Relative Change in Paired Data

1. **Introduction** **2**
2. **Mean of Ratios or Ratio of Means: which one to choose** **3**
3. **Statistical Tools for the Mean of Ratios** **6**
4. **Statistical Tools for the Ratio of Means** **7**
5. **Mean of Ratios *versus* Ratio of Means: the math** **8**

Special sections:

	**Core Concepts:** This chapter is focused on methods for relative inference for paired data. Sections that are conceptually at the core of that discussion will be highlighted.
	Some passages herein are for the **statistical aficionado**, and can be skipped by others.
	**Nerd alert…** There are some passage herein that are particularly nerdy, and can usually be skipped. The photo to the left will be your warning.

Section One: Introduction

We will use here an example that was used in the paired data chapter, before and after story of woody vegetation cover (Table 1).

Table 1. Pre- and Post-treatment woody vegetation cover, with incremental changes[1] and relative changes[2] noted for each plot.

Plot	PRE	POST	Post-Pre	Post/Pre
Plot 01	1	7	6	7.00
Plot 02	30	42	12	1.40
Plot 03	17	34	17	2.00
Plot 04	58	84	26	1.45
Plot 05	9	27	18	3.00
Mean	**23**	**38.8**	**15.8**	**2.97**

Note that the mean of the differences (column four) is the same as the difference of the means: $\left(\bar{y}_{Post} - \bar{y}_{\text{Pre}}\right) = 38.8 - 23 = 15.8.$ This equality falls apart with relative change, since in that case $\left(\bar{y}_{Post} / \bar{y}_{\text{Pre}}\right) = 38.8 / 23 = 1.69,$ which is quite different from the mean of the ratios (2.97). That difference in outcomes will play a central role in these notes.

In much ecology work, interest is often in relative, not incremental change. "Did the controlled burn reduce fuel loads by at least 50%? "After a burn, has the abundance of some favored plant increased by at least 50%?" So here we consider how to approach making confidence intervals for relative change.

For clarity, we reiterate that we are in this conversation assuming paired data. We pointed out in the Introduction that for incremental change, the average of the differences is the same thing as the difference of the averages. But, switching to relative change, we noted that the mean of the ratios and the ratio of the means are *not* the same.

---

[1] We usually take care to order the subtraction to yield a positive number for the mean; that only to make the story telling easier.

[2] All the plots showed an increase, so we chose Post/Pre for story-telling purposes also.

Section Two: Mean of Ratios or Ratio of Means: which one to choose

Core Concepts: Understanding that they are meaningfully different, and then knowing how to decide between them is core…

We will here use the same artificial data (presented in two different scenarios) to illustrate how you might go about choosing between the Mean of Ratios (MoR) and the Ratio of Means (RoM).

Suppose a landscape of interest (Microsizia National Park) is only 50 hectares in size; for controlled burn purposes, NPS fire ecologists have divided the park into five 10 ha plots, each subdivided into 10 1ha subplots. The intent of the controlled burn is to change the habitat to favor a particular species of plant (endangered, say, or a native being diminished in abundance by exotics). In the depiction below (Figure D.1), a yellow subplot indicates presence of the target species; white indicates total absence. For example, plot 1 was 20% covered pre-burn and 50% covered post-burn, more than double. Keep in mind that the data represent the entire population, so we are here discussing parameters perfectly known, not just sample estimates thereof.

Figure D.1 Presence and absence of a targeted plant species, on each sublplot of Microsizia National Park.

