Use of a null assumption to re-analyze data collected through a rolling cohort subject to selection bias due to informative censoring

Mark F. Reeder

Air Force Institute of Technology

Abstract

A novel method of estimating selection bias due to informative censoring for a rolling cohort utilizing matches is demonstrated for a recently published, and highly influential, study. The core reason for the bias is related to the principle that those with Covid-19 symptoms are prevented from obtaining a vaccine. The study explored the efficacy of the BNT162b2 mRNA Covid-19 vaccine in a nationwide mass vaccination setting in Israel, and was conducted over a 44-day study period spanning December 2020 through early February 2021. The present approach utilizes the published data to establish a population entering the study, and the number of matches which exist at the end of the study period is also a known value. Since the time-wise distribution of those censored was not made available, two different distribution patterns were compared. Simple probability rules were applied to estimate those who rolled out of the cohort on any given day. Under those circumstances, a null assumption (i.e., that the vaccine has no effect) for the exposed group clearly led to a different value for the outcome of interest, demonstrating the effect of the bias. The expected value from the null assumption could then be compared to the actual measured value to yield efficacy. It was found that the time-wise distribution of those rolling out strongly influenced the level of bias. Discussion of the process includes a brief overview of entities known from the Kaplan-Meier method of data presentation, from which several relevant inferences may be drawn. First, a transparent listing of the number of matches which both enter and exit the cohort for all possible combination of days within the study period would lead to a more robust approximation of the selection bias due to informative censoring. The literature suggests that such a practice be implemented and also that the estimate would likely be further improved if hazard ratios were modeled, and this effect was confirmed herein. Two other examples from the literature are explored and lead to a question: Is it possible that a similar selection bias due to informative censoring might also be responsible, at least in part, for the difference between research studies which report high vaccine efficiency and relatively poor outcomes in real-world environments? If so, the use of a booster to improve outcomes is called into question. All studies and reports on Covid-19 vaccine efficiency should be assessed to determine the significance of selection bias due to informative censoring and provide a fully-transparent description of how this often-neglected bias is taken into account.

Introduction

Recently, the article, “BNT162b2 mRNA Covid-19 Vaccine in a Nationwide Setting” by Dagan et al. was published in the New England Journal of Medicine (April 15, 2021). It describes outcomes of a large-scale study of the campaign in Israel to vaccinate a large population against Covid-19. The study matched newly vaccinated (“exposed”) people with unvaccinated (“controls”) in a 1:1 ratio according to demographic and clinical characteristics. Each study group included 596,691 persons. Its influence on subsequent studies is profound as it has been cited over 300 times in a period of only a few months.[1]

The study utilized a “rolling cohort” as described in the authors’ supplemental material. This approach was required since vaccination initiated in Israel on December 20, 2020, (calendar) Day 1 of the study period. Each match was tracked within the 44-day study period, and outcomes were compared by counting symptom occurrence, hospitalization, severe Covid-19 cases, and death. Matches accumulated over the entirety of the study, and the number of matches established on each day was provided by the authors in the form of a bar graph in their Figure S1 of the supplemental material, which is shown herein as Figure 1. Since it was one of the first research articles to describe a large-scale Covid-19 vaccination effort, it has been widely referenced not only by researchers, but also by government decision-makers, and -through the media- the worldwide general public.

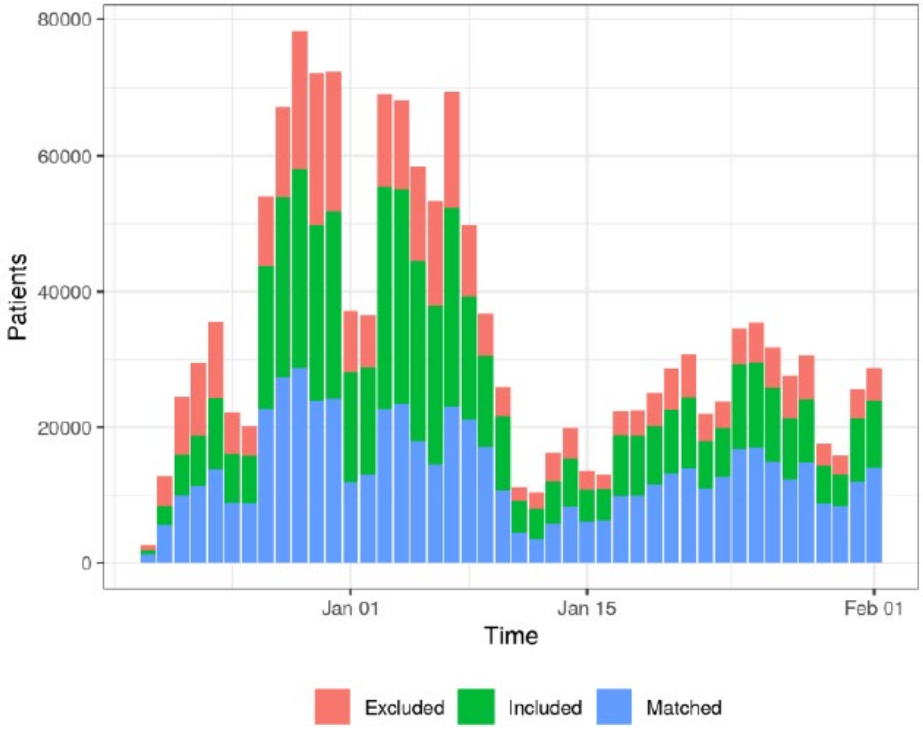

Figure 1. Excerpt from the supplemental material of Dagan et al. (their Figure S1). Matches are shown in blue.

In a traditional study, one would begin with the same number of matched people (e.g., 596,618) in each group, and the researchers would document how the exposed group fared in comparison to the control group. From the initiation of the study period until its end, one would compile the number of events (e.g., documented infection, occurrence of symptoms, hospitalization, severe disease, and death). One could then report a direct comparison of the two groups to evaluate the effectiveness of the treatment (e.g., a vaccine).

Unlike a traditional study, the number of matches for that of Dagan et al. changed significantly with time. In addition to accumulating matches each day, there was another way this change was manifested. A large portion of the matched people (259,941 people, which is 44%) who were initially in the unvaccinated control group were vaccinated within the 44-day study period. When vaccination of a person initially in the unvaccinated cohort occurred, the follow-up ended for the matched pair as the unvaccinated person “rolled out” (hereafter used without quotes) of the control cohort. In some cases, a

new unvaccinated match for a person who was initially in the (unvaccinated) cohort was found and counted as another (new) match. The study authors refer to this step as “re-matching,” and 86,601 individuals were re-matched during the study period.

There are two other reasons why the follow-up can end prior to the end of the study period. As they state on p. 1413, “For each person, follow-up ended at the earliest of the following events: occurrence of an outcome event, death unrelated to Covid-19, vaccination (for unvaccinated controls), vaccination of the matched control (for vaccinated persons), or the end of the study period.” The “occurrence of an outcome event” is a vague statement, which is also addressed later in the discussion. It is certainly clear that ‘death of one of the matched persons due to Covid-19’ would be an outcome which would end follow-up, and that is what is assumed for the purpose of assessing the selection bias. Otherwise, if follow-up simply concluded after one of the matched individuals merely experienced Covid-19 symptoms (also described as an “outcome event”), one would expect that there would be zero deaths reported in either arm. The second, “death unrelated to Covid-19,” is much easier to interpret. Both of these types of outcomes lead to only a very small percentage of the overall reason for follow-up to end ahead of the study period conclusion when compared to the 44% which terminated because the initially unvaccinated person attained a vaccine.

The majority of clinicians, health officials, and the general public are most interested in the efficacy of a vaccine for preventing the most severe outcomes. Thus, one of the most critical figures, presenting the number of deaths in each group (Fig. 2E from Dagan et al.), is repeated in Figure 2 below. It is useful as a benchmark to note that the cumulative number of deaths reported therein is 32 for the unvaccinated cohort while it is only 9 for the vaccinated cohort. These two values (32 and 9) are important to bear in mind as this re-analysis of the data is undertaken. Dagan et al. report that the estimated differences in preventing death from Covid-19 was “72% for days 14 through 20 (after one dose) and increases to 84% for days 21 through 27 (gradually shifting between the first and second vaccination doses).” It happens that the cumulative reduction in deaths reflected in Figure 2 is also 72% as (32-9)/32=0.72. So it stands to reason that cumulative number of deaths is a reasonable parameter to consider and offers the advantage of simplicity of interpretation. Undoubtedly, many who read the article will give those two values outsized attention. While it happens to be that (32-9)/32 is equal to 0.72, the authors used a Kaplan-Meier estimator to compute risk ratios for specific time periods after vaccination. Extensive Kaplan-Meier life tables are provided by the authors on pp. 46-54 of their supplemental material. Attributes and limitations of the Kaplan-Meier approach to data analysis will be forthcoming in the discussion, but it is appropriate to first lay the groundwork for why such a discussion is useful and necessary.

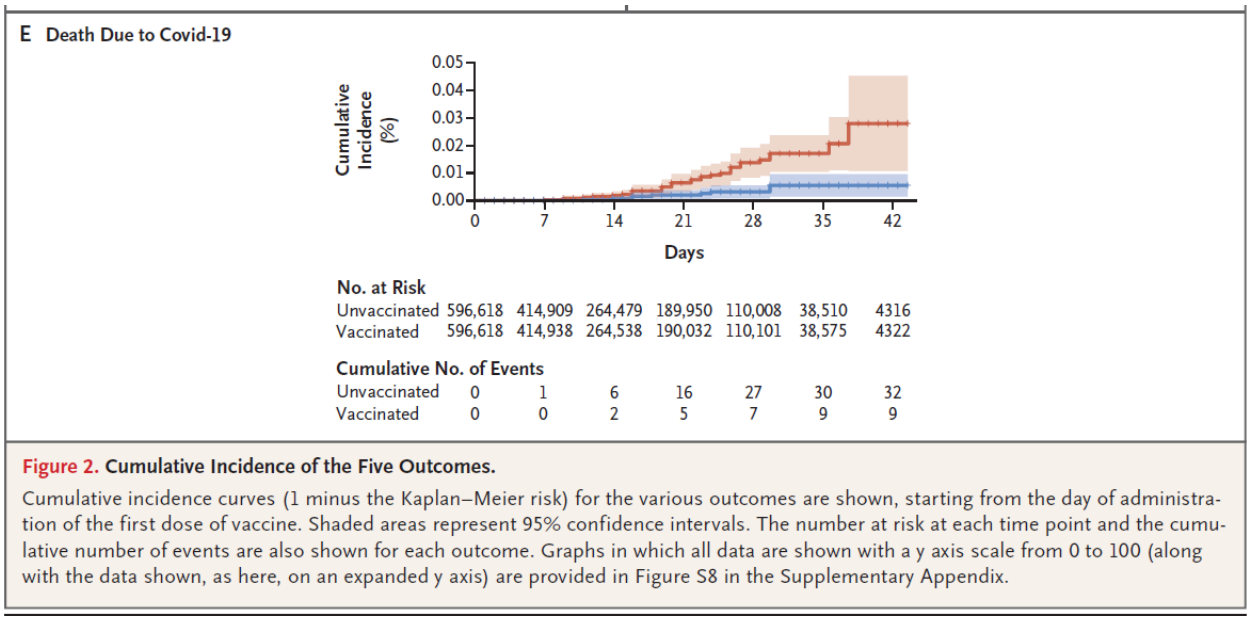

Figure 2. Cumulative Incidence of the Five Outcomes.

Cumulative incidence curves (1 minus the Kaplan–Meier risk) for the various outcomes are shown, starting from the day of administration of the first dose of vaccine. Shaded areas represent 95% confidence intervals. The number at risk at each time point and the cumulative number of events are also shown for each outcome. Graphs in which all data are shown with a y axis scale from 0 to 100 (along with the data shown, as here, on an expanded y axis) are provided in Figure S8 in the Supplementary Appendix.

Figure 2. Number of deaths due to Covid-19 reported as an outcome by Dagan et al. in the body of their NEJM article. (permission to reproduce requested)

The large number of people who roll out of the control cohort contributes to a peculiar bias when another important fact is considered: people who have Covid-19 symptoms are not permitted to obtain the vaccine. Specific guidance from the CDC is as follows from their "Frequently-asked-questions" webpage: "Question: Can I get vaccinated against COVID-19 while I am currently sick with COVID-19? Answer: No. People with COVID-19 who have symptoms should wait to be vaccinated until they have recovered from their illness and have met the criteria for discontinuing isolation; those without symptoms should also wait until they meet the criteria before getting vaccinated. This CDC guidance also applies to people who get COVID-19 before getting their second dose of vaccine."[2] Media reports provided similar guidance along with common-sense reasoning in early 2021, during the Israeli vaccination program reported by Dagan et al. For example, one media outlet reports: "If there's a chance you have COVID, however, then it's a different story. If you are having upper respiratory symptoms, the first thing you should do is get tested for COVID-19, Dr. Mandal says. 'For one thing, if you do have COVID-19 or are awaiting test results, you should immediately self-isolate, and definitely shouldn't expose the person giving you the shot….', '…If you currently have the virus, then getting vaccinated will not be immediately helpful as the body takes time to mount an immune response,' says Dr. Eudene Harry MD, a board-certified emergency medicine physician in Orlando, Florida."[3]

This difficulty arising from this process leads a "selection bias" due to "informative censoring," which has been identified as an often-neglected cause of bias in the literature,[4] despite several studies documenting and modeling its effect. Campigotto and Weller used simulation studies to evaluate bias when combined with Kaplan-Meier analysis. Their simulations, using proportions of informative censoring ranging from 1% to 40%, showed that, "the magnitude of the bias depends primarily on the proportion of patients who are informatively censored and secondarily on the hazard ratio between those who are informatively censored and those who remain in the study." They recommend that: "If informative censoring cannot be avoided, then all patients should be observed until progression, and sensitivity analyses should be used as appropriate."[5]

It is very easy to understand how such a bias might play a role in assessing vaccine efficacy in the Dagan et al. study by way of a counter-example. Suppose that the recommendation were for individuals to immediately seek the vaccine as quickly as possible if Covid-19 symptoms appear. How many deaths would be expected in the unvaccinated group for a study using the rules for ending the follow-up set in place by Dagan et al.? The answer is zero. After all, there would be enough time (hours, presumably) for an unvaccinated person to obtain the vaccine. Once that occurs, they would roll out of the unvaccinated cohort, and follow-up of the matched pair would end. Therefore, even if the person in this example ultimately died, their death would not be tallied using the rules set in place for this study. In a similar vein, one could argue that the number counted as severe Covid-19 cases and hospitalizations could be driven to zero for those labelled "unvaccinated" under such a circumstance if the response were sufficiently fast.

A direct example applicable to the actual circumstances also provides guidance if one considers two opposite scenarios. For each, assume that the individual in question is matched to someone outside the example discussion (in other words, they would not be matched to one another). 1) Suppose a person is matched into the *unvaccinated* group on (calendar) Day 5, experiences Covid-19 symptoms on Day 9, and dies 18 days later on Day 27. In this case, the individual would have been restricted from receiving the vaccine, and hence from rolling out of the unvaccinated cohort, from Days 9 through 26. Barring the extremely highly unlikely possibility that this person's match in the vaccinated group happens to die sooner than Day 27, the outcome would be counted as a death in the unvaccinated group. 2) Now suppose that the same progression occurs for a person in the *vaccinated* group – joining the group by being vaccinated on (calendar) Day 5, experiencing symptoms on Day 9, and dying on Day 27. This death would only be counted if the person's match remains unvaccinated over this period. In these two situations, the possibility of the unvaccinated person rolling out is the same as that of the match for the vaccinated person on Days 6 through 8. However, from Days 9 through 27, the restriction of the unvaccinated control from being vaccinated due to symptomatic Covid-19 leads to the potential for a significant bias in counting deaths.

