



Regression to the mean

Regression to the mean is a technical way of saying that things tend to even out over time. The sprinter who breaks the personal best record will probably run closer to their more typical time on the next race. The medical treatment that achieves stunning results on its first trial will probably not be as effective on its second trial. Regression to the mean refers to this tendency of a variable that is highly distinct from its norm to return to "its normal" over repeated tests. On average, observations tend to cluster around the mean,^[1] whether or not they come after a really unusual value. It only becomes obvious when a strange result (e.g., a hole-in-one in golf) is followed by something much more ordinary (like a double-bogey).

In medicine

Unfortunately, much of the effects claimed by [alternative medicine](#) can often be explained simply as regression to the mean. When Aunt Jane's acne gets better after rubbing mint leaves on her face, that's "[anecdotal evidence](#)" based almost entirely on regression to the mean. Many symptoms will come and go in an apparently random fashion. For example, if recorded in an objective way headaches tend to disappear without the aid of any treatment over time. People seek treatment when their symptoms are particularly severe, when they are at their respective "peak" of symptoms. Regression to the mean, therefore, suggests that if symptoms are excessively severe this week, then next week they should be less severe simply by random fluctuations. If treatment is only sought when these symptoms are at their worst there will almost always be a [coincidental](#) recovery, which can appear even if the treatment has no effect whatsoever.

A placebo control group in controlled trials removes the effect of regression to the mean. Both groups, on average, experience a tendency to regress to the mean (i.e., get better). If, compared to a placebo control group, the treatment group shows a statistically significant decrease of symptoms, even more of a decrease than the placebo group, the decrease can be attributed to the treatment, rather than simply regression to the mean.

Testing

For example, if a researcher gave a large group of students a test and selected the students who performed among the top 5%, these students would be likely to score worse, on average, if re-tested. Similarly, the bottom 5% would be likely to score better on a retest.^[2] In either case, the extremes of the distribution are likely to "regress to the mean" due to simple luck and natural random variation in the results.

Sports

One way of thinking about "regression to the mean" is in terms of sports performance. To win a football championship, for example, it is not enough only to be a good team — one needs to be both good *and* lucky. The team at the top of the standings in mid-season is likely to have been both good and lucky to that point, but cannot count on still being lucky for the rest of the season. For this reason, the team that is at the top of the standings at mid-season is more likely to drop in standings than to remain at the top.

This observation has been tagged the "Sports Illustrated Jinx". The jinx states that a player or team featured on the cover of a sports magazine such as *Sports Illustrated* is likely to have a disappointing year the following season (or even a disappointing game the following week). But if you think about it, a player is only likely to make the cover once, and to only make the cover for some surprisingly good performance -- something truly spectacular that

that requires not only their superlative skill, but also lots of luck to beat the superlative skill of their competitors.

Athletes on the cover of *Sports Illustrated* are likely to be at the very top of their game, and when you're at the top of your game the most likely direction to move next is *down*. The next year, although the player may still be as skilled, they will not be as lucky, and therefore they are more likely to have performances closer to their normal.