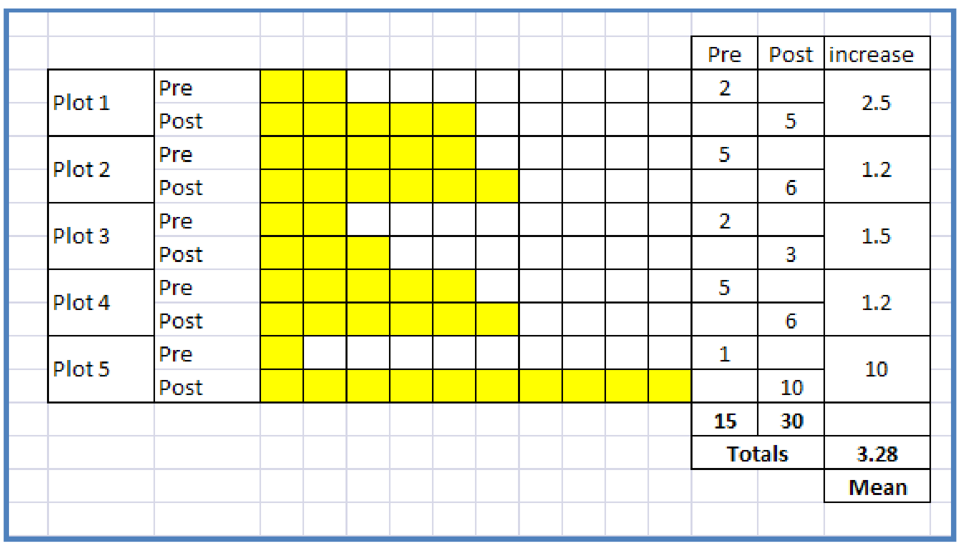

On the landscape scale, a total of 15 of the subplots were covered prior to treatment, and 30 were covered post-burn, a doubling in cover (by the ratio of the totals). The average coverage prior to treatment is 15/5 = 3, and after 30/5 = 6, so one can speak in terms of the ratio of totals or the ratio of means (RoM): they are the same.

Yet if you measure the increase separately in each plot, and then average the results (the mean of the ratios (MoR)), one gets the impression that the coverage has more than tripled. These results differ dramatically. In this tinkertoy example, it is clear that the abundance has doubled across the landscape, leading to the ratio of the means being the best choice (of parameter to estimate, and hence statistic to choose)

Second Illustration (counting things)

In another study in Microsizia National Park, researchers were interested in the increase in weight of the Microsiza black bear (a very unusual subspecies) from spring to summer, since their main food source are the plants from the above study. Their ability to put on weight over a season can be astounding, almost unbelievable. By great good fortune for our purposes, there happens to live one such bear in each of the five 10 ha plots in the park. The table below shows their spring and fall weights (in pounds).

Table D.1 Spring and fall weight (pounds) of bears in Microsizia National Park

	Spring Weight	**Fall Weight**	**Ratio (F/S)**	**Difference (F-S)**
Bear 1	200	500	2.5	300
Bear 2	500	600	1.2	100
Bear 3	200	300	1.5	100
Bear 4	500	600	1.2	100
Bear 5	100	1000	10	900
	300	**600**	**MoR = 3.25**	**300**
	Averages			

Here, the individual bear is clearly the important unit, and the average response *per bear* is relevant. If relative change is of interest, then the mean of the ratios (MoR = 3.28, more than triple) is the pertinent statistic to use.

In summary, the key to choosing between RoM and MoR is the following. If change in the population as a whole matters (e.g. coverage of plants in Microsizia National Park), while the individual sampling units are simply a convenient way to generate estimates, then the Ratio of Means is likely the better candidate. On the other hand, if the individual sampling units matter (e.g. bears), and it is not of interest to think in terms of change for the whole population[3], then the Mean of Ratios is the way to go.

[3] Thinking about the total weight gain in the whole population of bears makes my brain hurt.

The following schematic summarizes the process of choosing between incremental and relative change, and, given the latter, between MoR and RoM.:

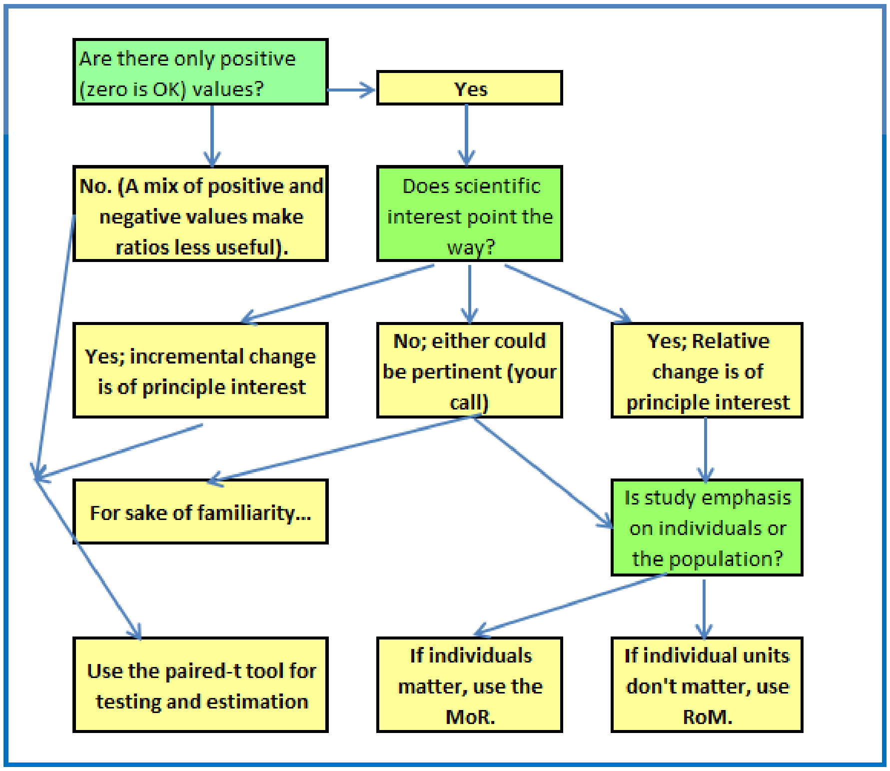

We now discuss statistical tools appropriate for each of these choices, and begin with the mean of individual ratios (MoR).

Section Three: Statistical Tools for the Mean of Ratios.

Let us consider first averaging the ratios (MoR). This is unlikely to be used as much at the RoM, but it turns out that we have already studied the statistical aspects of its use in previous sections, and so this discussion will be short, and can be summarized in just two points:

Having calculated a ratio for the values within each pair, the data analysis problem has now been reduced to that of a mean from a single sample (of ratios). This is precisely analogous to the process in the classical paired *t*[4]. So this content has been covered in the chapter **Methods for Paired Data**

In short: he one-sample *t*-distribution can be used if arguments hold that the mean from the sample will be approximately Normal. Failing that, bootstrapping can be used. The interactive Excel tool **options for paired data** that accompanies this book can be used to do the calculations.

As it happens, inference for the ratio of means (RoM) is also straightforward. There are no central limit type arguments for the ratio of two means[5], and so bootstrapping is the go-to method for making a confidence interval. Re-read Section Four if you need a reminder of how bootstrapping works. The interactive Excel tool **CIs for paired data** that accompanies this book can be used to do the calculations.

Considerations when forming ratios. It is useful to think through story-telling consequences when choosing how to construct ratios. Here are our thoughts on some of those points. These are not statistical dicta, only aesthetic suggestions. Here, for convenience, we take the pairing structure as before (pre) and after (post) some intervention.

If the data show an increase, the forming ratios as post/pre will serve better than the other way around. If the size of the increase is less than double (e.g. if the average is 1.75), we recommend using 100% x (post/pre – 1); this would yield an answer as a percentage increase (75% in this example). If the size is more than that, simply use suitable language of multiplication. For instance, if the mean is 3.5, simply refer to a 3.5-fold increase or 3.5 times as much (or language like that).

If the data show a decrease, one could directly use post/pre or 100% x (post/pre), but the reporting language can get snarly. For example if post/pre = 0.20, one can correctly say that there is now 20% as much as formerly. Here we find 100% x (1-pre/post) easier to report. For this example, we would report an 80% decline.

[4] In the classical paired *t*, one studies a single sample of incremental differences rather than relative.

[5] Well, there are some (e.g. if the two means each come from a Krakow distribution, their ratio asymptotically follows the snickerdoodle distribution). There. Did that help? In other words, there exist a few oddball special cases, but their applicability to data like that gathered in fire ecology work is likely to be quite limited.

Section Four: Statistical Tools for the Ratio of Means (RoM)

There is no Central Limit Theorem for the ratio of two means, meaning that there is no reliable way to predict what the distribution will be. In addition, although there does exist as standard error formula, it is only approximate, and is usually delivered with warnings explaining circumstances under which it does not work well.