This bias is acknowledged by Dagan et al. and described as "a selection bias due to informative censoring." They address the concern with respect to event outcomes, including severe Covid-19 and death, and their description is as follows: "To assess a possible selection bias that could stem from informative censoring, whereby controls who are vaccinated feel well around the time of vaccination, we performed a sensitivity analysis in which they were kept in the unvaccinated group for a period of time that was set differently for each outcome (Figure S7 of their supplement and Table S5). This analysis showed results similar to the main analysis, which suggests that such bias was small in our analysis." (p. 1422) Additional specific details are also provided, as follows: "Table S5 and Figure S7 show the results of the sensitivity analysis in which data for persons who were enrolled as controls and were then vaccinated were censored at a delay (a number of days after the vaccination date, depending on the outcome). The estimates are similar to those of the main analysis." (p. 1417)

To clarify, Dagan et al. present the outcomes without any correction of bias due to informative censoring in the body of their article as the "main analysis." Meanwhile, the outcomes with bias taken into account in the fashion they describe, presumably with the goal of eliminating the bias, are presented in their supplemental material.

Some specific guidance for their approach was given in their 'Methods' section: "We performed an additional sensitivity analysis to assess the potential for selection bias due to informative censoring. In this analysis, data on controls who were subsequently vaccinated were censored only after 7 days plus the median time from documented Covid-19 diagnosis to the outcome being studied." Thus, when death was considered as the outcome event, the delay in censoring can be computed using Dagan et al.'s Figure S3E, replicated here as Figure 3. It can be seen that the median time described is 11 days. Therefore, the total delay period for the sensitivity analysis for the outcome of death is 7+11 = 18 days.

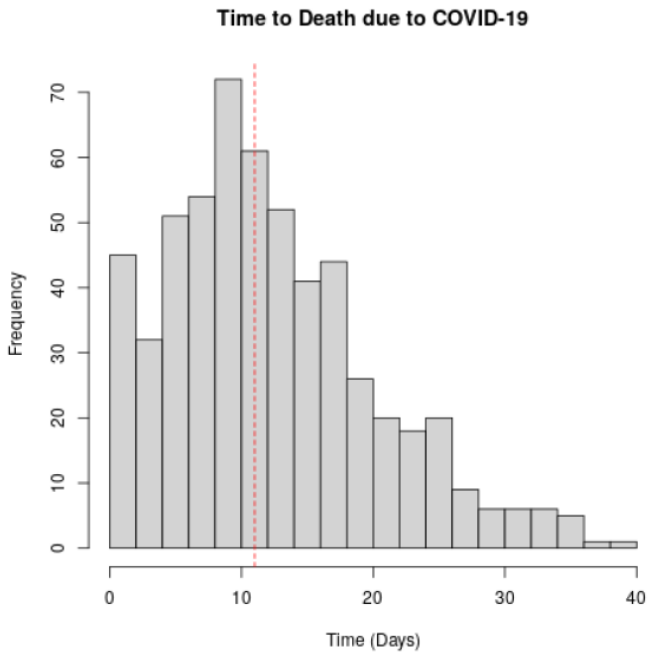

Figure 3. This histogram, reproduced from Dagan et al., shows "the distribution of time (in days) from the date the first positive PCR test was sampled until incidence of the different outcomes: hospitalization due to COVID-19 and death due to COVID-19. A dashed vertical red line in each plot shows the median." The median time between PCR-positive test and death is depicted as 11 days. (permission to reproduce requested)

Considering that a person must be free of Covid-19 symptoms on the day they are vaccinated, this exploration of the bias due to informative censoring made by Dagan et al. implies that a person could have (1) acquired the virus prior to vaccination, (2) presented no symptoms at the time of vaccination, and (3) died within 18 days of vaccination. In other words, the sensitivity analysis presumed that death would occur within 18 days of presenting symptoms. Furthermore, this assumption implies that the vaccine would have no effect for the first 18 days after it is administered. Is it possible that such an assumption could lead to a circular argument?

The analysis of this bias by Dagan et al. provided charts in Fig S7 ("Cumulative Incidence Curves when Delaying Censoring of Vaccinated Controls"). The specific chart for "Death due to Covid-19" is replicated in Figure 4 where it can be seen that the cumulative number of deaths for both unvaccinated and vaccinated groups rise. For the unvaccinated group the increase is from 32 to 39 whereas for the vaccinated group the rise is from 9 to 20. A first approximation from this figure suggests via a coarse approximation that the effectiveness in preventing death is lower - 49% [simply from (39-20)/39] - than what was given in the main body of the article (72%). Arguably, the most surprising outcome from this

analysis is the increase of deaths by 7 due to Covid-19 counted in the unvaccinated cohort, especially considering that the individuals were free of Covid-19 symptoms on the day of vaccination.

One clarifying comment on how the analysis proceeded is also provided in the caption of Dagan et al.'s Table S5, stating "This analysis does not allow vaccinated controls to re-enroll as exposed, so individuals do not contribute time to both study groups." While this sentence is a bit confusing to read at first, it is descriptive and necessary. Of the 259,941 controls who rolled out of the unvaccinated cohort, 86,601 (or 33.3%) were "rolled into" the vaccinated cohort and re-matched with a new unvaccinated person.

There is, however, an unanswered, and very important, question. How many of the 7 individuals whose deaths were attributed to the unvaccinated cohort in the "delayed-censoring sensitivity analysis" were originally attributed to the vaccinated cohort in the analysis of the main body of the article? To wit, if all 7 were re-matched and counted as part of the 9 deaths in Figure 2, one would reach the conclusion that 18 new deaths were counted in the vaccinated cohort due to the 18-day delay in censoring controls. On the other hand, if none of these 7 individuals were "rolled into" the vaccinated cohort in the main analysis, there would be only 11 additional deaths due to the delay in censoring. In the context of the present study, such a difference is large. The authors make the assumption that exposure to the vaccine has no effect in the first several days after vaccination, so in the supplement they count the 7 in the unvaccinated cohort. This is arguably an example of aforementioned circular logic. Aside from this ambiguity, there is a possibility that suggests that the vaccine itself might be a confounding factor shortly after vaccination. For example, recent data from Bernal et al. show that study participants vaccinated with BNT162b2 "had a higher odds of testing positive for Covid-19 in the first 9 days after vaccination."[6]

The number of deaths counted in the vaccinated cohort is more than doubled (increasing from 9 to 20), but this is not surprising given the rules regarding censoring. The authors assessed the difference in outcome by stating in reference to all outcomes, "This analysis showed results similar to those of the main analysis, which suggests that any such bias was small in our analysis." Yet a simple measure of efficacy of the vaccine in preventing death over the 44-day period shows a reduction from 72% to 49% in the analysis accounting for bias. Also significant is that the error bars shown by the authors overlap when the delayed-censoring approach is used whereas they do not overlap in the main analysis. The words "small" and "similar" are subjective, and such an assertion would be certainly strengthened by a more objective rationale.

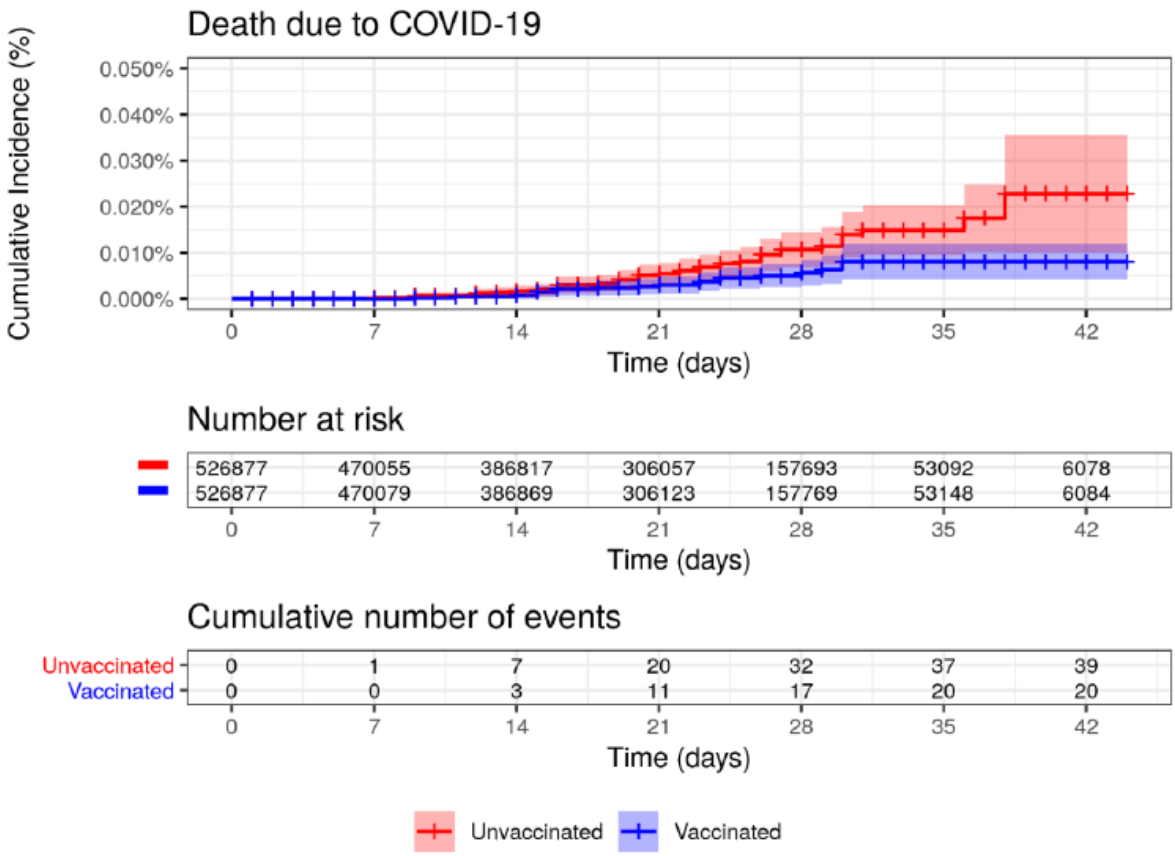

Legend: Cumulative incidence curves (one minus the Kaplan-Meier risk) for the COVID-19 related death outcome when delaying the censoring of vaccinated controls. Shaded areas are 95% confidence intervals. The tables below the curve show the number at risk at each time point and the cumulative number of events. Vaccine effectiveness estimates are included in Table S5.

Figure 4. A copy of figure S7E from Dagan et al. which indicated cumulative deaths of 39 for the unvaccinated cohort and 20 for the vaccinated cohort. (permission to reproduce requested)

One purpose of the present communication is to demonstrate that a different approach, with several advantages, may be utilized to approximate the effectiveness of the vaccine described in Dagan et al. in preventing deaths. Rather than choosing a specific number of days to delay censoring, one may compute the apparent effectiveness of a "null vaccine" given the rules restricting vaccination only to individuals free of Covid-19 symptoms. This value quantifies the effect of the bias alone and may then be compared to the measured outcome, for example as a ratio, and used as a means of assessment of efficacy. However, in order to implement the approach, one must also utilize the information regarding those rolling out of the unvaccinated cohort over time. Figure 1 does not provide the time-wise distribution for those rolling *out* of the unvaccinated cohort nor does any information in the remainder Dagan et al. article. The prospect that the Kaplan-Meier tables does not fully inform the reader regarding the distribution of those rolling out requires a lengthier argument and is addressed in the discussion herein.

A second purpose of the present communication is to show that the distribution has a substantial influence on this "apparent effectiveness," and therefore it should be included in future similar studies of this nature. A broader point is that data transparency is essential not only for the scientific community but also for decision-makers and the general public.

To begin, an approximation of the number of matches for each day from Dagan et al.'s Fig S1 (see also Figure 1) is depicted in Figure 5. For the purposes of explaining the effect of the bias and the approach, it is not necessary to utilize the exact values of that of Dagan et al. Nonetheless, the distribution was established to result in a total of 596,618 matches.

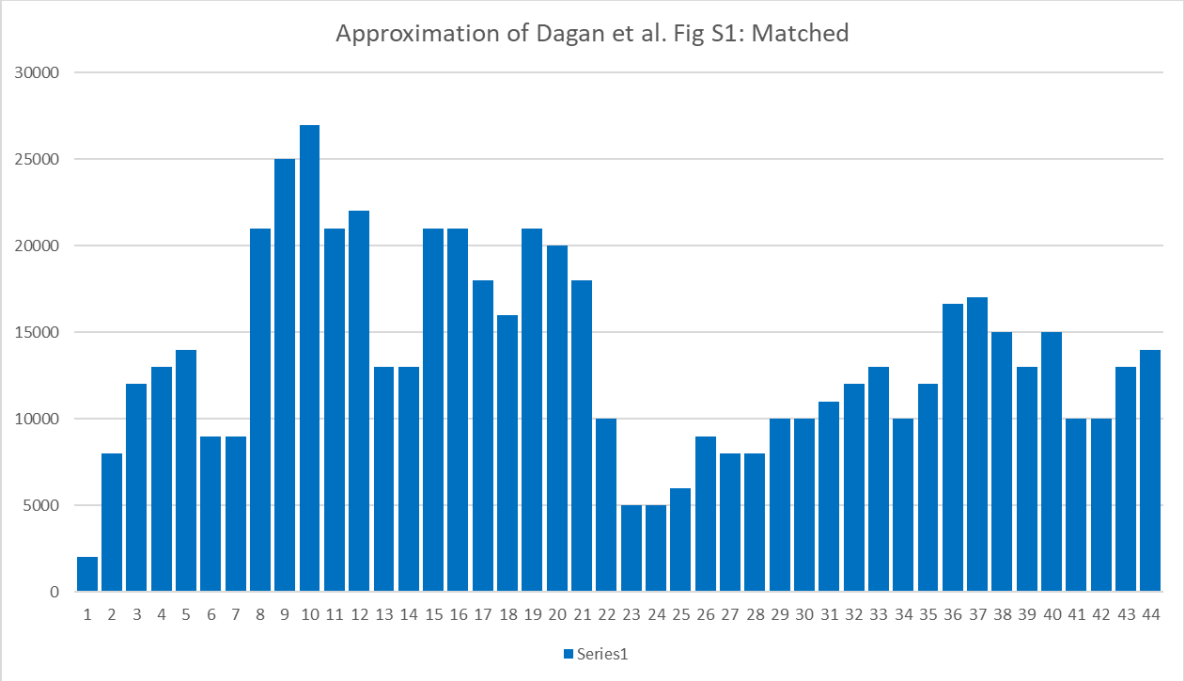

Figure 5. Approximation of data from Dagan et al. (Fig. S1), matched persons only.

A cumulative number of matches provides some added insight into how their study was conducted. This is given in Figure 6. It is important to understand that the number of people in each group for any given day differs from the cumulative number of matches. This occurs to due to those who are roll out of each group (i.e., are censored).

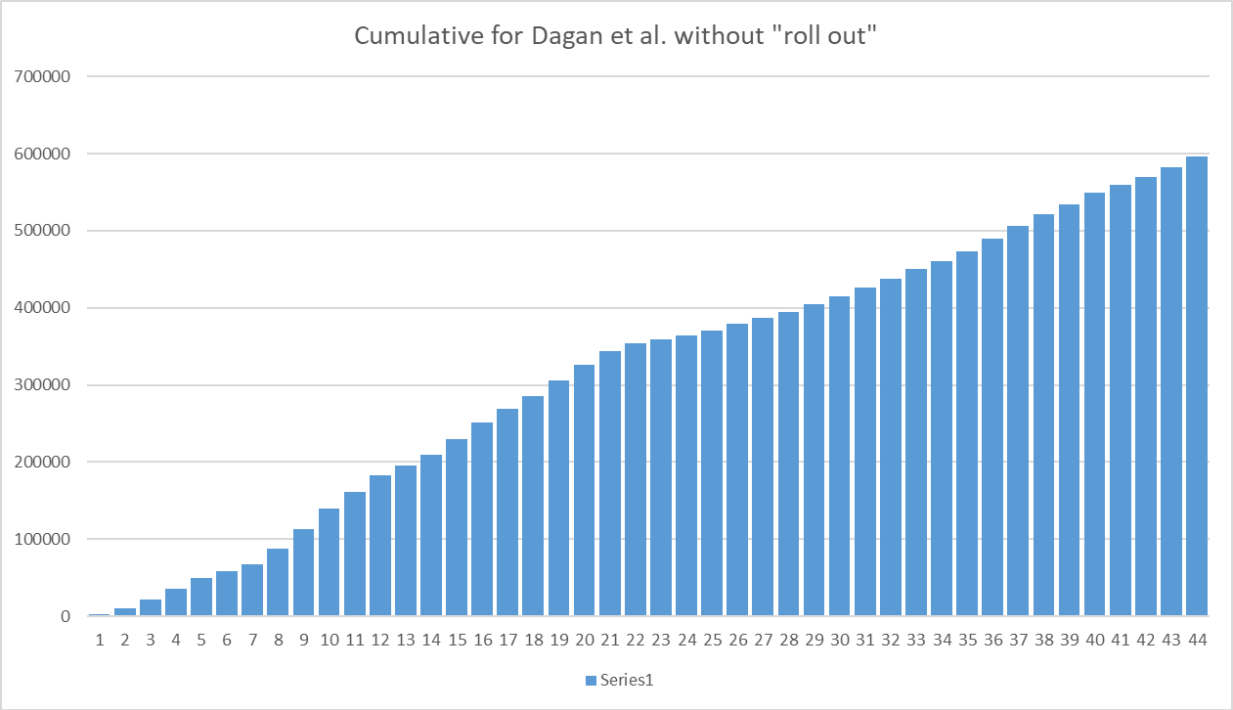

Figure 6. Cumulative number of matches for Dagan et al., by day if there were no censored matches.

