

Mathematical analysis of “mathematical analysis” of a vitamin D COVID-19 trial

November 17, 2020 in [sophistry](#) | Tags: [Córdoba study](#), [Manolis Kellis](#), [Randomized Controlled Trial](#), [Vitamin D](#)

Lior Pachter

Division of Biology and Biological Engineering &
Department of Computing and Mathematical Sciences California Institute of Technology

Abstract

A [recently published](#) pilot study on the efficacy of 25-hydroxyvitamin D3 (calcifediol) in reducing ICU admission of hospitalized COVID-19 patients, concluded that the treatment “seems able to reduce the severity of disease, but larger trials with groups properly matched will be required to show a definitive answer”. In a [follow-up](#) paper, Jungreis and [Kellis](#) re-examine this so-called “Córdoba study” and argue that the authors of the study have undersold their results. Based on a reanalysis of the data in a manner they describe as “rigorous” and using “well established statistical techniques”, they urge the medical community to “consider testing the vitamin D levels of all hospitalized COVID-19 patients, and taking remedial action for those who are deficient.” Their recommendation is based on two claims: in an examination of unevenness in the distribution of one of the comorbidities between cases and controls, they conclude that there is “no evidence of incorrect randomization”, and they present a “mathematical theorem” to make the case that the effect size in the Córdoba study is significant to the extent that “they can be confident that if assignment to the treatment group had no effect, we would not have observed these results simply due to chance.”

Unfortunately, the “mathematical analysis” of Jungreis and Kellis is deeply flawed, and their “theorem” is vacuous. Their analysis cannot be used to conclude that the Córdoba study shows that calcifediol significantly reduces ICU admission of hospitalized COVID-19 patients. Moreover, the Córdoba study is fundamentally flawed, and therefore there is nothing to learn from it.

The Córdoba study

The Córdoba study, described by the authors as a pilot, was ostensibly a randomized controlled trial, designed to determine the efficacy of 25-hydroxyvitamin D3 in reducing ICU admission of hospitalized COVID-19 patients. The study consisted of 76 patients hospitalized for COVID-19 symptoms, with 50 of the patients treated with calcifediol, and 26 not receiving treatment. Patients were administered “standard care”, which according to the authors consisted of “a combination of hydroxychloroquine, azithromycin, and for patients with pneumonia and NEWS score 5, a broad spectrum antibiotic”. Crucially, admission to the ICU was determined by a “Selection Committee” consisting of intensivists, pulmonologists, internists, and members of an ethics committee. The Selection Committee based ICU admission decisions on the evaluation of several criteria, including presence of comorbidities, and the level of dependence of patients according to their needs and clinical criteria.

The result of the Córdoba trial was that only 1/50 of the treated patients was admitted to the ICU, whereas 13/26 of the untreated patients were admitted (p -value = 7.7×10^{-7} by Fisher’s exact test). This is a minuscule p -value but it is meaningless. Since there is no record of the Selection Committee deliberations, it is impossible to know whether the ICU admission of the 13 untreated patients was due to their previous high blood pressure comorbidity. Perhaps the 11 treated patients with the comorbidity were not admitted to the ICU because they were older, and the Selection Committee considered their previous higher blood pressure to be more “normal” (14/50 treatment patients were over the age of 60, versus only 5/26 of the untreated patients).

Figure 1: Table 2 from [1] showing the comorbidities of patients. It is reproduced by virtue of [1] being published open access under the CC-BY license.

The fact that admission to the ICU could be decided in part based on the presence of co-morbidities, and that there was a significant imbalance in one of the comorbidities, immediately renders the study results meaningless. There are several other problems with it that potentially confound the results: the study did not examine the Vitamin D levels of the treated patients, nor was the untreated group administered a placebo. Most importantly, the study numbers were tiny, with only 76 patients examined. Small studies are notoriously problematic, and are known to produce large effect sizes [9]. Furthermore, sloppiness in the study does not lead to confidence in the results. The authors state that the “rigorous protocol” for determining patient admission to the ICU is available as Supplementary Material, but there is no Supplementary Material distributed with the paper.

Prognostic factors for COVID-19 at baseline.

