmatch fullmatching and compare results for subclasses and balance with the class example using optmatch through MatchIt.

[Solution for Review Question 7](#)

Week 2-- Matching Methods Part 2 (implementation); Potential Outcomes and Study Design

expectations



- In the spring of 2021 there will be only one problem set and it will be due Week 9 of the quarter, ~~by 5pm on Monday, May 24.~~
- There are weekly practice problems, with solutions, available in the “Course Readings, Files and Examples” section.

expectations



- In the spring of 2021, you're going to have access to our weekly lectures. You are in charge of pacing. For example, if you are eager to jump ahead then you can go back to older versions of the website and look at lectures.
- The only “hard” deadlines are the problem set and (if you're taking the class for 3 credits) a presentation.
- Please! Please! Please feel free to reach out to ~~us~~ to talk and set up time to go over ideas/questions/your-projects. It's super weird for us to not get to know you in person.

OCCASIONAL NOTES

Chocolate Consumption, Cognitive Function,
and Nobel Laureates

Franz H. Messerli, M.D.

Dietary flavonoids, abundant in plant-based foods, have been shown to improve cognitive function. Specifically, a reduction in the risk of dementia, enhanced performance on some cognitive tests, and improved cognitive function in elderly patients with mild impairment have been associated with a regular intake of flavonoids.^{1,2} A subclass of flavonoids called flavanols, which are widely present in cocoa, green tea, red wine, and some fruits, seems to be effective in slowing down or even reversing the reductions in cognitive performance that occur with aging. Dietary flavanols have also been shown to improve endothelial function and to lower blood pressure by causing vasodilation in the peripheral vasculature and in the brain.^{3,4} Improved cognitive performance with the administration of a cocoa polyphenolic extract has even been reported in aged Wistar-Unilever rats.⁵

Since chocolate consumption could hypothetically improve cognitive function not only in individuals but also in whole populations, I wondered whether there would be a correlation between a country's level of chocolate consumption and its population's cognitive function. To my knowledge, no data on overall national cognitive function are publicly available. Conceivably, however, the total number of Nobel laureates per capita could serve as a surrogate end point reflecting the proportion with superior cognitive function and thereby give us some measure of the overall cognitive function of a given country.

METHODS

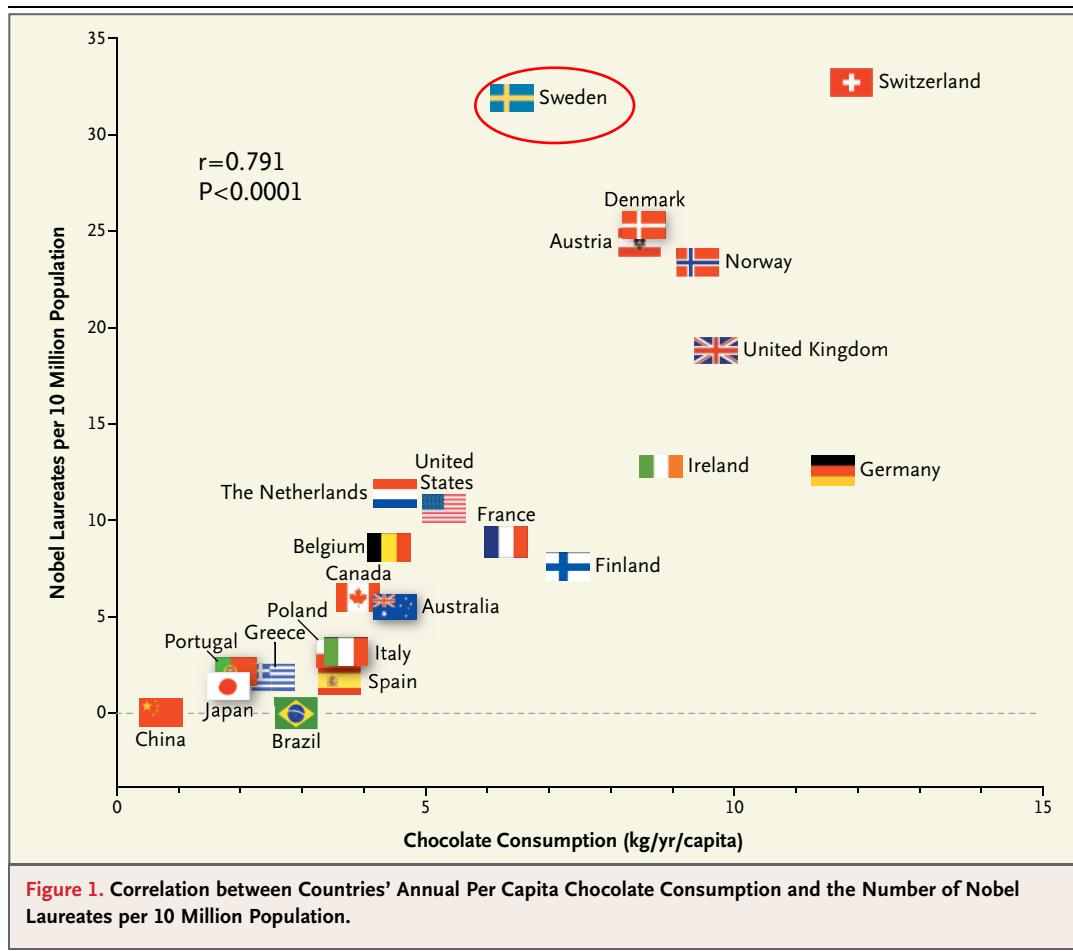
A list of countries ranked in terms of Nobel laureates per capita was downloaded from Wikipedia (http://en.wikipedia.org/wiki/List_of_countries_by_Nobel_laureates_per_capita). Be-