Results Section I. Assessment of bias for a "neutral" distribution of those rolling out of the unvaccinated group

To assess the effect of individuals rolling out of the unvaccinated cohort, one must first establish a time-wise distribution of those who start in the unvaccinated group but become vaccinated within the 44-day study period and roll out of that group. Ideally, one would use the actual data, but for the purpose of demonstrating the effect on bias, one may also propose a specific time-wise distribution and follow it to an outcome. A reasonable place to begin is to simply assume that a constant multiplied by the number of matches from the previous day would yield a reasonable, and neutral, estimate of the number of those rolling out of the unvaccinated cohort. As expected, the number works out be slightly over 44% in order to achieve the cumulative total of 259,941 who roll out. This distribution is shown in Figure 7. As one would anticipate, the distribution in this figure resembles Figure 5, except that the columns appear shifted one day to the right and the vertical scale is altered. Since the number of deaths due to Covid-19 and other causes is expected to be much smaller than the 259,941 who roll out due to vaccination, the effect of those individuals on the bias are neglected in the present discussion.

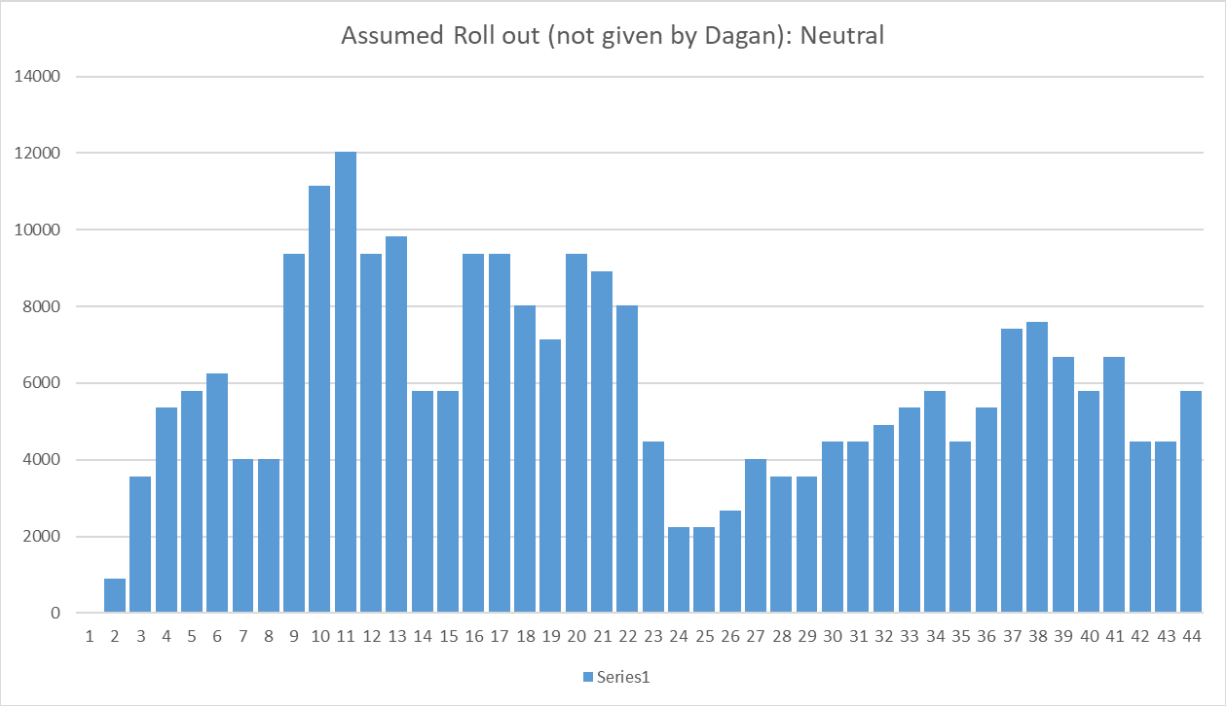

Figure 7. Assumed time-wise distribution of those rolling out of the unvaccinated cohort – neutral case.

With this "neutral" distribution treated as 'known,' one may then compute the number of people in each cohort for any given day of the trial. This result is given for the neutral distribution in Figure 8.

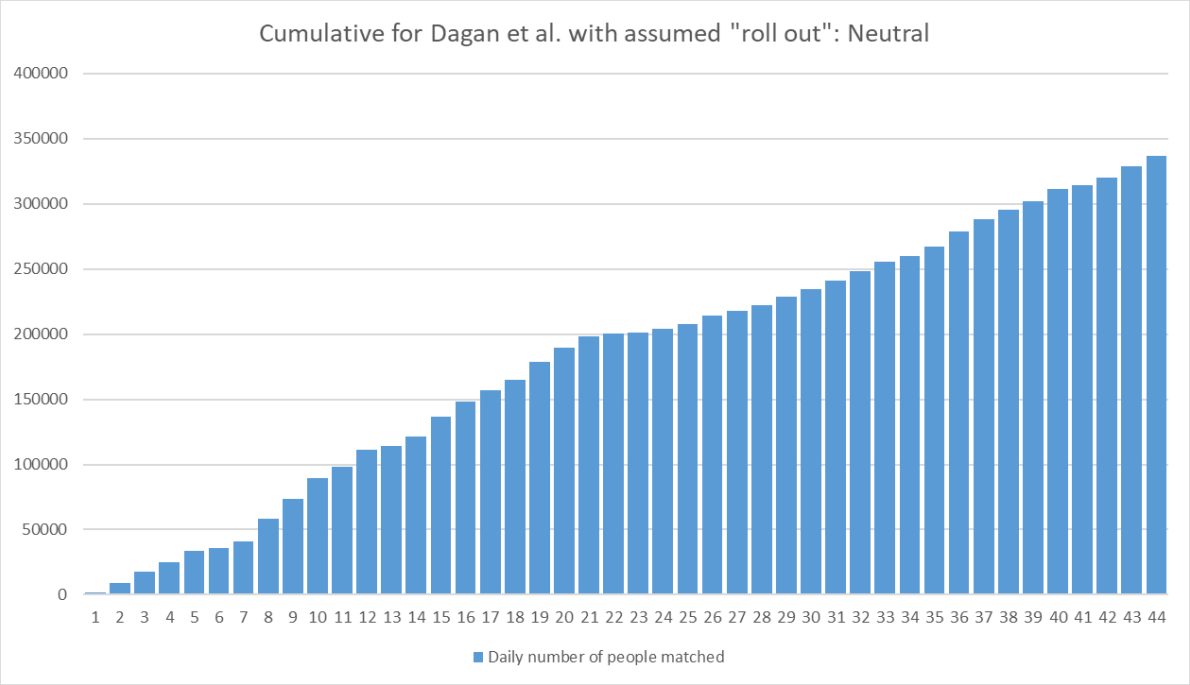

Figure 8. Number of people in each group by day for the neutral distribution.

The next step in the process of demonstrating the bias is matching the number of symptomatic persons in the unvaccinated cohort to the value of 3,607 reported by Dagan et al. One could set parameters which allow for sub-groups or time intervals to have different likelihoods of having symptomatic Covid-19 to fully quantify bias. However, it is sufficient for the purpose of demonstrating bias to assume that the number of symptomatic people is simply proportional to the number of persons in each group on any given day. This approach has the advantage of being easy to understand and straightforward. To determine the value used for the proportion, it is helpful to set up a spreadsheet. This is given in Figure 9. The value at the top of the sixth column (0.004718, shown in red) was determined to provide the proper proportion. This value is multiplied by number of matches for each day, resulting in the values in the sixth column. These values are summed to yield the total of 3,607 for the unvaccinated cohort, listed at the bottom of that column. Another, arguably slightly more appropriate, way to approach the number of symptomatic people might be to subtract those who were already counted as symptomatic on previous days. This approach was tested and showed a negligible difference due to the low number of symptomatic people compared to the total number of matches in the study (less than one percent). Thus, the simpler approach is used herein, and its outcome is shown in Figure 10. Once more, it is not surprising that the distribution resembles that of Figure 8 since a simple proportional constant was used. Only the vertical scale is different.

Figure 9. Spreadsheet used to prepare plots for the demonstration of bias for the neutral case.

			0.44616		0.0004718	1	0.02408		0.02408	Fraction of vaccinated	
	Matched number	Cumulative people	Number vaccianted	Number of people	Number of symptomatic	Number of new symptomatic	Deaths for 18-day-delay	Percent of prior-day "controls"	Deaths for 18-day-delay	whose match remains	Deaths for 18-day-delay
Trial	of people	who are matched	"neutral" approach	in each matched group	people for each day	people for each day	Unvaccinated	vaccinated each day	Vaccinated	in group	Vaccinated
Day	(control and exposed)	(before "rolling")	(Constant = 0.44616)		(unvaccinated group)	(vaccinated group)	(constant set to yield 32)	(This is the "rolling"	*before roll of match*	(before death is counted)	*after roll*
	(est. from Fig. S1 of Dagan)				(constant set to yield 3607)	Effectiveness = nil when 1.0)		portion of the cohort)		(No "rolling" on Day 1)	of matched unvaxed
1	2000	2000	0	2000	0.0	0.0	0.00	0.00	0	N/A	
2	8000	10000	892	9108	0.9	0.9	0.00	44.62	0.00	0.04	
3	12000	22000	3569	17538	4.3	4.3	0.00	39.19	0.00	0.08	
4	13000	35000	5354	25184	8.3	8.3	0.00	30.53	0.00	0.12	
5	14000	49000	5800	33384	11.9	11.9	0.00	23.03	0.00	0.16	
6	9000	58000	6246	36138	15.8	15.8	0.00	18.71	0.00	0.21	
7	9000	67000	4015	41123	17.0	17.0	0.00	11.11	0.00	0.25	
8	21000	88000	4015	58107	19.4	19.4	0.00	9.76	0.00	0.28	
9	25000	113000	9369	73738	27.4	27.4	0.00	16.12	0.00	0.30	
10	27000	140000	11154	89584	34.8	34.8	0.00	15.13	0.00	0.36	
11	21000	161000	12046	98538	42.3	42.3	0.00	13.45	0.00	0.41	
12	22000	183000	9369	111168	46.5	46.5	0.00	9.51	0.00	0.47	
13	13000	196000	9816	114353	52.4	52.4	0.00	8.83	0.00	0.51	
14	13000	209000	5800	121553	54.0	54.0	0.00	5.07	0.00	0.55	
15	21000	230000	5800	136753	57.3	57.3	0.00	4.77	0.00	0.56	
16	21000	251000	9369	148383	64.5	64.5	0.00	6.85	0.00	0.58	
17	18000	269000	9369	157014	70.0	70.0	0.00	6.31	0.00	0.61	
18	16000	285000	8031	164983	74.1	74.1	0.00	5.11	0.00	0.64	
19	21000	306000	7139	178844	77.8	77.8	0.00	4.33	0.00	0.66	
20	20000	326000	9369	189475	84.4	84.4	0.02	5.24	0.02	0.67	0.00
21	18000	344000	8923	198552	89.4	89.4	0.10	4.71	0.10	0.69	0.01
22	10000	354000	8031	200521	93.7	93.7	0.20	4.04	0.20	0.71	0.02
23	5000	359000	4462	201059	94.6	94.6	0.29	2.23	0.29	0.72	0.05
24	5000	364000	2231	203829	94.9	94.9	0.38	1.11	0.38	0.72	0.08
25	6000	370000	2231	207598	96.2	96.2	0.41	1.09	0.41	0.72	0.10
26	9000	379000	2677	213921	97.9	97.9	0.47	1.29	0.47	0.71	0.13
27	8000	387000	4015	217905	100.9	100.9	0.66	1.88	0.66	0.71	0.20
28	8000	395000	3569	222336	102.8	102.8	0.84	1.64	0.84	0.71	0.30
29	10000	405000	3569	228767	104.9	104.9	1.02	1.61	1.02		0.42
30	10000	415000	4462	234305	107.9	107.9	1.12	1.95	1.12		0.52
31	11000	426000	4462	240844	110.5	110.5	1.26	1.90	1.26		0.64
32	12000	438000	4908	247936	113.6	113.6	1.30	2.04	1.30		0.71
33	13000	451000	5354	255582	117.0	117.0	1.38	2.16	1.38		0.78
34	10000	461000	5800	259782	120.6	120.6	1.55	2.27	1.55		0.90
35	12000	473000	4462	267320	122.6	122.6	1.69	1.72	1.69		1.03
36	16618	489618	5354	278584	126.1	126.1	1.78	2.00	1.78		1.13
37	17000	506618	7414	288170	131.4	131.4	1.87	2.66	1.87		1.23
38	15000	521618	7585	295585	136.0	136.0	2.03	2.63	2.03		1.37
39	13000	534618	6692	301893	139.5	139.5	2.15	2.26	2.15		1.49
40	15000	549618	5800	311093	142.4	142.4	2.26	1.92	2.26		1.60
41	10000	559618	6692	314400	146.8	146.8	2.28	2.15	2.28		1.64
42	10000	569618	4462	319939	148.3	148.3	2.28	1.42	2.28		1.65
43	13000	582618	4462	328477	150.9	150.9	2.32	1.39	2.32		1.66
44	14000	596618	5800	336677	155.0	155.0	2.36	1.77	2.36		1.68
	596618		**259941**		**3607**	**3607**	**32.0**				**19.3**
	See Dagan Fig 1		See Dagan Fig 1		See Dagan Fig. 2B	Artificially set to same as	Matches Fig. 2E				This value is the expected
						unvaccinated group.	This ^^ is the total number of deaths				number of deaths for
						Could be altered to match	in the unvaccinated group.				comparison to 9 in Fig. 2E
						Fig. 2B if the constant in row 1	The constant in row 1 is altered				for the given distribution of
						is set equal to 0.6623	until this value is reached.				those "rolling out" of the
											unvaccinated cohort.