Poor prognosis risk factor	Group receiving Calcifediol (n = 50)	Group without Calcifediol (n = 26)	IC 95 %	P
≥ 60 years	14 (28 %)	5 (19.23 %)	-0.11 - 0.28	0.40
Previous lung disease	4 (8%)	2 (7.69 %)	-0.12 - 0.13	0.96
Previous Chronic kidney disease	0	0	-	-
Previous Diabetes mellitus	3 (6%)	5 (19.23 %)	-0.30 - 0.03	0.08
Previous High blood pressure	11 (24.19 %)	15 (57.69 %)	-0.58 - -0.13	0.002
Previous Cardiovascular disease	2 (4%)	1 (3.85 %)	-0.09 - 0.09	0.97
Immunosuppressed & transplanted	6 (12 %)	1 (3.85 %)	-0.03 - 0.20	0.24
At least one prognostic bad risk factor ^a	24 (48 %)	16 (61.54 %)	-0.37 - 0.10	0.26
PaO ₂ /FiO ₂ (mean +/-SD)	346.57 +/- 73.38	334.62 +/- 66.33	-22.29 - 46.19	0.49
C-reactive protein (mg/L) (mean +/-SD)	82.93 +/- 62.74	94.71 +/- 63.64	-42.15 - 18.59	0.44
LDH (U/L)(mean +/-SD)	308.12 +/- 83.83	345.81 +/- 108.57	-82.46 - 7.08	0.10
D-Dimer (ng/mL) (mean +/-SD)	650.92 +/- 405.61	1333.54 +/- 2570.50	-360.29 - 1725.53	0.19
Lymphocytes < 800/ μ L	10 (20 %)	6 (23.08 %)	-0.16 - 0.23	0.75
Ferritin (ng/mL) (mean +/-SD)	691.04 +/- 603.54	825.16 +/- 613.95	-166.31 - 434.55	0.36
IL-6 (22/48) (pg/mL) (mean +/-SD)	28.88 +/- 75.05	19.54 +/- 19.45	-41.88 - 23.19	0.41

SD: Standard Deviation.

^a Patients with at least one of the following risk factors (age >60, previous lung disease, chronic kidney disease, diabetes mellitus, hypertension, cardiovascular disease or Immunosuppressed and transplanted patients).

There is also an embarrassing typo: Fisher's exact test is referred to twice as "Fischer's test". To err once in describing this classical statistical test may be regarded as misfortune; to do it twice looks like carelessness.

A pointless statistics exercise

The Córdoba study has not received much attention, which is not surprising considering that by the authors' own admission it was a pilot that at best only motivates a properly matched and powered randomized controlled trial. Indeed, the authors mention that such a trial (the COVIDIOL trial), with data being collected from 15 hospitals in Spain, is underway. Nevertheless, Jungreis and Kellis [3], apparently mesmerized by the $7.7 * 10^{-7}$ p-value for ICU admission upon treatment, felt the need to "rescue" the study with what amounts to faux statistical gravitas. They argue for immediate consideration of testing Vitamin D levels of hospitalized patients, so that "deficient" patients can be administered some form of Vitamin D "to the extent it can be done safely". Their message has been noticed; only a few days after [3] appeared the authors' tweet to promote it has been retweeted more than 50 times [8].

Jungreis and Kellis claim that the p-value for the effect of calcifediol on patients is so significant, that in and of itself it merits belief that administration of calcifediol does, in fact, prevent admission of patients to ICUs. To make their case, Jungreis and Kellis begin by acknowledging that imbalance between the treated and untreated groups in the previous high blood pressure comorbidity may be a problem, but claim that there is "no evidence of incorrect randomization." Their argument is as follows: they note that while the p-value for the imbalance in the previous high blood pressure comorbidity is 0.0023, it should be adjusted for the fact that there are 15 distinct comorbidities, and that just by chance, when computing so many p-values, one might be small. First, an examination of Table 2 in [1] (Figure 1) shows that there were only 14 comorbidities assessed, as none of the patients had previous chronic kidney disease. Thus, the number 15 is incorrect. Second, Jungreis and Kellis argue that a Bonferroni correction should be applied, and that this correction should be based on 30 tests ($=15 * 2$). The reason for the factor of 2 is that they claim that when testing for imbalance, one should test for imbalance in both directions. By applying the Bonferroni correction to the p-values, they derive a "corrected" p-value for previous high blood pressure being imbalanced between groups of 0.069. They are wrong on several counts in deriving this number. To illustrate the problems we work through the calculation step-by-step:

The question we want to answer is as follows: given that there are multiple comorbidities, is there is a significant imbalance in *at least* one comorbidity. There are several ways to test for this, with the simplest being Šidák's correction [10] given by

$$q = 1 - (1 - m)^n,$$

where m is the minimum p-value among the comorbidities, and n is the number of tests. Plugging in $m = 0.0023$ (the smallest p-value in Table 2 of [1]) and $n = 14$ (the number of comorbidities) one gets 0.032 (note that the Bonferroni correction used by Jungreis And Kellis is the Taylor approximation to the Šidák correction when m is small). The Šidák correction is based on an assumption that the tests are independent. However, that is certainly not the case in the Córdoba study. For example, having at least one prognostic factor is one of the comorbidities tabulated. In other words, the p-value obtained is conservative. The calculation above uses $n = 14$, but Jungreis and Kellis reason that the number of tests is $30 = 15 * 2$, to take into account an imbalance in either the treated or untreated direction. Here they are assuming two things: that two-sided tests for each comorbidity will produce double the p-value of a one-sided test, and that two sided tests are the "correct" tests to perform. They are wrong on both counts. First, the two-sided Fisher exact test does not, in general produce a p-value that is double the 1-sided test. The study result is a good example: 1/49 treated patients admitted to the ICU vs. 13/26 untreated patients produces a p-value of $7.7 * 10^{-7}$ for both the 1-sided and 2-sided tests. Jungreis and Kellis do not seem to know this can happen, nor understand why; they go to great lengths to explain the importance of conducting a 1-sided test for the study result. Second, there is a strong case to be made that a 1-sided test is the correct test to perform for the comorbidities. The concern is not whether there was an imbalance of any sort, but whether the imbalance would skew results by virtue of the study including too many untreated individuals with comorbidities. In any case, if one were to give Jungreis and Kellis the benefit of the doubt, and perform a two sided test, the corrected p-value for the previous high blood pressure comorbidity is 0.06 and not 0.069.

The most serious mistake that Jungreis and Kellis make, however, is in claiming that one can accept the null hypothesis of a hypothesis test when the p-value is greater than 0.05. The p-value they obtain is 0.069 which, even if it is taken at face value, is not grounds for claiming, as Jungreis and Kellis do, that "this is not significant evidence that the assignment was not random" and reason to conclude that there is "no evidence of incorrect randomization". That is not how p-values work. A p-value less than 0.05 allows one to reject the null hypothesis (assuming 0.05 is the threshold chosen), but a p-value above the chosen threshold is not grounds for accepting the null. Moreover, the corrected p-value is 0.032 which is certainly grounds for rejecting the null hypothesis that the randomization was random.

Correction of the incorrect Jungreis and Kellis statistics may be a productive exercise in introductory undergraduate statistics for some, but it is pointless insofar as assessing the Córdoba study. While the extreme imbalance in the previous high blood pressure comorbidity is problematic because patients with the comorbidity may be more likely to get sick and require ICU admission, the study was so flawed that the exact p-value for the imbalance is a moot point. Given that the presence of comorbidities, not just their effect on patients, was a factor in determining which patients were admitted to the ICU, the extreme imbalance in the previous high blood pressure comorbidity renders the result of the study meaningless *ex facie*.

A definition is not a theorem is not proof of efficacy

In an effort to fend off criticism that the comorbidities of patients were improperly balanced in the study, Jungreis and Kellis go further and present a “theorem” they claim shows that there was a minuscule chance that an uneven distribution of comorbidities could render the study results not significant. The “theorem” is stated twice in their paper, and I’ve copied both theorem statements verbatim from their paper:

Theorem 1 *In a randomized study, let p be the p-value of the study results, and let q be the probability that the randomization assigns patients to the control group in such a way that the values of $P_{prognostic}(Patient)$ are sufficiently unevenly distributed between the treatment and control groups that the result of the study would no longer be statistically significant at the 95% level after p controlling for the prognostic risk factors. Then $q < \frac{p}{0.05}$.*

According to Jungreis and Kellis, $P_{prognostic}(Patient)$ is the following: “There can be any number of prognostic risk factors, but if we knew what all of them were, and their effect sizes, and the interactions among them, we could combine their effects into a single number for each patient, which is the probability, based on all known and yet-to-be discovered risk factors at the time of hospital admission, that the patient will require ICU care if not given the calcifediol treatment. Call this (unknown) probability $P_{prognostic}(Patient)$.”

The theorem is restated in the Methods section of Jungreis and Kellis paper as follows:

Theorem 2 *In a randomized controlled study, let p be the p-value of the study outcome, and let q be the probability that the randomization distributes all prognostic risk factors combined sufficiently unevenly between the treatment and control groups that when controlling for these prognostic risk p factors the outcome would no longer be statistically significant at the 95% level. Then $q < \frac{p}{0.05}$.*

While it is difficult to decipher the language the “theorem” is written in, let alone its meaning (note Theorem 1 and Theorem 2 are supposedly the same theorem), I was able to glean something about its content from reading the “proof”. The mathematical content of whatever the theorem is supposed to mean, is the definition of conditional probability, namely that if A and B are events with $P(B) > 0$, then

$$P(A|B) := \frac{P(A \cap B)}{P(B)}.$$

To be fair to Jungreis and Kellis, the “theorem” includes the observation that

$$P(A \cap B) \leq P(A) \Rightarrow P(A|B) \leq \frac{P(A)}{P(B)}.$$

This is not, by any stretch of the imagination, a “theorem”; it is literally the definition of conditional probability followed by an elementary inequality. The most generous interpretation of what Jungreis and Kellis were trying to do with this “theorem”, is that they were showing that the p-value for the study is so small, that it is small even after being multiplied by 20. There are less generous interpretations.