The solution is conceptually simple: bootstrapping. Here we repeat, for your convenience, some of the material from our chapter **Methods for Paired Data** as an introduction/reminder of the principles of bootstrapping.

Bootstrapping for a mean from a single sample proceeds as follows. It is based on the idea that your data themselves are the best representation you have of the population whence they came.

Conceptual algorithm.

(1) Write down each of your *n* sample values an infinite number of times[6]. This will generate a population (of infinite size) from which to do the simulation.
(2) Draw a random sample of size *n* from that population. Write down the mean.
(3) Repeat (2) a large number of times (B = 1000 will suffice[7]).
(4) Sort the values from smallest to largest.

Then, for B = 1000, number 26 (isolating 2.5% that are smaller) and number 975 (isolating 2.5% that are larger) in that sorted list are the lower and upper bounds on a 95% confidence interval[8]. For an 80% interval, use numbers 101 and 900 (isolating the outer 20% and capturing the central 80%) from that list. Stunningly simple, except for step (1).

What step (1) implies, if you think about it a bit, is that any one datum in your original sample could appear one time, twice, thrice (less likely) or even zero times in any given bootstrap sample. How can we trick the computer into doing that without having to face infinity? Read on.

Practical Algorithm

Replace steps (1) and (2) above with: draw a random sample of size *n* from your original sample, *sampling with replacement*. The next two steps are the same.

What does "sampling with replacement" mean, and why does it work? First, the "what." Column 4 of Table 1 has five values, namely the five differences (post-fire minus pre-fire) in vegetation

[6] You might immediately call "bullshit" on this statement, and you would be correct. Hold your breath, and wait. The practical form of the algorithm provides a solution to this seemingly impossible step.
[7] The number of times you repeat this is the size B of the bootstrap simulation. It is important to know that should you choose B = 1000 or B = 100,000, you are still studying inference for a sample of size n.
[8] This is the so-called empirical bootstrap. There are more sophisticated approaches (a popular one being the so-called bias-corrected, accelerated (BCA) bootstrap method that employ some clever machinations on that simulated distribution. We won't be discussing those here; they are not necessary for the core ideas.

cover values (the single sample being studied). The data are: 6, 12, 17, 26, and 18. Independently randomly select numbers from that list (independent selection allows a given number to be repeated in your list). For instance you could[9] get 12, 6, 18, 6, and 17, yielding an average of 11.8. The datum “6” appeared twice in the bootstrap sample.

Why does this work? By sampling with replacement, we perfectly mimic what step (2) of the conceptual algorithm does if you could actually write down all the values an infinity of times (step (1)). A simple computational trick replaces infinity! Gratitude to Brad Efron, the inventor of the bootstrap, for this ground-changing idea.

The Excel app **CIs for paired data** That accompanies these notes will perform the bootstrapping calculations for a ratio of means.

This next section is somewhat nerdy, but pretty important: it is worth wrestling with if you are working with relative change for paired data.

Section Five: Mean of Ratios *versus* Ratio of Means: the math

Let’s suppose (without loss of generality) that we want to form ratios (for either statistic) that are oriented as post/pre (recall that our simplifying pretext here is that we are studying data from before (pre) and after (pose) some intervention.

Let $R_i = \frac{Y_{i,post}}{Y_{i,pre}}$ be the ratio measured on each plot (where $Y_{i,post}$ is the post-treatment measurement, and $Y_{i,pre}$ the pre-treatment measurement on plot *i*).

Then the mean of the ratios is $\bar{R} = \frac{1}{n}\sum_{i=1}^{n} R_i$.

In general, $\frac{\bar{Y}_{post}}{\bar{Y}_{pre}} \neq \bar{R}$; in fact, one can show that for strictly positive values such as these, $\frac{\bar{Y}_{post}}{\bar{Y}_{pre}} \leq \bar{R}$. We will return to this inequality later to discuss circumstances when they are in fact equal.

[9] Ken *did* get precisely this using a random number generator.

Important Caveat: Data must be strictly non-negative[10]. It does not make sense to consider ratios of numbers where one is positive, and the other negative[11]. If you are going to average the ratios (MoR), be aware that a zero in the column that will live in the denominator will blow up your arithmetic. Fortunately, more often than not, RoM will be of more interest, and the stray zero here or there will not hurt.