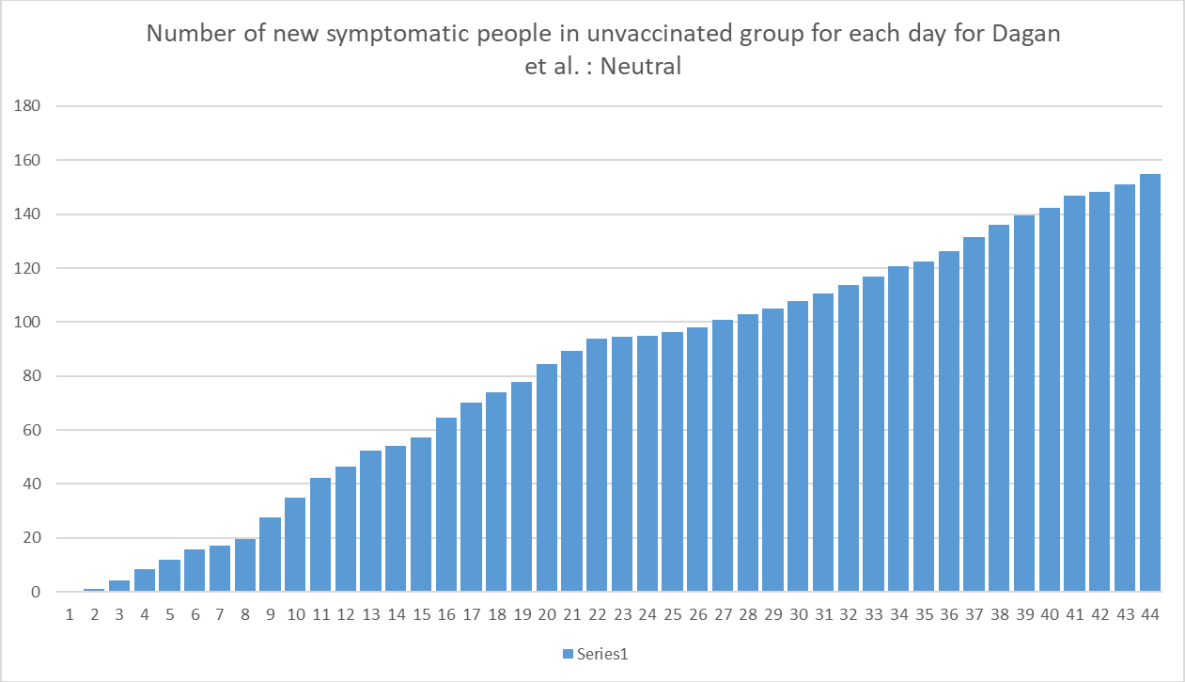

Figure 10. New symptomatic people in the unvaccinated cohort each day, as projected using the 'neutral' model. If there were no rolling out due to vaccination of matches, this would also represent the number of symptomatic vaccinated persons (with a "null vaccine").

At this point, one could apply a correction factor in the form of either a constant or a more sophisticated expression, to match the number of symptomatic people in the vaccinated group as well. However, the focus in the present process is to establish the effect of a "null vaccine," and so the number of symptomatic "vaccinated" people is taken to be equal to the values for those who are unvaccinated. Thus, the values in Column 7 match those of Column 6.

The next step is to estimate both the proportion of symptomatic people who die and the time after which symptoms appear that death would occur. One might develop a sophisticated model by accounting for factors in various subgroups or by applying a possible distribution of delay times between appearance of symptoms and death. Here, it is sufficient to demonstrate the effect of bias with an exceptionally straightforward approach. Simply take the number of people who are newly symptomatic on each day and assume that they each have equal likelihood of dying precisely 18 days later.

The choice of an 18-day differential between symptom appearance and death is analogous to the value utilized by Dagan et al. in their delayed censoring approach and can be supported on other grounds as well. First, consider the plot from Dagan et al. reproduced in Figure 11 which shows a histogram for the number of days from number of days from a PCR-positive test and hospitalization. The median is 10 days shorter than that for death (see Figure 3). If one assumes that symptoms, which would prevent vaccination, initiate 8 days prior to hospitalization, one computes an 18-day differential. Early data in the pandemic from Wuhan specifically state their population had an 18.5-day median delay between illness onset and death.[7] Faes et al. report a mean time between symptom onset and hospitalization in Belgium of 5.74 days and that the median hospital length-of-stay for patients that die is 12.2 days for the working-age population, yielding a total time from symptom onset to death of 17.9 days. So, an assumed 18-day differential between symptom onset and death may be justifiable. Faes et al. do

continue that the median hospital length-of-stay is quite a bit lower (5.7) for elderly patients and there is also variability from country to country.[8]

Even though the choice of an 18-day differential between symptom onset and death appears a reasonable estimate, one might also justify other choices. What would the outcome be if the delay time for delayed censoring were selected by Dagan et al. to be 14 days or 22 days? The histogram presented in Figure 3 shows that a non-negligible number of deaths occurred more than 18 days after the first PCR test was given. It is quite reasonable to expect that deaths may be undercounted in the vaccinated group, despite their effort applied via the delayed-censoring analysis. Additional information regarding the patient population, especially among those whose deaths were newly counted in the sensitivity analysis with delayed censoring, would be helpful to determine whether an 18-day delay is an appropriate choice. In addition, data analogous to that compiled by Faes et al. specific to Israel, if available, would be a welcome addition to inform the choice.

In a more advanced analysis using the null vaccine approach, one might also venture to utilize a time-wise distribution of death occurrence for any given symptomatic person, as opposed to choosing a singular delay (18 days in this analysis) for all deaths to occur. Herein, the advantage once more is simplicity, and it serves the purpose of illustrating the effect of the bias brought about due to individuals being free of Covid-19 symptoms on the day they chose to obtain the vaccine.

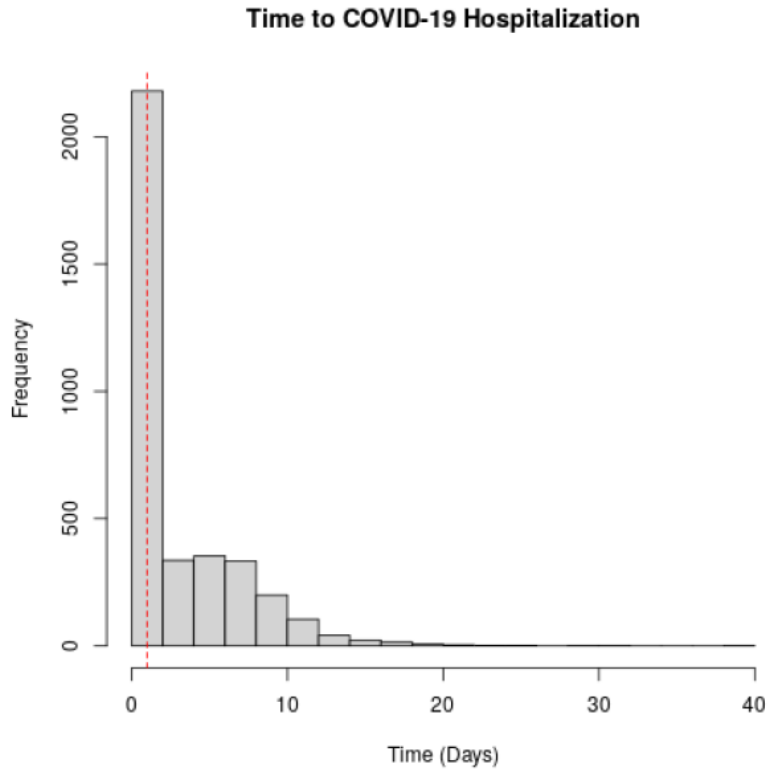

Figure 11. This histogram shows "the distribution of time (in days) from the date the first positive PCR test was sampled until incidence of the different outcomes: hospitalization due to COVID-19 and death due to COVID-19. A dashed vertical red line in each plot shows the median." The median time between a PCR-positive test and hospitalization is one day. (permission to reproduce requested)

To determine the proper proportionality constant to apply to the number of symptomatic people, one can utilize the value for the total number of deaths in the control group (32.0). The proportionality constant works out to be 0.02408 for this 'neutral' case, and is shown in red above the eighth column in Figure 9. Since the number of deaths is low, it is useful to include "fractional deaths" to maintain accuracy. It may be helpful to consult once more the number of newly symptomatic people in the unvaccinated group from the same spreadsheet, and this is given in Figure 10. It should make sense that the number grows with time since the number of matches also grow monotonically under the assumed 'neutral' circumstances. Given that there is a necessary delay between symptom appearance and death, one might expect that some of those who become symptomatic near the end of the trial die after the follow-up has ended after Day 44.

The number of deaths, computed under this set of assumptions, on each given day is given in Figure 12. Due to the proportionality constant applied to symptoms and the 18-day delay, the shape of the curve resembles that of Figure 10 except that there is a pronounced (18-day) shift to the right and the vertical scale has changed. If one sums all the values represented in the bars of Figure 12, the result is 32.0 deaths.

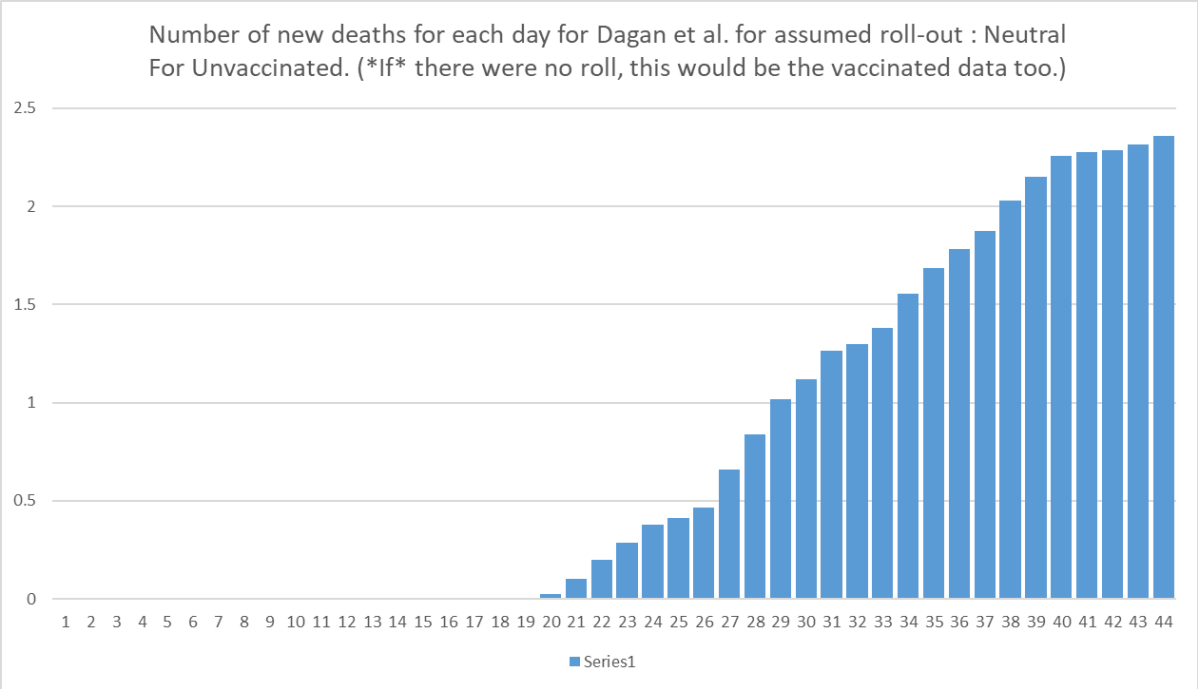

Figure 12. Number of deaths per day given fractionally in the unvaccinated cohort when roll out of due to vaccination is neutral.

At this point in the procedure, the net effect of the bias can finally be addressed. The fraction of individuals who roll out of the unvaccinated cohort on any given day can be found by dividing the number of people who roll out on a given day by the number of people in that cohort on the *prior* day. (It is appropriate to use the value from the prior day for the denominator because a person cannot both join the unvaccinated cohort and roll out of it on the same day.) This fraction also represents the likelihood that follow-up will end for any person in the vaccinated cohort, whether symptomatic or not,

as well. This fraction is given in terms of a percentage in Column 9. It is notable that the percentage diminishes with time as the denominator (i.e., the number of matches) grows with time.

The fraction of people who remain in the vaccinated cohort for any given day can be simply computed as one minus the fraction who roll out. Then, to compute an estimate of the possibility of remaining in the cohort for multiple days, one can simply multiply those fractions together, following simple laws of probability. For example, suppose a person is in a group on Day 1. Let's then assume there is a 60% chance of that person remaining in the group through Day 2, a 70% chance of remaining in the group through Day 3, and a 50% chance of the person remaining in that group through Day 4. One can compute the likelihood of the person remaining in the group as 0.60* 0.70 *0.50 = 0.21. In other words, there would be a 21 percent chance that person would still be in the group after Day 4 in this brief example.

For the present situation, there is an 18-day progression from symptom appearance to death, so there are eighteen terms multiplied together instead of three in the above example. This fraction is shown in the eleventh column for each day of interest. Note that the values do not extend past Day 26 since the assumption is made that an 18-day difference between symptom appearance and death is in place. In column 12, the number of deaths is computed for the vaccinated cohort with the bias taken into account. The difference is depicted in Figure 13.

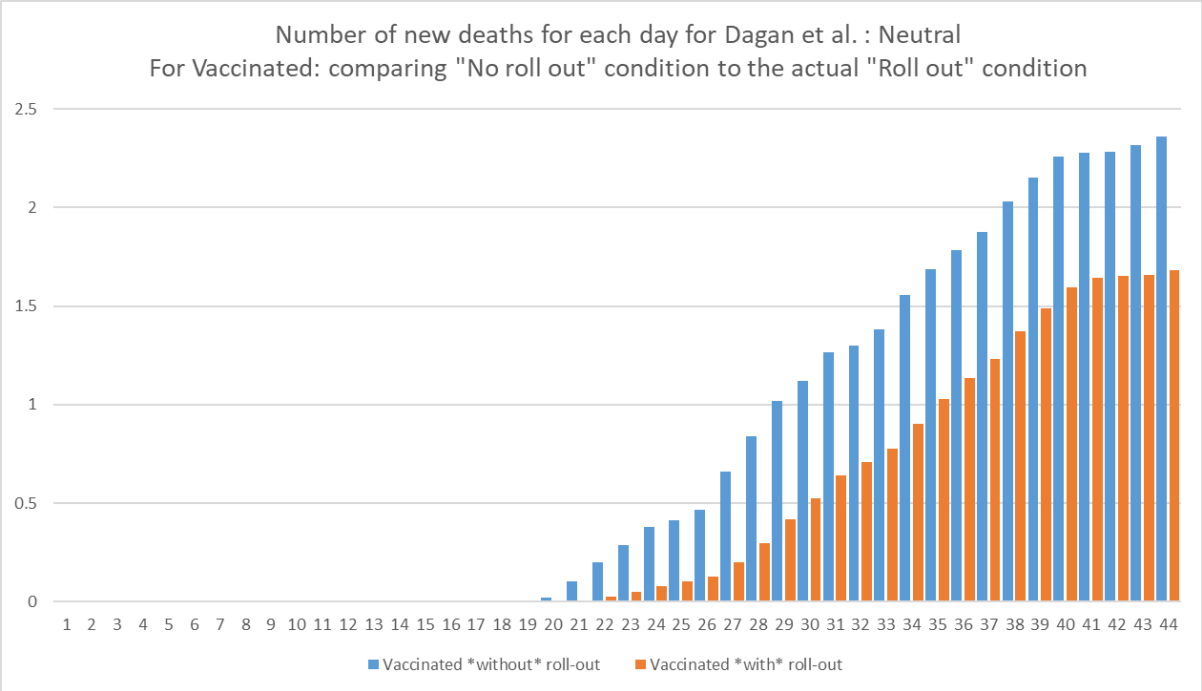

*Figure 13. The effect of the bias due to rolling out by unvaccinated persons is illustrated for the vaccinated cohort for this "neutral" case. Deaths per day are shown. The sum of the *with* roll-out cases (orange) is 19.3 rather than 32.0.*

By summing the (orange) columns for the vaccinated cohort in Figure 13, the expected value for a null vaccine under the conditions where selection bias due to informative censoring is present can be computed, under the assumption of a uniform 18-day delay between symptom appearance and death. The summed output is 19.3 shown also in the spreadsheet of Figure 9 at the bottom of Column 13. This value may then be compared to the actual value of 9 found in the study. The ratio of (19.3-9)/19.3 yields

a coarse approximation of 54% efficacy for the vaccination effectiveness under the described circumstances using the null vaccine approach. Bearing in mind that the goal is to establish an "expected number of deaths" for comparison to the actual number (9), if one were to allow the vaccine to "take credit for" a lower value for the number of symptomatic people, the expected number of deaths would drop from 19.3 to a lower value, which in turn would lead to a *reduced* value of efficacy for preventing deaths.