cause the population of a country is substantially higher than its number of Nobel laureates, the numbers had to be multiplied by 10 million. Thus, the numbers must be read as the number of Nobel laureates for every 10 million persons in a given country.

All Nobel Prizes that were awarded through October 10, 2011, were included. Data on per capita yearly chocolate consumption in 22 countries was obtained from Chocosuisse (www.chocosuisse.ch/web/chocosuisse/en/home), Theobroma-cacao (www.theobroma-cacao.de/wissen/wirtschaft/international/konsum), and Caobisco (www.caobisco.com/page.asp?p=213). Data were available from 2011 for 1 country (Switzerland), from 2010 for 15 countries, from 2004 for 5 countries, and from 2002 for 1 country (China).

RESULTS

There was a close, significant linear correlation ($r=0.791$, $P<0.0001$) between chocolate consumption per capita and the number of Nobel laureates per 10 million persons in a total of 23 countries (Fig. 1). When recalculated with the exclusion of Sweden, the correlation coefficient increased to 0.862. Switzerland was the top performer in terms of both the number of Nobel laureates and chocolate consumption. The slope of the regression line allows us to estimate that it would take about 0.4 kg of chocolate per capita per year to increase the number of Nobel laureates in a given country by 1. For the United States, that would amount to 125 million kg per year. The minimally effective chocolate dose seems to hover around 2 kg per year, and the dose-response curve reveals no apparent ceiling on the number of Nobel laureates at the highest chocolate-dose level of 11 kg per year.



DISCUSSION

The principal finding of this study is a surprisingly powerful correlation between chocolate intake per capita and the number of Nobel laureates in various countries. Of course, a correlation between X and Y does not prove causation but indicates that either X influences Y, Y influences X, or X and Y are influenced by a common underlying mechanism. However, since chocolate consumption has been documented to improve cognitive function, it seems most likely that in a dose-dependent way, chocolate intake provides the abundant fertile ground needed for the sprouting of Nobel laureates. Obviously, these findings are hypothesis-generating only and will have to be tested in a prospective, randomized trial.

The only possible outlier in Figure 1 seems to be Sweden. Given its per capita chocolate consumption of 6.4 kg per year, we would predict that Sweden should have produced a total of

about 14 Nobel laureates, yet we observe 32. Considering that in this instance the observed number exceeds the expected number by a factor of more than 2, one cannot quite escape the notion that either the Nobel Committee in Stockholm has some inherent patriotic bias when assessing the candidates for these awards or, perhaps, that the Swedes are particularly sensitive to chocolate, and even minuscule amounts greatly enhance their cognition.