 

Wow! A double alert: Skip if you wish…

First, let us examine the formula for the mean of ratios in a certain way $\bar{R}=\frac{1}{n}\sum_{i=1}^{n}R_i=\frac{1}{n}\left(R_1+R_2+...+R_n\right)=\frac{1}{n}R_1+\frac{1}{n}R_2+...+\frac{1}{n}R_n$. This is nothing particularly exciting, showing only that we can portray the average as the sum of weighted data values, here with each weight chosen to be equal. This generalizes to a so-called "weighted average[12]", wherein weights can be chosen arbitrarily[13], aside from the condition that the weights sum to 1.

The MoR for the abundance of plants in Microsizia National Park is

$$\bar{R}=\frac{1}{5}(2.5)+\frac{1}{5}(1.2)+\frac{1}{5}(1.5)+\frac{1}{5}(1.2)+\frac{1}{5}(10).$$

Now let us suppose we decide to weight each value by the initial value rather than weighting them equally. The uncorrected weights are 2, 5, 2, 5, and 10, respectively. They sum to 24, so the weighted average is

$$\bar{R}_w=\frac{2}{24}(2.5)+\frac{5}{24}(1.2)+\frac{2}{24}(1.5)+\frac{5}{24}(1.2)+\frac{10}{24}(10).$$

[10] We say, "non-negative" instead of "positive" in order to acknowledge that zero can be an acceptable datum.
[11] In principle, one could have strictly negative values (or "non-positive" to be nerdily correct), but surely you can find a way to rescale the data into positive values. Contemplating ratios of negative numbers makes our heads hurt.
[12] In principle, the classical mean is also a weighted average (with the choice being to have equal weights); that said, the term "weighted average" is used by convention to refer to cases where the weights are chosen to be unequal.
[13] Arbitrary does not imply capricious; weights are usually chosen for some reasonably sound purpose.

No intuition there. But let me re-frame it in symbols.

$$\bar{R}_w = \frac{x_1}{T_x}\left(\frac{y_1}{x_1}\right) + \frac{x_2}{T_x}\left(\frac{y_2}{x_2}\right) + \frac{x_3}{T_x}\left(\frac{y_3}{x_3}\right) + \frac{x_4}{T_x}\left(\frac{y_4}{x_4}\right) + \frac{x_5}{T_x}\left(\frac{y_5}{x_5}\right),$$

where *y* denotes numerator terms, *x* denominator, and T_x denotes the total of *x*. There. Is that better? Not so much? Let's keep at it.

If you look at each term, you can see where we can cancel the "*x*" bits. Let me do that, and then clean it up some.

$$\bar{R}_w = \frac{1}{T_x}(y_1) + \frac{1}{T_x}(y_2) + \frac{1}{T_x}(y_3) + \frac{1}{T_x}(y_4) + \frac{1}{T_x}(y_5) = \frac{1}{T_x}\sum y_i = \frac{T_y}{T_x}.$$

One more step. If we divide both the numerator and denominator by the sample size, clearly the value of that ratio does not change, but we can change its representation.

$$\frac{T_y/n}{T_x/n} = \frac{\bar{y}}{\bar{x}}.$$

Hence the RoM is mathematically just a weighted MoR, with weights chosen to be based on the denominator term. The difference between the two choices can be thought of in terms of how to weight each datum. Weight them equally (which makes sense in the case of the bears and seals), and you get the Mean of the Ratios. Weight them by the values in the control group (chosen by you; they appear in the denominator), and you get the Ratio of the Means.

We indicated earlier that $\frac{\bar{Y}_{post}}{\bar{Y}_{pre}} \leq \bar{R},$ implying that the two can sometimes be equal. Sparing you the proofs, we will assert here two different conditions when this will be so.

(1) If all the relative changes are identical (e.d. each plot or subject increases by the same amount), then $\frac{\bar{Y}_{post}}{\bar{Y}_{pre}} = \bar{R}$; further,

(2) If all the denominator terms (the pre-treatment data in these notes) are the same, then $\frac{\bar{Y}_{post}}{\bar{Y}_{pre}} = \bar{R}.$