Results, Section II. "Optimized" distribution to increase bias

One can also set up a time-wise distribution for those who roll out of the control group which leads to a considerably lower value for expected number of deaths for a null vaccine. For the purposes of this discussion, the distribution is described as "optimized" only because it would lead to the appearance that a "null vaccine" would increase the bias. One such example of a distribution is given in Figure 14. Here it can be seen that there are two relatively large groupings of those rolling out. Note that the vertical scale is larger than that of Figure 7. The spreadsheet associated with the output is given in Figure 15 and is analogous to the information given in Figure 9.

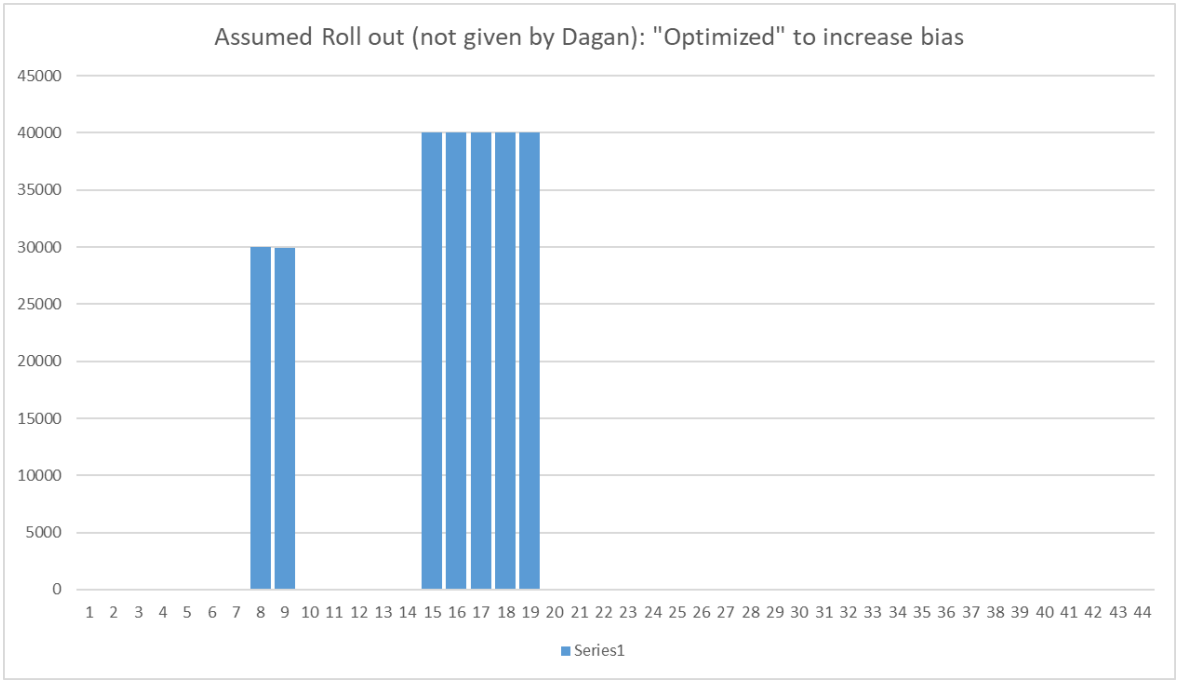

Figure 14. Assumed time-wise roll out distribution for a nearly "optimized" outcome in the sense that the expected number of deaths form a null vaccine would be lower than the neutral case.

Figure 15. Spreadsheet used to prepare plots for the demonstration of bias for the "optimized" case.

					0.00062118	1	0.02654		0.02654	Fraction of vaccinated	
	Matched number	Cumulative people	Number vaccianted	Number of people	Number of symptomatic	Number of symptomatic	Deaths for 18-day-delay	Percent of prior-day "controls"	Deaths for 18-day-delay	whose match remains	Deaths for 18-day-delay
Trial	of people	who are matched	"strategic" approach	in each matched group	people for each day	people for each day	Unvaccinated	vaccinated each day	Vaccinated	in group	Vaccinated
Day	(control and exposed)	(before "rolling")			(unvaccinated group)	(vaccinated group)	(constant set to yield 32)	(This is the "rolling"	*before roll of match*	(before death is counted)	*after roll*
	(est. from Fig. S1 of Dagan)				(constant = 0.0004718)	Effectiveness = nil when 1.0)		portion of the cohort)		(No "rolling" on Day 1)	of matched unvaxed people
1	2000	2000	0	2000	0.0	0.0	0.00	0.00	0.00	N/A	
2	8000	10000	0	10000	1.2	1.2	0.00	0.00	0.00	0.05	
3	12000	22000	0	22000	6.2	6.2	0.00	0.00	0.00	0.02	
4	13000	35000	0	35000	13.7	13.7	0.00	0.00	0.00	0.02	
5	14000	49000	0	49000	21.7	21.7	0.00	0.00	0.00	0.02	
6	9000	58000	0	58000	30.4	30.4	0.00	0.00	0.00	0.02	
7	9000	67000	0	67000	36.0	36.0	0.00	0.00	0.00	0.02	
8	21000	88000	30000	58000	41.6	41.6	0.00	44.78	0.00	0.02	
9	25000	113000	29941	53059	36.0	36.0	0.00	51.62	0.00	0.03	
10	27000	140000	0	80059	33.0	33.0	0.00	0.00	0.00	0.07	
11	21000	161000	0	101059	49.7	49.7	0.00	0.00	0.00	0.07	
12	22000	183000	0	123059	62.8	62.8	0.00	0.00	0.00	0.07	
13	13000	196000	0	136059	76.4	76.4	0.00	0.00	0.00	0.07	
14	13000	209000	0	149059	84.5	84.5	0.00	0.00	0.00	0.07	
15	21000	230000	40000	130059	92.6	92.6	0.00	26.84	0.00	0.07	
16	21000	251000	40000	111059	80.8	80.8	0.00	30.76	0.00	0.09	
17	18000	269000	40000	89059	69.0	69.0	0.00	36.02	0.00	0.14	
18	16000	285000	40000	65059	55.3	55.3	0.00	44.91	0.00	0.21	
19	21000	306000	40000	46059	40.4	40.4	0.00	61.48	0.00	0.39	
20	20000	326000	0	66059	28.6	28.6	0.03	0.00	0.03	1.00	0.00
21	18000	344000	0	84059	41.0	41.0	0.16	0.00	0.16	1.00	0.00
22	10000	354000	0	94059	52.2	52.2	0.36	0.00	0.36	1.00	0.01
23	5000	359000	0	99059	58.4	58.4	0.58	0.00	0.58	1.00	0.01
24	5000	364000	0	104059	61.5	61.5	0.81	0.00	0.81	1.00	0.01
25	6000	370000	0	110059	64.6	64.6	0.96	0.00	0.96	1.00	0.02
26	9000	379000	0	119059	68.4	68.4	1.10	0.00	1.10	1.00	0.02
27	8000	387000	0	127059	74.0	74.0	0.96	0.00	0.96	1.00	0.03
28	8000	395000	0	135059	78.9	78.9	0.87	0.00	0.87		0.06
29	10000	405000	0	145059	83.9	83.9	1.32	0.00	1.32		0.09
30	10000	415000	0	155059	90.1	90.1	1.67	0.00	1.67		0.11
31	11000	426000	0	166059	96.3	96.3	2.03	0.00	2.03		0.14
32	12000	438000	0	178059	103.2	103.2	2.24	0.00	2.24		0.15
33	13000	451000	0	191059	110.6	110.6	2.46	0.00	2.46		0.17
34	10000	461000	0	201059	118.7	118.7	2.14	0.00	2.14		0.20
35	12000	473000	0	213059	124.9	124.9	1.83	0.00	1.83		0.25
36	16618	489618	0	229677	132.3	132.3	1.47	0.00	1.47		0.31
37	17000	506618	0	246677	142.7	142.7	1.07	0.00	1.07		0.41
38	15000	521618	0	261677	153.2	153.2	0.76	0.00	0.76		0.76
39	13000	534618	0	274677	162.5	162.5	1.09	0.00	1.09		1.09
40	15000	549618	0	289677	170.6	170.6	1.39	0.00	1.39		1.39
41	10000	559618	0	299677	179.9	179.9	1.55	0.00	1.55		1.55
42	10000	569618	0	309677	186.2	186.2	1.63	0.00	1.63		1.63
43	13000	582618	0	322677	192.4	192.4	1.72	0.00	1.72		1.72
44	14000	596618	0	336677	200.4	200.4	1.81	0.00	1.81		1.81
	596618		**259941**		**3607**	**3607**	**32.0**				**12.0**
	See Dagan Fig 1		See Dagan Fig 1		See Dagan Fig. 2B	Artificially set to same as	Matches Fig. 2E				This value is the expected
						unvaccianted group.	This ^^ is the total number of deaths				number of deaths for
						Could be altered to match	in the unvaccinated group.				comparison to 9 in Fig. 2E
						Fig. 2B if constant = 0.6623	The constant in row 1 is altered				for the given distribution of
							until this value is reached.				those "rolling out" of the
											unvaccinated cohort.

The first item to examine for the optimized case is the resulting number of cumulative matches, which is given in Figure 16. The difference between this “optimized” output and that of the “neutral” case in Figure 8 is rather profound. One easy way to quantify the difference is to look specifically at Day 19. The neutral case shows 178,844 people in each group on Day 19 whereas the “optimized” case leads to only 46,059 in each group (see also the spreadsheet data in Figure 9 and Figure 15). Only those who are free of Covid-19 symptoms are permitted to roll out of the unvaccinated cohort, so, as one might anticipate, the proportion of those who are symptomatic in the unvaccinated cohort on that day would be higher for the optimized case.

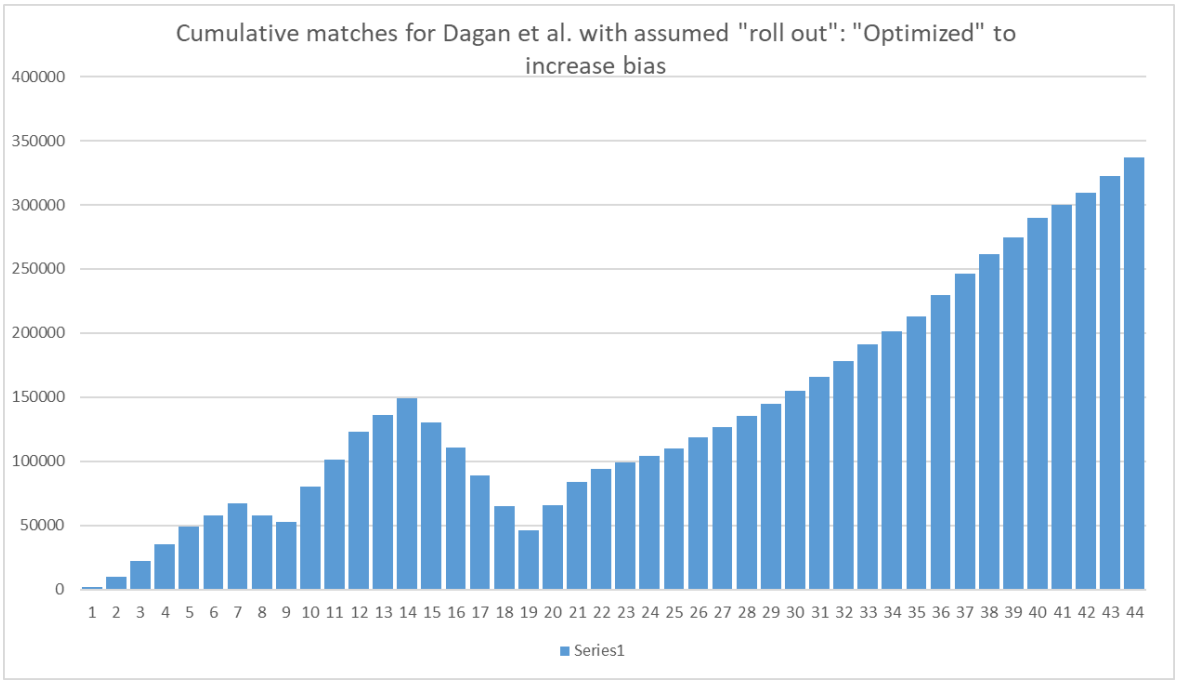

Figure 16. Number of people in each group by day for the “optimized” distribution.

In essence, the “optimized” distribution filters out many more of the people who were matched in the first portion of the study. If one proceeds to compute the number of symptomatic people in the unvaccinated group using the same constant of proportionality approach, one can produce Figure 17. Note that the actual constant used to compute the number of people in the unvaccinated cohort who become symptomatic is adjusted upward from 0.0004718 (from Figure 9) up to 0.00062118 (from Figure 15) so that the same value of 3,607 symptomatic patients in the control group is reached. In other words, the control output is used as a benchmark, or calibration parameter, for the present analysis.

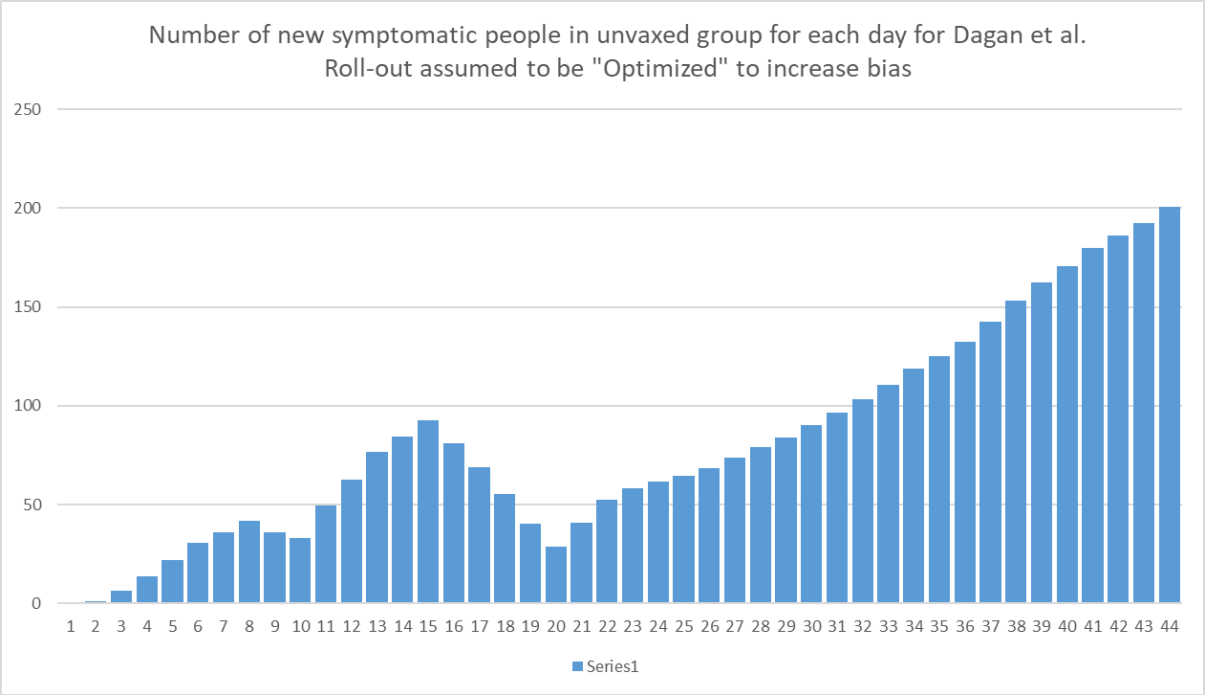

Figure 17. New symptomatic people in the unvaccinated cohort each day for the "optimized" roll out distribution. If there were no rolling out due to vaccination of matches, this would also represent the number of symptomatic vaccinated persons (with a "null vaccine").

The same uniform 18-day delay between symptom appearance and death for a fraction of individuals in the unvaccinated cohort is used in this demonstration to work out the bias. Note that, once more, the proportionality constant must be adjusted (from 0.02408 to 0.02654) to maintain 32.0 deaths in the unvaccinated group. The output is given in Figure 18. This same constant is then used with vaccinated group as well since a “null vaccine” is still assumed for this “optimized” distribution.