A second hypothesis, reverse causation — that is, that enhanced cognitive performance could stimulate countrywide chocolate consumption — must also be considered. It is conceivable that persons with superior cognitive function (i.e., the cognoscenti) are more aware of the health benefits of the flavanols in dark chocolate and are therefore prone to increasing their consumption. That receiving the Nobel Prize would in itself increase chocolate intake countrywide seems unlikely, although perhaps celebratory events associated with this unique

honor may trigger a widespread but most likely transient increase.

Finally, as to a third hypothesis, it is difficult to identify a plausible common denominator that could possibly drive both chocolate consumption and the number of Nobel laureates over many years. Differences in socioeconomic status from country to country and geographic and climatic factors may play some role, but they fall short of fully explaining the close correlation observed.

STUDY LIMITATIONS

The present data are based on country averages, and the specific chocolate intake of individual Nobel laureates of the past and present remains unknown. The cumulative dose of chocolate that is needed to sufficiently increase the odds of being asked to travel to Stockholm is uncertain. This research is evolving, since both the number of Nobel laureates and chocolate consumption are time-dependent variables and change from year to year.

CONCLUSIONS

Chocolate consumption enhances cognitive function, which is a sine qua non for winning the

Nobel Prize, and it closely correlates with the number of Nobel laureates in each country. It remains to be determined whether the consumption of chocolate is the underlying mechanism for the observed association with improved cognitive function.

Dr. Messerli reports regular daily chocolate consumption, mostly but not exclusively in the form of Lindt's dark varieties.

Disclosure forms provided by the author are available with the full text of this article at NEJM.org.

From St. Luke's–Roosevelt Hospital and Columbia University, New York.

This article was published on October 10, 2012, at NEJM.org.

1. Nurk E, Refsum H, Drevon CA, et al. Intake of flavonoid-rich wine, tea, and chocolate by elderly men and women is associated with better cognitive test performance. *J Nutr* 2009;139:120-7.
2. Desideri G, Kwik-Urbe C, Grassi D, et al. Benefits in cognitive function, blood pressure, and insulin resistance through cocoa flavanol consumption in elderly subjects with mild cognitive impairment: the Cocoa, Cognition, and Aging (CoCoA) Study. *Hypertension* 2012;60:794-801.
3. Corti R, Flammer AJ, Hollenberg NK, Lüscher TF. Cocoa and cardiovascular health. *Circulation* 2009;119:1433-41.
4. Sorond FA, Lipsitz LA, Hollenberg NK, Fisher ND. Cerebral blood flow response to flavanol-rich cocoa in healthy elderly humans. *Neuropsychiatr Dis Treat* 2008;4:433-40.
5. Bisson JF, Nejdi A, Rozan P, Hidalgo S, Lalonde R, Messaoudi M. Effects of long-term administration of a cocoa polyphenolic extract (Acticoa powder) on cognitive performances in aged rats. *Br J Nutr* 2008;100:94-101.

DOI: 10.1056/NEJMon1211064

Copyright © 2012 Massachusetts Medical Society.

AN NEJM APP FOR IPHONE

The NEJM Image Challenge app brings a popular online feature to the smartphone. Optimized for viewing on the iPhone and iPod Touch, the Image Challenge app lets you test your diagnostic skills anytime, anywhere. The Image Challenge app randomly selects from 300 challenging clinical photos published in NEJM, with a new image added each week. View an image, choose your answer, get immediate feedback, and see how others answered. The Image Challenge app is available at the iTunes App Store.

whether there would be a correlation between a country's level of chocolate consumption and its population's cognitive function." Using the success of a country in winning Nobel Prizes as a surrogate for "the proportion with superior cognitive function" in a country, he analyzed the relationship between the number of Nobel laureates per capita in a country with that country's per capita chocolate consumption.

Messerli reported "a close, significant linear correlation ($r=0.791$, $p<0.0001$) between chocolate consumption per capita and the number of Nobel laureates per 10 million persons in a total of 23 countries." The relationship was even stronger when Sweden, the home of the Nobel Prize, was removed from the calculations, as it appeared to have won more Nobel prizes than expected based on its chocolate consumption. Switzerland, on the other hand, "was the top performer in terms of both the number of Nobel Laureates and chocolate consumption." (It should perhaps be noted at this point that Messerli, a hypertension expert who lives in New York City, was born in Switzerland and reports in his disclosure statement that he consumes chocolate daily, "mostly but not exclusively in the form of Lindt's dark varieties.")