Does Vitamin D intake reduce ICU admission?

There has been a lot of interest in Vitamin D and its effects on human health over the past decade [2], and much speculation about its relevance for COVID-19 susceptibility and disease severity. One interesting result on disease susceptibility was published recently: in a study of 489 patients, it was found that the relative risk of testing positive for COVID-19 was 1.77 times greater for patients with likely deficient vitamin D status compared with patients with likely sufficient vitamin D status [7]. However, definitive results on Vitamin D and its relationship to COVID-19 will have to await larger trials. One such trial, a large randomized clinical trial with 2,700 individuals sponsored by Brigham and Women’s Hospital, is currently underway [4]. While this study might shed some light on Vitamin D and COVID-19, it is prudent to keep in mind that the outcome is not certain. Vitamin D levels are confounded with many socioeconomic factors, making the identification of causal links difficult. In the meantime, it has been suggested that

it makes sense for individuals to maintain reference nutrient intakes of Vitamin D [6]. Such a public health recommendation is not controversial.

As for Vitamin D administration to hospitalized COVID-19 patients reducing ICU admission, the best one can say about the Córdoba study is that nothing can be learned from it. Unfortunately, the poor study design, small sample size, availability of only summary statistics for the comorbidities, and imbalanced comorbidities among treated and untreated patients render the data useless. While it may be true that calcifediol administration to hospital patients reduces subsequent ICU admission, it may also not be true. Thus, the follow-up by Jungreis and Kellis is pointless at best. At worst, it is irresponsible propaganda, advocating for potentially dangerous treatment on the basis of shoddy arguments masked as “rigorous and well established statistical techniques”. It is surprising to see Jungreis and Kellis argue that it may be unethical to conduct a placebo randomized controlled trial, which is one of the most powerful tools in the development of safe and effective medical treatments. They write “the ethics of giving a placebo rather than treatment to a vitamin D deficient patient with this potentially fatal disease would need to be evaluated.” The evidence for such a policy is currently non-existent. On the other hand, there are plenty of known risks associated with excess Vitamin D [5].

References

1. Marta Entrenas Castillo, Luis Manuel Entrenas Costa, José Manuel Vaquero Barrios, Juan Francisco Alcalá Díaz, José López Miranda, Roger Bouillon, and José Manuel Quesada Gomez. Effect of calcifediol treatment and best available therapy versus best available therapy on intensive care unit admission and mortality among patients hospitalized for COVID-19: A pilot randomized clinical study. *The Journal of steroid biochemistry and molecular biology*, 203:105751, 2020.
2. Michael F Holick. Vitamin D deficiency. *New England Journal of Medicine*, 357(3):266–281, 2007.
3. Irwin Jungreis and Manolis Kellis. Mathematical analysis of Córdoba calcifediol trial suggests strong role for Vitamin D in reducing ICU admissions of hospitalized COVID-19 patients. *medRxiv*, 2020.
4. JoAnn E Manson. <https://clinicaltrials.gov/ct2/show/nct04536298>.
5. Ewa Marcinowska-Suchowierska, Małgorzata Kupisz-Urbańska, Jacek Łukaszkiwicz, Paweł Płudowski, and Glenville Jones. Vitamin D toxicity—a clinical perspective. *Frontiers in endocrinology*, 9:550, 2018
6. Adrian R Martineau and Nita G Forouhi. Vitamin D for COVID-19: a case to answer? *The Lancet Diabetes & Endocrinology*, 8(9):735–736, 2020.
7. David O Meltzer, Thomas J Best, Hui Zhang, Tamara Vokes, Vineet Arora, and Julian Solway. Association of vitamin D status and other clinical characteristics with COVID-19 test results. *JAMA network open*, 3(9):e2019722–e2019722, 2020.
8. Vivien Shotwell. <https://tweetstamp.org/1327281999137091586>.
9. Robert Slavin and Dewi Smith. The relationship between sample sizes and effect sizes in systematic reviews in education. *Educational evaluation and policy analysis*, 31(4):500–506, 2009.
10. Lynn Yi, Harold Pimentel, Nicolas L Bray, and Lior Pachter. Gene-level differential analysis at transcript-level resolution. *Genome biology*, 19(1):53, 2018.