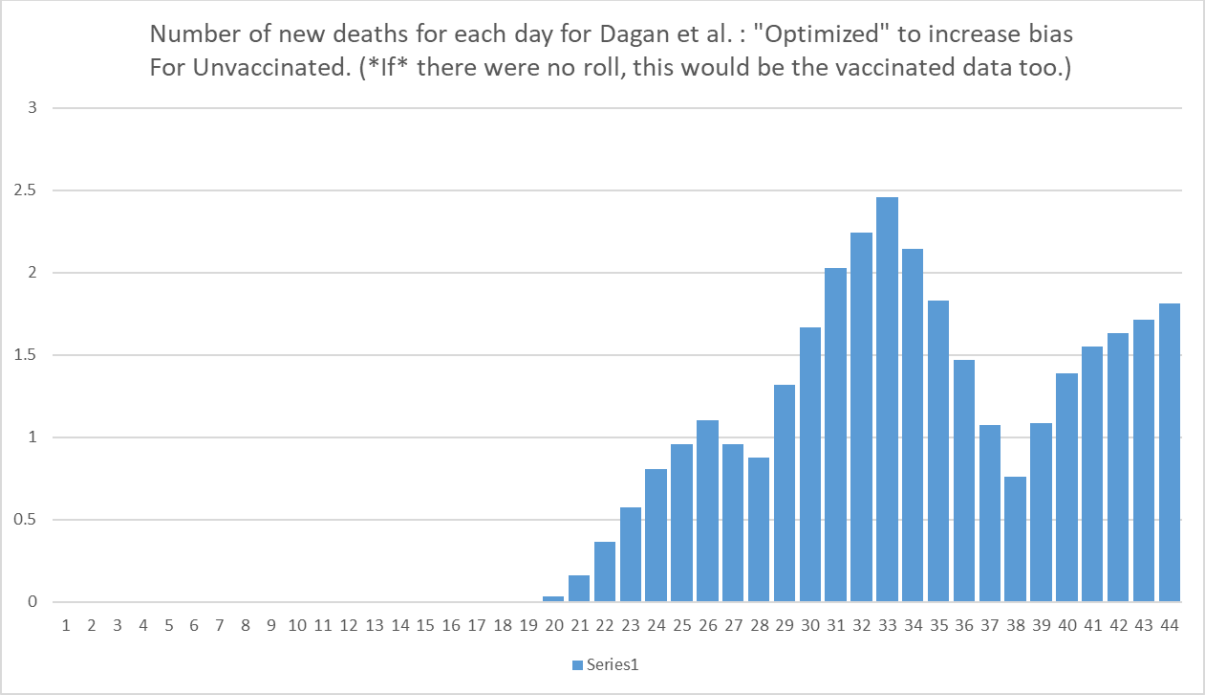

Figure 18. Number of deaths given fractionally in the unvaccinated cohort when roll out of due to vaccination is "optimized".

The same process is followed using the law of probabilities to eliminate those in the vaccinated group whose "match" rolls out of the unvaccinated group. With that, the effect for the "optimized" case using the same uniform 18-day delay between symptom initiation and death may be determined. The results, in the form of a comparison to illuminate the bias, are shown in Figure 19. Summing the total number of deaths for this "optimized" group leads to a value of only 12.0. In turn, a coarse estimate of the vaccine efficacy under this distribution and enumerated set of assumptions would be only 25 % [(12-9)/12)]. It is rather remarkable that such a strong influence on the "apparent effectiveness of a null vaccine" can be achieved without any assumptions being made about the health, age, or demographics of the people in the trial. Rather, simply the choice of the time-wise distribution of those rolling out, combined with simple laws of probability to provide an estimate, is demonstrated to play an outsized role on the interpretation of the outcome.

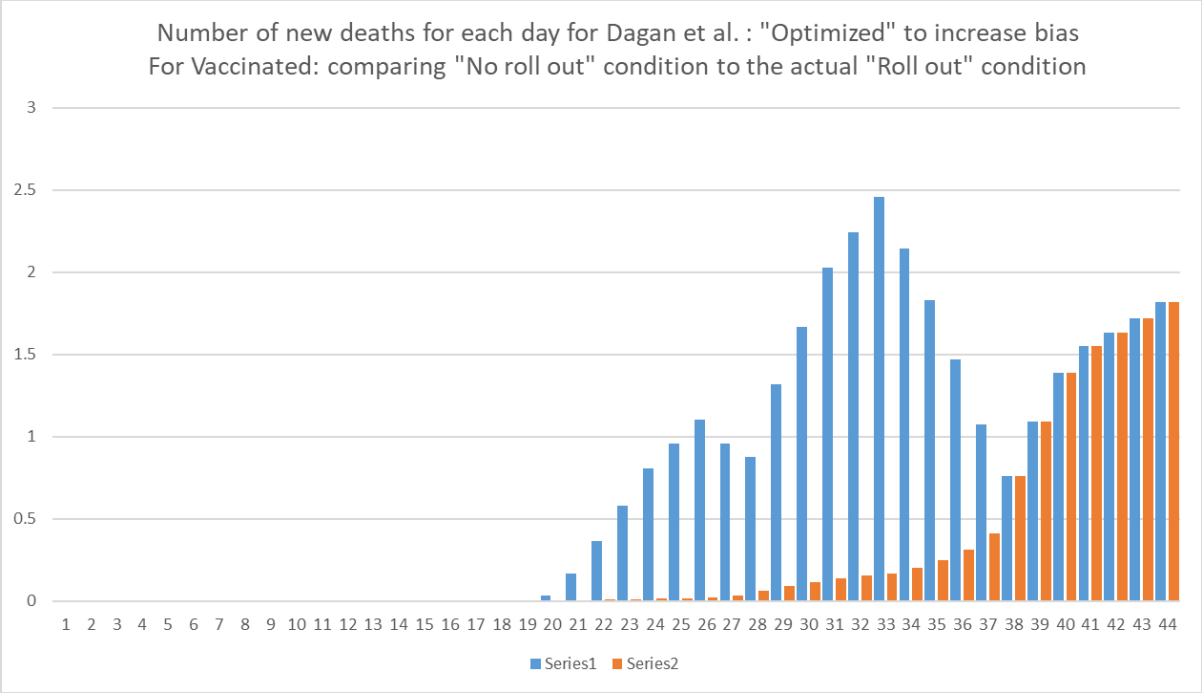

*Figure 19. The effect of the bias due to rolling out by unvaccinated persons is illustrated for the vaccinated cohort for this "optimized" case. Deaths 0per day are shown. The sum of the *with* roll-out cases (orange) is 12.0 rather than 32.0.*

It should also be pointed out that the bias could also be reduced substantially by adjusting the time-wise distribution of those rolling out of the trial. If all, or nearly all, of the rolling out occurs near the end of the trial, the expected value approaches 32. However, it is a mathematical certainty that the expected value must be lower than 32.

Campigotto and Weller specifically state that "the hazard ratio between the group of patients who remain on study and the group of patients who are informatively censored" is a secondary factor to consider (after the primary influence of the proportion who are informatively censored). For Dagan et al. one important consideration is that older individuals tended to be vaccinated sooner than younger people. In a Jerusalem Post story dated 7 January 2021, the Israeli Health ministry is quoted as saying that about 70 percent of those over the age of 60 had the first dose of vaccine.[9] Age is correlated closely with the likelihood of dying from Covid-19 as the CDC reports that those aged 65-74 are 95 times more likely to die than those aged 18-29. Meanwhile, the likelihood of attaining a case of Covid-19 is similar throughout all age groups.[10] With this in mind, consider an adjusted model which allows those in the first 14 (calendar) days of the trial to have twice the likelihood of dying compared to those who enter the trial on (calendar) day 15 or later should they become symptomatic. For the purpose of this model, also assume that symptoms occurrence remains unchanged, and additionally utilize the "optimized" time-wise distribution for the roll out. The spreadsheet under this model, adjusted to crudely take hazard ratio into account, is given in Figure 20.

					0.00062118	1	0.0188		0.0188	Fraction of vaccinated	
	Matched number	Cumulative people	Number vaccianted	Number of people	Number of symptomatic	Number of symptomatic	Deaths for 18-day-delay	Percent of prior-day "controls"	Deaths for 18-day-delay	whose match remains	Deaths for 18-day-delay
Trial	of people	who are matched	"strategic" approach	in each matched group	people for each day	people for each day	Unvaccinated	vaccinated each day	Vaccinated	in group	Vaccinated
Day	(control and exposed)	(before "rolling")			(unvaccinated group)	(vaccinated group)	(constant set to yield 32)	(This is the "rolling"	*before roll of match*	(before death is counted)	*after roll*
	(est. from Fig. S1 of Dagan)				(constant = 0.0004718)	Effectiveness = nil when 1.0)	HR is 2*days 1-14, 1*days 15-26	portion of the cohort)	HR is 2*days 1-14, 1*days 15-26	(No "rolling" on Day 1)	of matched unvaxed people
1	2000	2000	0	2000	0.0	0.0	0.00	0.00	0.00	N/A	
2	8000	10000	0	10000	1.2	1.2	0.00	0.00	0.00	0.05	
3	12000	22000	0	22000	6.2	6.2	0.00	0.00	0.00	0.02	
4	13000	35000	0	35000	13.7	13.7	0.00	0.00	0.00	0.02	
5	14000	49000	0	49000	21.7	21.7	0.00	0.00	0.00	0.02	
6	9000	58000	0	58000	30.4	30.4	0.00	0.00	0.00	0.02	
7	9000	67000	0	67000	36.0	36.0	0.00	0.00	0.00	0.02	
8	21000	88000	30000	58000	41.6	41.6	0.00	44.78	0.00	0.02	
9	25000	113000	29941	53059	36.0	36.0	0.00	51.62	0.00	0.03	
10	27000	140000	0	80059	33.0	33.0	0.00	0.00	0.00	0.07	
11	21000	161000	0	101059	49.7	49.7	0.00	0.00	0.00	0.07	
12	22000	183000	0	123059	62.8	62.8	0.00	0.00	0.00	0.07	
13	13000	196000	0	136059	76.4	76.4	0.00	0.00	0.00	0.07	
14	13000	209000	0	149059	84.5	84.5	0.00	0.00	0.00	0.07	
15	21000	230000	40000	130059	92.6	92.6	0.00	26.84	0.00	0.07	
16	21000	251000	40000	111059	80.8	80.8	0.00	30.76	0.00	0.09	
17	18000	269000	40000	89059	69.0	69.0	0.00	36.02	0.00	0.14	
18	16000	285000	40000	65059	55.3	55.3	0.00	44.91	0.00	0.21	
19	21000	306000	40000	46059	40.4	40.4	0.00	61.48	0.00	0.39	
20	20000	326000	0	66059	28.6	28.6	0.05	0.00	0.05	1.00	0.00
21	18000	344000	0	84059	41.0	41.0	0.23	0.00	0.23	1.00	0.00
22	10000	354000	0	94059	52.2	52.2	0.51	0.00	0.51	1.00	0.01
23	5000	359000	0	99059	58.4	58.4	0.82	0.00	0.82	1.00	0.02
24	5000	364000	0	104059	61.5	61.5	1.14	0.00	1.14	1.00	0.02
25	6000	370000	0	110059	64.6	64.6	1.35	0.00	1.35	1.00	0.02
26	9000	379000	0	119059	68.4	68.4	1.56	0.00	1.56	1.00	0.03
27	8000	387000	0	127059	74.0	74.0	1.35	0.00	1.35	1.00	0.05
28	8000	395000	0	135059	78.9	78.9	1.24	0.00	1.24		0.09
29	10000	405000	0	145059	83.9	83.9	1.87	0.00	1.87		0.13
30	10000	415000	0	155059	90.1	90.1	2.36	0.00	2.36		0.16
31	11000	426000	0	166059	96.3	96.3	2.87	0.00	2.87		0.20
32	12000	438000	0	178059	103.2	103.2	3.18	0.00	3.18		0.22
33	13000	451000	0	191059	110.6	110.6	1.74	0.00	1.74		0.12
34	10000	461000	0	201059	118.7	118.7	1.52	0.00	1.52		0.14
35	12000	473000	0	213059	124.9	124.9	1.30	0.00	1.30		0.18
36	16618	489618	0	229677	132.3	132.3	1.04	0.00	1.04		0.22
37	17000	506618	0	246677	142.7	142.7	0.76	0.00	0.76		0.29
38	15000	521618	0	261677	153.2	153.2	0.54	0.00	0.54		0.54
39	13000	534618	0	274677	162.5	162.5	0.77	0.00	0.77		0.77
40	15000	549618	0	289677	170.6	170.6	0.98	0.00	0.98		0.98
41	10000	559618	0	299677	179.9	179.9	1.10	0.00	1.10		1.10
42	10000	569618	0	309677	186.2	186.2	1.16	0.00	1.16		1.16
43	13000	582618	0	322677	192.4	192.4	1.22	0.00	1.22		1.22
44	14000	596618	0	336677	200.4	200.4	1.29	0.00	1.29		1.29
	596618		**259941**		**3607**	**3607**	**32.0**				**8.9**
	See Dagan Fig 1		See Dagan Fig 1		See Dagan Fig. 2B	Artificially set to same as	Matches Fig. 2E				This value is the expected
						unvaccianted group.	This ^^ is the total number of deaths				number of deaths for
						Could be altered to match	in the unvaccinated group.				comparison to 9 in Fig. 2E
						Fig. 2B if constant = 0.6623	The constant in row 1 is altered				for the given distribution of
							until this value is reached.				those "rolling out" of the
											unvaccinated cohort.

Figure 20. Spreadsheet used to demonstrate bias for the "optimized" case with those entering the study from Days 1-14 assumed to have double the chance of dying from Covid-19, compared to those entering later.

This model demonstrates that a specific time-wise distribution of those rolling out combined with a crude, but justifiable, age-associated model can lead to a value of 8.9 expected deaths in the "vaccinated" group - even if the vaccine does nothing at all (i.e., it is a null vaccine). Since this value is very nearly the 9 deaths found in the actual case, this analysis provides some credence to the notion that Dagan et al.'s findings might actually support the notion that the vaccine they studied has no effect of deaths which occur due to Covid-19. To produce this output, note that the constant above columns 9 and 11 is 0.0188. This value is multiplied by 2 for Days 1-14 and multiplied by 1 for Days 15 onward. Otherwise, the process is the same which was followed for the prior two examples in Figure 9 and Figure 15.

The underlying reason can perhaps be better understood with reference to Figure 21. The reduction of the multiplier from 2 down to 1 is observable in the "No roll out" dataset in that there is a marked step down from Day 32 to Day 33, in contrast with the data shown in Figure 19. This outcome, reducing the number of expected deaths from 12.0 to 8.9, also exemplifies the finding of Campigotto and Weller that hazard ratio can is a secondary factor in determining the influence of selection bias due to informative censoring. A point of emphasis is that this outcome does not mean that the actual vaccine effectiveness studied by Dagan et al. is definitively zero; rather, it suggests that additional transparency of the data is required in order to rule such a possibility.

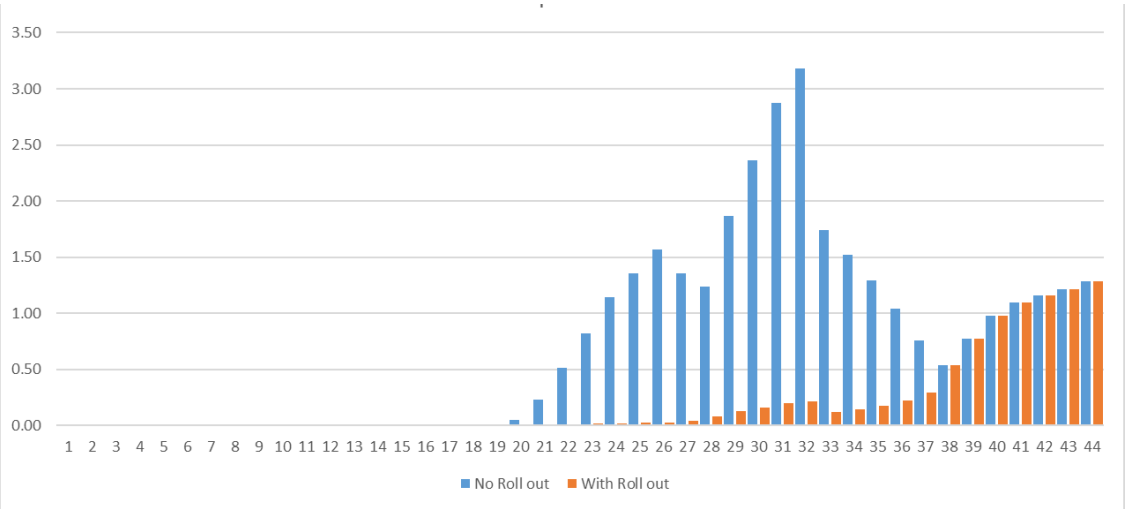

Figure 21. Number of new deaths for each day if one assumes that the chances of dying if Covid-19 symptoms occur is twice as high for those entering the study during Days 1-14 than for Days 15-end.