Messerli duly points out that correlation does not prove causation, but, he writes, "since chocolate consumption has been documented to improve cognitive function, it seems most likely that in a dose-dependent way, chocolate intake provides the abundant fertile ground needed for the sprouting of Nobel laureates. Obviously, these findings are hypothesis-generating only and will have to be tested in a prospective, randomized trial."

Regarding Sweden's status as an outlier, Messerli writes that "one cannot quite escape the notion that either the Nobel Committee in Stockholm has some inherent patriotic bias when assessing the candidates for these awards or, perhaps, that the Swedes are particularly sensitive to chocolate, and even

The author is a Forbes contributor. The opinions expressed are those of the writer.

LARRY HUSTEN'S POPULAR POSTS

[Study Finds Sodium-Potassium Ratio Strongly Tied to Mortality and CV Disease](#) 34,222 views

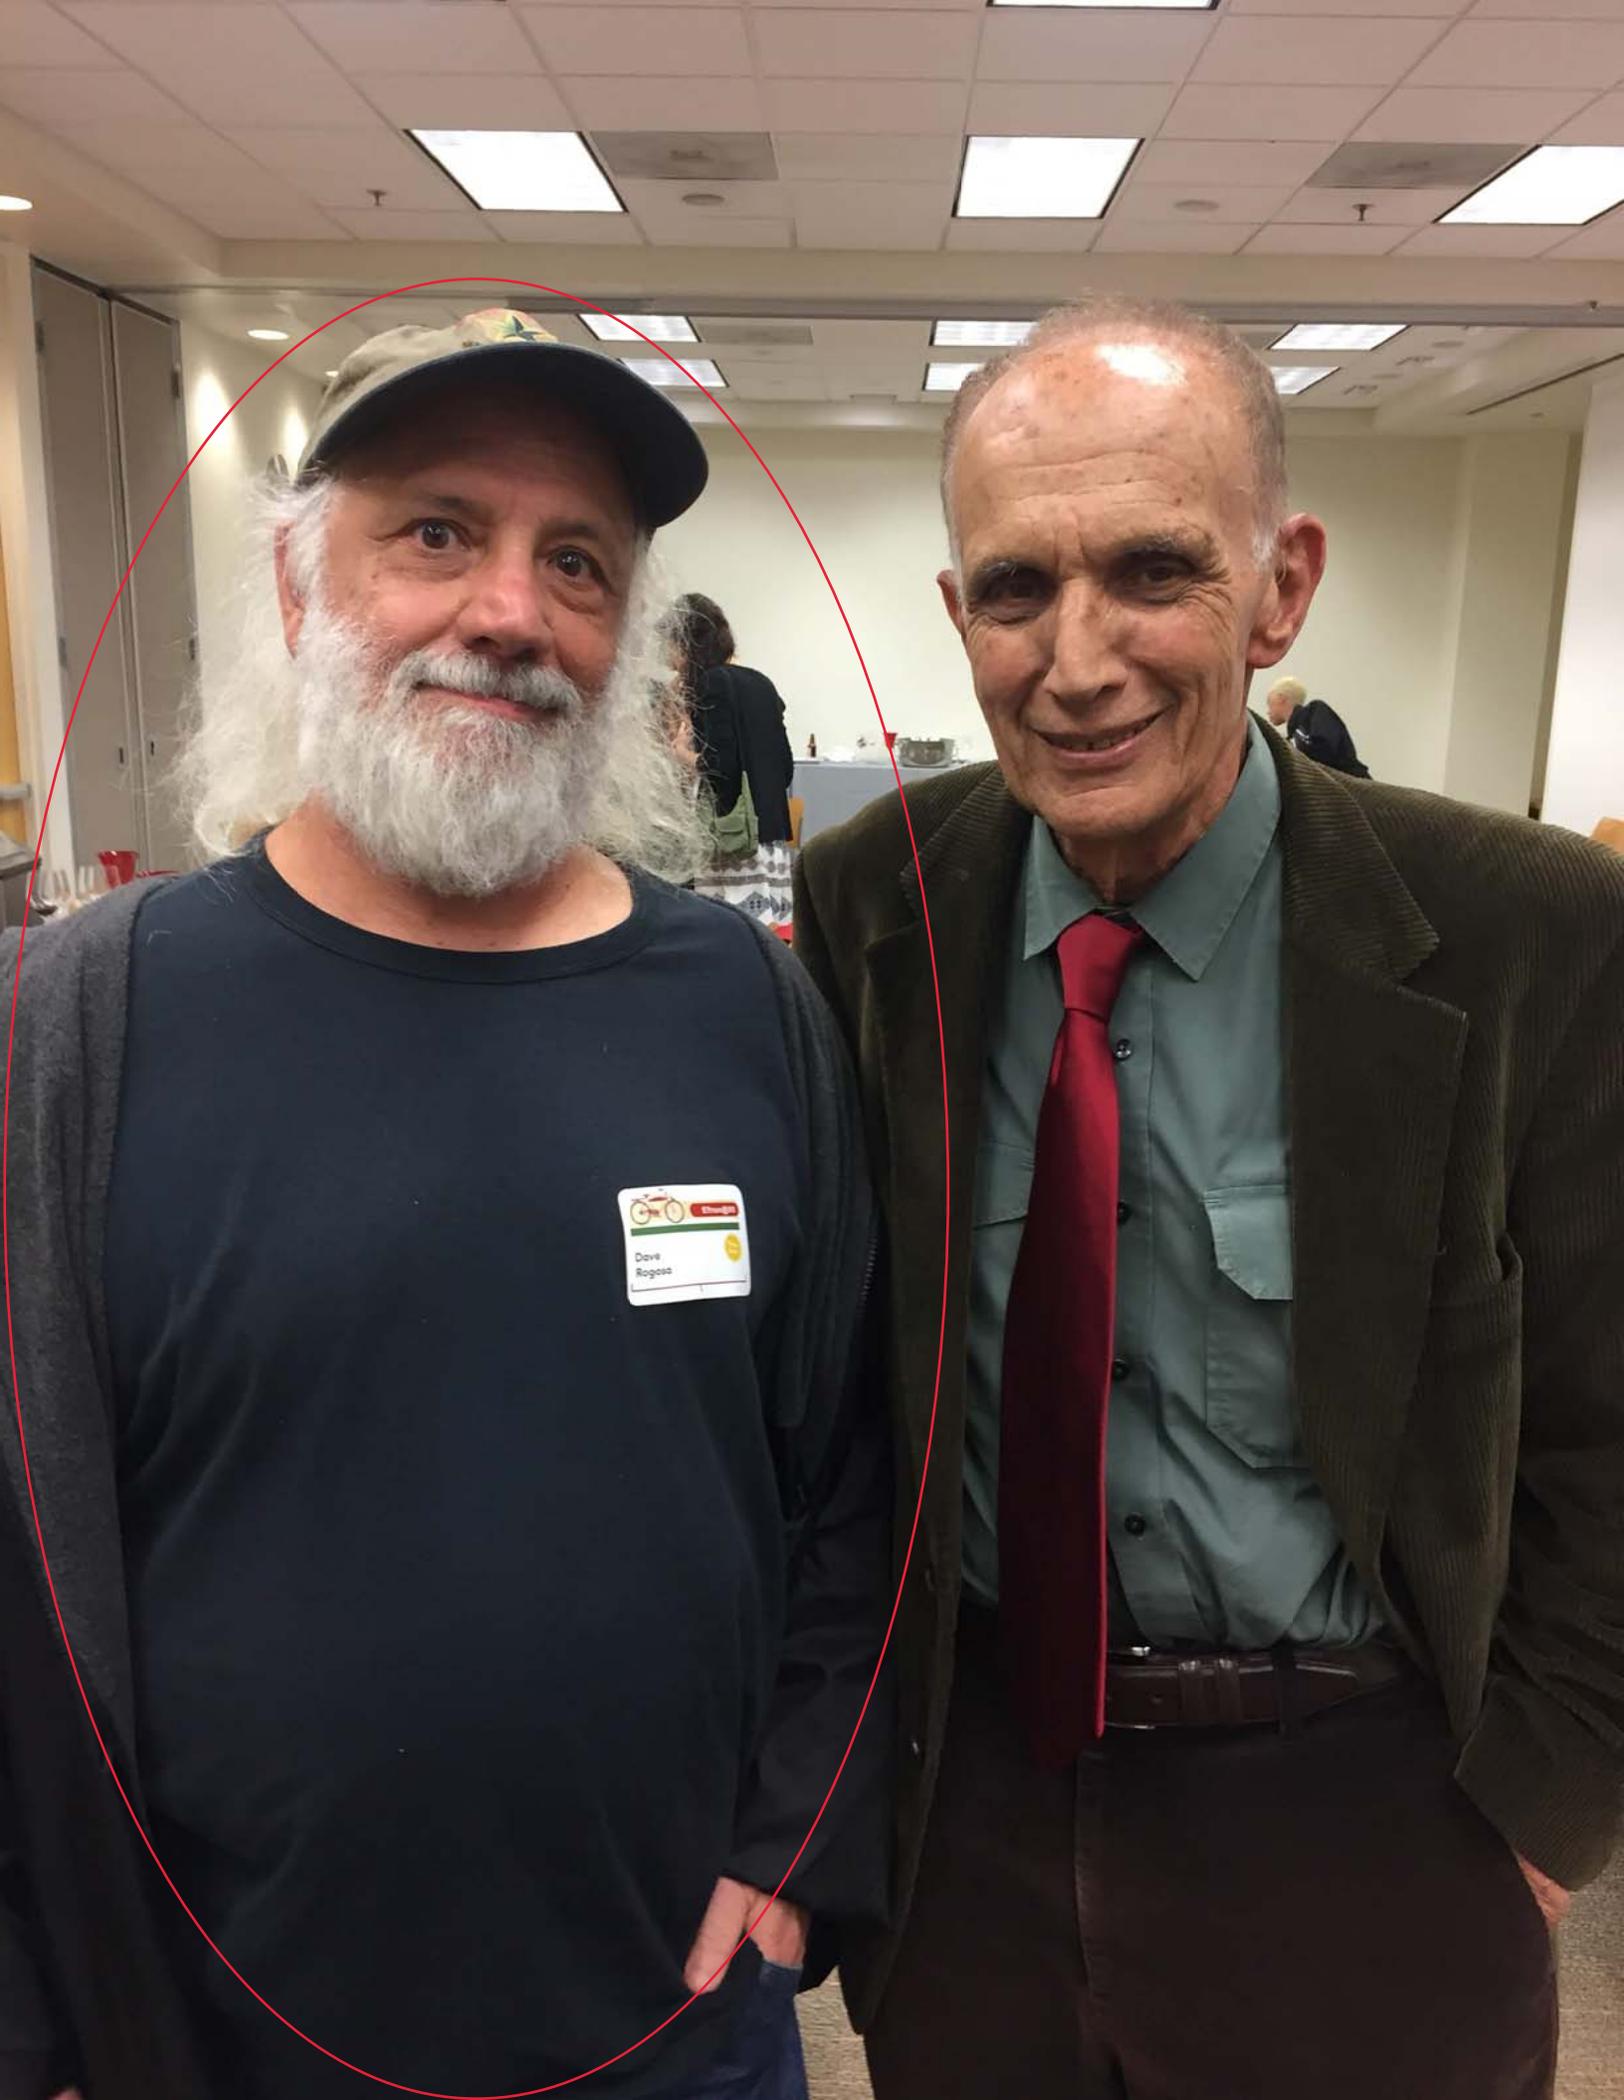
[Live Blog-- FDA Advisory Panel on Rivaroxaban \(Xarelto\)](#) 32,372 views

[FDA Approves Another New Weight Loss Drug](#) 25,768 views

[Ohio Hospital And Cardiology Group Pay \\$4.4 Million To Settle Charges Over Unnecessary PCIs](#) 25,160 views

[Resveratrol and Fraud](#) 23,437 views

MORE FROM LARRY HUSTEN



 Dave Rogosa