Discussion

These three examples confirm that (1) selection bias due to informative censoring can have an important effect, (2) the time-wise distribution of those rolling out influence the level of bias, and (3) the hazard ratio does have a secondary effect on the level of bias. A follow-on question is: could the extensive supplementary data provided in the form of Kaplan-Meier life tables by Dagan et al. be utilized to fully address such a bias? The Kaplan-Meier approach yields values associated with a time difference, described as 'number censored,' after being counted in the group for a certain number of days whereas

the examination of the bias carried out in the previous section utilized events which occur at specific times, in an absolute sense. One can simplify the question, then, by considering the combination of an initial distribution showing when matches enter the study and a list of the number censored after a specific number of days, analogous to what is provided by Dagan et al. on a much larger scale. Dagan et al. also provided the total number of those who rolled out before the end of the study period, from which one can calculate how many matches are intact at its end.

To examine this question, consider the thought experiment illustrated in Figure 22. Herein, the concern is not related to what outcomes of any events might be. Rather, the goal is to assess whether more than one time-wise distribution of matched study participants is possible. In other words, with constraints applied by Dagan et al., could one deduce a unique distribution of the number who roll out of the study on any given day? If so, one would be able to rule out the possibility of ambiguity implied by the two different plots in Figure 8 and Figure 16.

The example includes a total of 18 matches formed over four days (4 on Day 1, 6 on Day 2, 6 on Day 3, and 2 on Day 4). After Day 4 (Day 5), we take as given that 9 matches remain intact. Our question can posed as this: If we are given that 1 match lasts (i.e., is censored after) 4 days, 3 matches last 3 days, 7 matches last 2 days, and 7 matches last for 1 day, is this associated with only one unique possible sequence of events?

To determine the answer, one can set up a system of equations where all possible combinations are expressed. It is certain that the lone 4-day event started from a match on Day 1 and lasted until Day 5. Likewise, it is certain that every match formed on Day 4 is censored after only 1 day. For this example, the remaining possible outcomes for matches are expressed in terms of 8 different variables (shown as a-h, on the left) while the constraints applied to the problem may be expressed as equations, of which there are 7 (shown on the right). Because the number of variables exceeds the number of constraints, the problem is under-constrained. Like Dagan et al., the number of matches formed is provided for each day, the number of matches left at the end of the study period is given. The number of matches which last (or are censored after) each allowable number-of-days is also given, and so this specification is analogous to what is provided in a Kaplan-Meier table, which is also provided by Dagan et al. That the system of equations is under-constrained suggests that more than one solution is possible.

Number of matches formed each day during a following 4-day study period (study ends on Day 5):

Day 1	Day 2	Day 3	Day 4	Day 5
4	6	6	2	(9 left - known)

Given: 18 matches, 9 matches left on Day 5
1 match lasts 4 days
3 matches last 3 days
7 matches last 2 days
7 matches last 1 day

Match outcomes expressed as 8 variables a through h
4-day: 1 match lasts from Day 1 to Day 5 (only option, known)
3-day: a matches lasts from Day 1 to Day 4
b matches lasts from Day 2 to Day 5
2-day: c matches last from Day 1 to Day 3
d matches last from Day 2 to Day 4
e matches last from Day 3 to Day 5
1-day: f matches last from Day 1 to Day 2
g matches last from Day 2 to Day 3
h matches last from Day 3 to Day 4
2 matches last from Day 4 to Day 5 (only option, known)

Rules expressed as equations for variables a through h
1) a + b = 3 (3 matches last 3 days)
2) c + d + e = 7 (7 matches last 2 days)
3) f + g + h + 2 = 7 (7 matches last 1 day)
4) 1 + a + c + f = 4 (4 matches formed on Day 1)
5) b + d + g = 6 (6 matches formed on Day 2)
6) e + h = 6 (6 matches formed on Day 3)
7) 1 + b + e + 2 = 9 (9 matches remain on Day 5)

There are 8 variables but only 7 equations. Hence the system is under-constrained.

Figure 22. Thought experiment to demonstrate that, given an identical initial daily distribution of matches, two different sequences in time can lead to the same information in a Kaplan-Meier table.

To fully demonstrate the proof, one may carry out calculations and perform linear algebra as needed. In Figure 23, two different solutions are shown. Each was determined by first assuming a value for the variable 'a' (the number of matches which start on Day 1 and end on Day 4) and then the solution follows. The essential point is that more than one histogram shape of accumulated matches can be produced.

Possible outcome 1: **Begin by setting a = 1**.
Carrying out the linear algebra leads to:
4-day: 1 match starts Day 1 and ends Day 5 (only option, known)
3-day: **1** match starts Day 1 and ends Day 4
2 matches start Day 2 and end Day 5
2-day: 1 match starts Day 1 and ends Day 3
2 matches start Day 2 and end Day 4
4 matches start Day 3 and end Day 5
1-day: 1 match starts Day 1 and ends Day 2
2 matches start Day 2 and end Day 3
2 matches start Day 3 and end Day 4
2 matches start Day 4 and end Day 5 (only option, known)

Possible outcome 2: **Begin by setting a = 2**.
Carrying out the linear algebra leads to:
4-day: 1 match starts Day 1 and ends Day 5 (only option, known)
3-day: **2** matches starts Day 1 and ends Day 4
1 match starts Day 2 and ends Day 5
2-day: 1 match starts Day 1 and ends Day 3
1 match starts Day 2 and ends Day 4
5 matches start Day 3 and end Day 5
1-day: 0 matches start Day 1 and end Day 2
4 matches start Day 2 and end Day 3
1 match starts Day 3 and ends Day 4
2 matches start Day 4 and end Day 5 (only option, known)

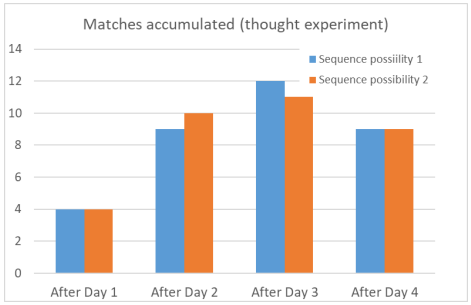

The under-constrained system means that more than 1 result for accumulated matches for intermediate days within the study (as specified).

Figure 23. Demonstration that more than one possible outcome is possible for the thought experiment described.

The two different temporal sequences shown can satisfy the same constraints (number of matches given each day, number of matches existing and the end, and censorship times consistent with what is given by a Kaplan-Meier table for the time period). Thus, more than one possible distribution of matches may exist. Moreover, consider that if yet another day were added to this thought experiment, the number of variables would increase by 5. Referring to Figure 20, one would need to account for Day 1 to Day 6, Day 2 to Day 6, Day 3 to Day 6, Day 4 to Day 6, and Day 5 to Day 6. However, only two added constraints would be required: one equation for the additional time-interval, and a second to account for the number of matches initiated on the new day. So even more time-wise distributions of matches are possible. As the number of days increases further (e.g., to 44), the number of possible distributions of outcomes increases multiplicatively.

This exercise undergirds the statements about the limitations of Kaplan-Meier tables carefully drawn out in the literature. A recent article by Rosen et al. expounds upon the issue central to the present discussion, as follows: "Despite widespread use of the Kaplan –Meier method, little attention has been paid to the proportion of participants censored over time… Few (trials) if any, disclose the number of patients censored at each time interval. Therefore, little is known about the pattern of censoring… In 2018, the Lancet Oncology editors mandated publications presenting Kaplan-Meier plots to include the number of censored patients at each time interval."[11] An example of the basis of these statements, and the decision made by the Lancet Oncology editorial committee, is well illustrated herein. The two different histograms of matches, shown in Figure 8 and Figure 16, leads to the difference in 'expected outcome' when the combined effects of informative censoring and temporal delays in outcomes are taken into account. The combination of an initial list of when matches are formed and a Kaplan-Meier life table is insufficient for assessing a bias which is affected by time in an absolute sense, as is the situation depicted herein.

There is however, some helpful information which can be gleaned from the Kaplan-Meier tables. Consider once more Figure 22 and note the two lines specified as the 'only option' for matches in the left column. If one considers the last line of the various Kaplan-Meier tables corresponding to the longest (44-day) differential in the Dagan et al. supplement, it can be seen that about 500-600 matches entered into the trial on Day 1 (depending on the precise table) persisted to the end of the study period. Estimating the Day 1 matches as about 2,000 would mean that 25-30% of those who enter the study on Day 1 remain through the trial. In the two examples given herein, the values are considerably lower (less than 5% for both the neutral case and the optimized case). With specific reference to the present study, one might conjecture that this discrepancy might be due to the nature of the population given the vaccine on Day 1 of the trial. For example, if those individuals who were relatively young and healthy (e.g., government or business leaders and their families) were given the vaccine on Day 1 in order to encourage participation by others, one might expect that matches for those individuals to be more likely to remain unvaccinated throughout the 44-day period. This approach toward reconciliation is difficult to extend to time delays of 2 days to 43 days, for the reasons depicted in Figure 22 and Figure 23.

This difference suggests that a better way of quantifying the bias using a null vaccine approach would be to include a more complete representation of information for both the Day that a match is formed along with the Day it is censored, for all combinations. For example, how many matches were formed on Day 5 were subsequently censored on Day 6 and, for example, on Day 22? From a full compendium of data,

one would not only be able to answer that question, but would also be able to compute with certainty how many matches formed on Day 5 were still uncensored by, say, Day 23. If this data were made readily available to researchers, it would enable a direct computation of the fraction of individuals who entered the study on, say, Day 5 and who remained on Day 23. As such, this would be a preferred method of determining values in Column 11 of the spreadsheets of Figure 9 and Figure 15, which by necessity herein relied on assumptions and probability laws.

The Kaplan-Meier tables from Dagan et al. also provide an interesting and insightful detail if one considers the 1-day delay time. Analogous to the line in Figure 22, it can be seen than anyone who entered the study on the last day would end up in this group. If one considers the number of symptomatic infections for both vaccinated and unvaccinated groups censored after only 1 day (p. 49, supplemental material), one can see that 227 symptomatic Covid-19 infections were counted for the unvaccinated group whereas only 90 symptomatic infections were counted for the vaccinated group. This excerpt of the Kaplan-Meier table is shown in Figure 24. The red-outlined boxes indicated the 1-day differential.

This large difference diminishes as the time-delay represented in the Kaplan-Meier life table increases. By a 7-day differential, the number of symptomatic people in the vaccinated group exceeds that of the unvaccinated group (195 to 190, as indicated by the blue-outlined boxes). The initial difference (227 vs. 90) is incompatible with the notion that the vaccine has no effect in only one day. One logical explanation is that the low value on Day 1 for the vaccinated gorup could be a consequence of the rule that only those without Covid-19 symptoms may receive the vaccine.

	Main Analysis										
	Unvaccinated					Vaccinated					
Time (Days)	Number at Risk	Number of Events	Discrete Time Hazard per 100,000	Number Censored	Cumulative Incidence	Number at Risk	Number of Events	Discrete Time Hazard per 100,000	Number Censored	Cumulative Incidence	Number at Risk
											Symptomatic SARS-CoV-2 Infection
1	596618	227	38	39287	0.000	596618	90	15	39277	0.000	526877
2	557104	222	40	32773	0.001	557251	116	21	32769	0.000	515683
3	524109	199	38	27499	0.001	524366	190	36	27482	0.001	505725
4	496411	198	40	25910	0.002	496694	154	31	25906	0.001	498817
5	470303	204	43	29775	0.002	470634	193	41	29762	0.001	491683
6	440324	179	41	26377	0.002	440679	165	37	26374	0.002	479330
7	413768	190	46	27743	0.003	414140	195	47	27720	0.002	468974

Figure 24. Excerpt of Kaplan-Meier life table from Dagan et al. Supplement, p. 49: symptomatic Covid-19 for time differences of 1 through 7 days. (permission to reproduce requested)

This raises a different, but arguably even more significant, issue from the delayed censoring discussion. To perform a fully fair comparison, both the vaccinated and unvaccinated groups should enter the trial with the same level of symptom occurrence. Since those who enter the vaccinated group are symptom-free the day of entry, it is essential that only those who are symptom-free should be allowed to enter the unvaccinated group as part of a match. The Day values in the Kaplan-Meier table of Figure 24 strongly suggests that Dagan et al. allowed unvaccinated people with symptoms to be matched with vaccinated individuals who were free of symptoms. It would be interesting to know how many of the

227 symptomatic infections in the unvaccinated group correspond to those who entered the study on the last day (calendar Day 43) of the study period.

If symptomatic people are entered the unvaccinated arm, they cannot roll out of it by attaining a vaccination. Since some with symptoms progress to more serious events, the increased values in the unvaccinated group signified by the one-day differential could be quite impactful and may permeate other aspects of the results. This consideration only reinforces how important it is for researchers to contemplate the ramifications of the CDC rule that vaccinations are permitted only for those who do not present Covid-19 symptoms. This is especially important to factor into analyses in which consider data is collected over a compressed time period.

Although the nature of the rolling cohort exacerbates the bias, this problem may be present to some degree in other studies conducted on Covid-19 vaccines. Start first with a thought experiment which uses a null vaccine as a basis for comparison: suppose a study has two groups which are unvaccinated and are tracked for symptoms from 1 February to 30 April with assumed constant infection rates over time. For Group 1, it is known that each individual is free of Covid-19 symptoms on 1 February. For Group 2, it is known that each individual is free of Covid-19 symptoms on both 1 February and on 1 March. Would it not be reasonable to expect that the percentage of people who experience symptoms by 30 April would be higher for Group 1 than for Group 2? Likewise, if individuals- who must be free of symptoms when vaccinated- are exposed to the vaccine on 1 March and tracked through 30 April, it would be more appropriate to compare the outcome to Group 2, not Group 1. In a similar vein, depending on the study, one must also consider that some individuals may not report symptoms until after the follow-up is concluded. So even if the same time difference is considered, a discussion of how bias is either eliminated or addressed should be included when outcomes for early time periods are compared to outcomes for later time periods.

Now examine a recently published study by Jara et al., which included clear and precise statements of the model assumptions in their associated supplementary material. One statement reads: “The chance to get the vaccine is the same for everyone in the time-predictor-dependent risk set is the same (sp). We think this is also a valid assumption for our analysis because the vaccination scheme was primarily based on individual’s age.”[12] One outcome described in their article is reproduced in Figure 25 and demonstrates an apparent increase in Covid-19 cumulative incidence when one dose of the two-dose vaccine studied (Sinovac) is administered but then a decrease in cumulative incidence of Covid-19 when two doses are administered. Is it possible that the restriction of the second dose vaccine to only those free of Covid-19 symptoms might lead to a selection bias due to informative censoring which favors the administration of two doses of the vaccine?

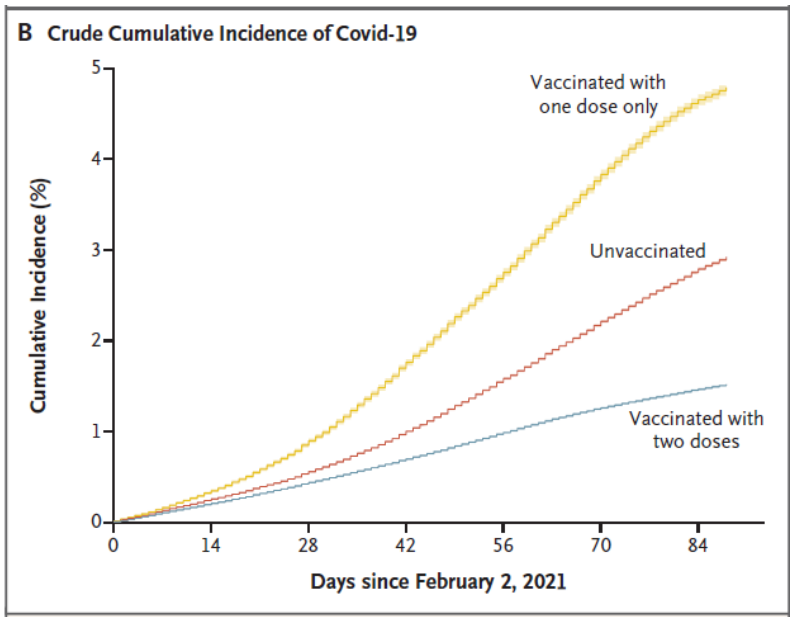

Figure 2. Vaccination Rollout and Crude Cumulative Incidence of Covid-19 in the Study Cohort.

Panel A shows the pace and coverage of the vaccination program among persons who received both doses of vaccine (first and second doses shown separately) or only one dose during the study period (February 2 through May 1, 2021). Panel B shows the crude cumulative incidence of Covid-19 during the study period among unvaccinated persons, among persons who had received only one dose of vaccine, and among persons who had received both doses of vaccine. The relatively high cumulative incidence of Covid-19 in the one-dose group should be interpreted with caution. As shown in Panel A, this group initiated vaccination approximately 40 days after the beginning of the vaccination campaign on February 2, 2021. Therefore, the incidence curve includes all cases that occurred from before vaccination up to 13 days after receipt of the first dose. Shading on the lines indicates 95% confidence intervals.

Figure 25. Reproduction of Figure 2B of Jara et al. (ref. 12) demonstrating that one dose of Sinovac appears to increase the cumulative incidence of Covid-19 while two doses appear to reduce the cumulative incidence of Covid-19.

A second example of an apparent increase in covid-19 cases after the first dose of vaccine (Pfizer-BioNTech BNT162b2) is provided by Bernal et al (ref. 6). The odds ratio is above unity within the time within which a second dose might typically be administered. Therefore, it seems reasonable that a second dose would not be administered to patients who, for whatever reason, are most likely to experience a confirmed case. Is it possible that this result might also be influenced by a selection bias due to informative censoring, which biases toward a perceived increased vaccine efficacy for those hwo receive two doses of vaccine?

Many have attributed relatively disappointing performance of the vaccines in real-world scenarios, compared to research trials, either to the occurrence of new variants or to declining vaccine effectiveness in time. Is it possible that an important contributor to this discrepancy is the “neglected cause of bias” (as phrased in ref. 4) in favor the vaccine for many studies?

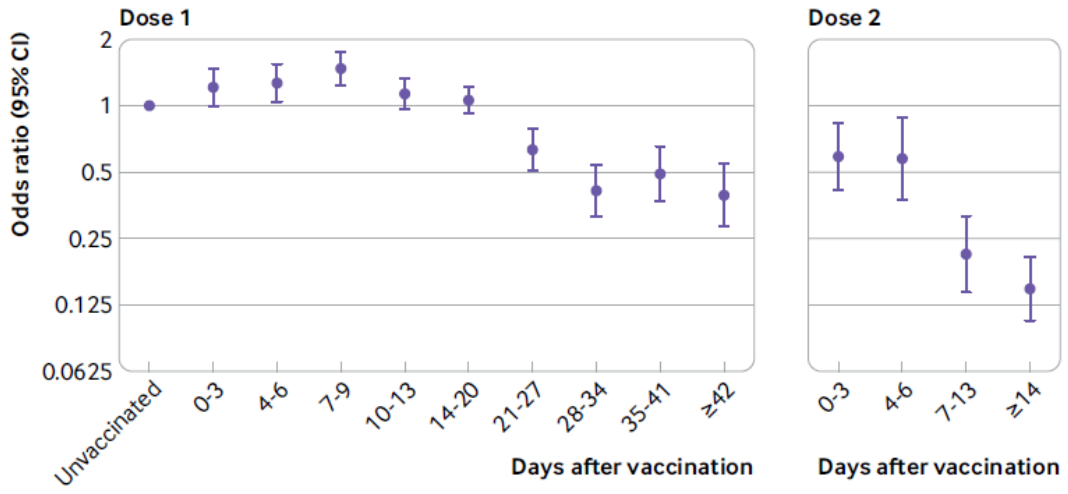

Fig 2 | Adjusted odds ratios for confirmed cases of covid-19 by interval after vaccination with Pfizer-BioNTech BNT162b2 before 4 January 2021 in those aged 80 years and older

Figure 26. Reproduction of odds ratio from ref. 6 showing that one dose of Covid vaccine leads to an apparent increase in cases through the first 20 days after vaccination while two doses lead to an apparent benefit.

The apparent uptick in Covid cases after one dose of vaccine demonstrated in these two examples (Jara et al. and Bernal et al.) raise another important question. Dagan et al. reported 2,389 symptomatic cases within the vaccinated group. Surely, many of those who were experienced symptoms prior to receiving a second dose, yet the article does not detail how those individuals fared compared to either the unvaccinated group or the fully vaccinated group. Is it possible that the authors might have chosen the phrase, "For each person, follow-up ended at the earliest of the following events: *occurrence of an outcome event*…" in order to allow the termination of follow-up once Covid-symptoms prevent an individual who has received *one* dose of vaccine from obtaining a second dose? If not, why would the authors avoid stating the condition as "occurrence of Death due to Covid-19" since follow-up would surely continue for all other outcomes?

*Was the phrase "occurrence of an outcome event" chosen specifically to obscure how those who experience symptoms and other outcomes after a **single dose** of vaccine fared in the study? In other words, did follow-up end prior to more serious outcomes for those individuals who received one dose of vaccine. One might rationalize that this approach would simplify the analysis to focus on "unvaccinated" compared to "fully vaccinated." However, such a choice would lead to serious ramifications with respect to selection bias due to informative censoring.*

Summary and conclusions

A very simple model demonstrates that the "selection bias due to informative censoring" can play an outsized role, influencing how one should interpret data in Covid-19 vaccination studies, especially those employing a rolling cohort. The bias originates from a very simple reason: People experiencing symptoms of Covid-19 are ineligible to be vaccinated.

The approach to demonstrate how one could estimate of bias for the rolling cohort described in the New England Journal of Medicine article by Dagan et al. utilized the concept of a "null vaccine" along with several very simple assumptions. The exercise included: (1) an assumed time-wise distribution of those rolling out of the unvaccinated cohort (not provided by Dagan et al.), (2) the number of newly symptomatic people on each day was assumed to be proportional to the number of people in each cohort with the total unvaccinated people who become symptomatic serving as a de facto calibration setting, and (3) a fixed 18-day delay between symptom appearance and death for the relatively small fixed percentage of people who are symptomatic with Covid-19 with each symptomatic individual having the same odds of dying. The two proportionality constants are selected to achieve values for the unvaccinated persons reported in Figure 2 of Dagan et al. This process was repeated for a second assumed 'optimized' time-wise distribution and the same assumed 18-day delay between symptom appearance and death for the (small) fraction of those who experience Covid-19 symptoms. A third example demonstrated the effect of an age-associated hazard ratio between those who enter the study early in the study period and those who enter later is applied.

The outcome of these three examples demonstrates that the selection bias due to informative censoring should be explored more deeply for Dagan et al. For the 'neutral' distribution, the estimated expected number of deaths under the assumptions given for the "null vaccine" worked out to be 19.3, compared to the 32.0 for the unvaccinated group. In other words, the "statistical rule" that only those who are free of Covid-19 symptoms can obtain a vaccine leads to a 40 percent reduction in the expected number of deaths implied by the "unvaccinated" label.

Moreover, the choice of time-wise distribution of those rolling out of the unvaccinated cohort (due to receiving the vaccine) strongly affects the expected number of deaths from the "null vaccine" in the particular rolling cohort described by Dagan et al. With the assumptions described herein, the value would be reduced from 32 to 12.0, as demonstrated in the second example. In turn, one could approximate vaccine effectiveness for preventing death over a 44-day period after vaccination as [1-(9 actual deaths/12.0 expected deaths)] = 25%. Interestingly, these two values are somewhat close to the values found by using the data in the supplemental material of Dagan et al. depending on whether one "counts" the 7 additional deaths reported in the unvaccinated group there [leading to 39-20)/39 = 49%] or rather in the vaccinated cohort [leading to (32-27)/32 = 16%]. This outcome is consistent with Campigotto and Weller, who found that the bias depends heavily on proportion of individuals who are censored.

Campigotto and Weller also found that hazard ratio influenced the bias in a secondary sense.[5] Such an effect was explored herein using an age-adjusted approach. Approximately twice as many individuals over age 60 were vaccinated in the first 14 days of the trial compared to the number vaccinated later and because age plays an outsized role in whether a symptomatic person with Covid-19 survives. To model this effect, a factor of 2 was applied to compute the number of deaths of symptomatic individuals for those who entered the trial in the first 14 days. This age-adjusted model of the hazard ratio, combined with the 'optimized' time-wise distribution, led to an outcome wherein the entire difference of 32 deaths for the unvaccinated group and 9 for the vaccinated group could be explained only by selection bias due to informative censoring. In other words, the set of circumstances and assumptions laid out in this the example would lead to a determination that the actual vaccine had zero efficacy.

The use of a "null vaccine" assumption to check for bias, using a statistical 'rule' that only those free of Covid-19 symptoms, might also be useful outside the scope of the rolling cohort utilized by Dagan et al. This bias is likely to play an outside role when (1) the overall study time is compressed, (2) vaccines are administered throughout the study period, and/or (3) two doses of vaccines are required – meaning that the individual in the test group must be Covid-symptom free at *two* specific times during the test period. Alternatively, one could take measures to ensure that only an unvaccinated subgroup -free of Covid-19 symptoms at the time at which the vaccination occurs for the test group- is used as the control. In either case, the focus should be placed on ensuring comparisons are conducted on a fair basis.

In a narrow sense, a key conclusion of this effort is that the research community would be served if Dagan et al. provide the time-wise distribution of those rolling out of the unvaccinated control cohort. This exercise demonstrates that there is not simply a possibility, but rather a mathematical certainty, that the data in arguably the most impactful figure of the article (2E, showing number of deaths in each group) is biased in favor of the vaccine. While the authors analyze the data in a fashion which avoids a direct comparison of the two values (32 vs. 9), they also state only that "any such bias was small" without sufficient quantification of what is meant by the word "small." The present exercise focuses on 'death' as an outcome, but any outcome with a delay from symptom origination, such as severe Covid-19 or hospitalization, would also be affected by this same bias.

Clearly, the level of selection bias due to informative censoring depends on data which has not been released. It is counter-intuitive that a widely referenced journal article on the subject of Covid-19 vaccine effectiveness (Dagan et al.) does not disclose the total number of hospitalizations and deaths due to Covid-19 which occurred during the time period of the study. Moreover, the total number of hospitalizations and deaths from all causes is withheld. Others have addressed this overarching ethical concern in a much broader sense.[13]

Clearly, data transparency is essential, both for thorough analyses and academic debate. Unfortunately, many of the most consequential works related to addressing the Covid-19 crisis maintain significant limitations on data availability, and perhaps that is the most important point to be made in the present discussion. Academic journals and government bodies should *require full transparency* of medical studies, especially for work related to Covid-19. If transparency cannot be ensured by journals (or by other entities), then the research community, governmental bodies, and the general public are cautioned to avoid relying upon them when making decisions about health and safety.

Acknowledgments

Working through this exercise of exploring the effects of various assumptions on bias for a specific study was an interesting challenge, and in many ways, enjoyable. One goal of this effort is to lay out a process that a broad audience, both within the medical field and outside of it, could follow.

The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government

[1] Dagan N., Barda, N., Kepten, E., Miron, O., Perchik, S., Katz, M.A., Hernan, M.A., Lipsitch, M., Reis, B., and Balicer, R., "BNT162b2 mRNA Covid-19 Vaccine in a Nationwide Mass Vaccination Setting," *The New England Journal of Medicine*, Vol. 384;15, pp. 1412- 1423, (with supplementary materials pp. 1-54.

[2] https://www.cdc.gov/coronavirus/2019-ncov/vaccines/faq.html , accessed 15 July 2021.

[3] https://www.bustle.com/wellness/can-you-get-covid-19-vaccine-sick-doctors, "Can You Get the COVID Vaccine If You're Sick? Don't Risk Infecting Someone Else" accessed 15 July 2021.

[4] Templeton, A.J., Amir, E., Tannock, I.F., "Informative censoring – a neglected cause of bias in oncology trials," Nature Reviews Clinical Oncology, April 2020, Vol. 17, pp. 327-328.

[5] Campigotto, F. and Weller, E., "Impact of informative censoring on the Kaplan-Meier estimate of progression-free survival in Phase II clinical trials," Journal of Clinical Oncology, 2014 Sep 20, 32(27):3068-3074.

[6] Bernal, J.L., Andrews, N., Gower, C., Robertson, C., Robertson, C., Stowe, J., Tessier, E., Simmons, R., Cottrell, S., Roberts, R., O'Doherty, M., Brown, K., Cameron, C., Stockton, D., McMenamin, J., Ramsay, M., "Effectiveness of the Pfizer-BioNTech and Oxford-AstraZenca vaccines on covid-19 related symptoms, hospital admissions and mortality inolder adults in England: test negative case-control study," *British Medical Journal*, 13 May 2021; 373:n1088.

[7] Zhou, F., Du, R., Fan, G., Liu, Y., Liu, Z., Xiang, J., Wang, Y., Song, B., Gu, X., Guan, L., Wei, Y., Li, h., Wu, X., Xu, J., Tu, S., Zhang, Y., Chen, H. Cao, B., "Clinical couse and risk factors for mortality of adult inpatients with COVID-19 in Wuhan, China: a retrospective cohort study," The Lancet, Vol. 395, pp. 1054-1062, March 28, 2020.

[8] Faes, C., Abrams, S., Van Beckhoven, D., Meyfroidt, G., Vlieghe, E., Hens, N., et al. "Time between Symptom Onset, Hospitalisation and Recovery or Death: Statistical Analysis of Belgian COVID-19 Patients," *Int. J. Environ. Res. Public Health,* 2020, 17, 7560; doi:10.3390/ijerph17207560.

[9] Jaffe-Hoffman, Maayan, "Coronavirus: Who got Israel's First Vaccinations," 7 January 2021, *The Jerusalem Post*.

[10] CDC website, https://www.cdc.gov/coronavirus/2019-ncov/covid-data/investigations-discovery/hospitalization-death-by-age.html, referenced 13 August 2021.

[11] Rosen, K., Prasad, V., and Chen, E.Y. "Censored patients in Kaplan-Meier plots of cancer drugs: an empirical analysis of data sharing," *European Journal of Cancer*, Vol. 141, December 2020, pp. 152-161.

[12] Jara, A., Undurraga, E.A., Gonzalez, C., Paredes, F., Fontecille, T., Jara, G., Pizarro, A., Acevedo, J., Leo, K., Leon, F., Sans, C., Leighton, P., Suarez, P., Garcia-Escorza, H., Araos, R., "Effectiveness of an Inactivated SARS-Cov-2 Vaccine in Chile," *The New England Journal of Medicine*, 7 July 2021, DOI:10.1056/NEJMoa2107715.

[13] "Peter Doshi: Pfizer and Moderna's '95% effective' vaccines – we need more details and the raw data," *The BMJ Opinion*, 4 Jan 2021 (with follow up on 5 Feb 2021), https://blogs.bmj.com/bmj/2021/01/04/peter-doshi-pfizer-and-modernas-95-effective-vaccines-we-need-more-details-and-the-raw-data/.