11/30 Prof Kevin Lang's title and abstract:

Title: "How Credible is the Credibility Revolution?"

Abstract:
Suppose we observe that a null hypothesis is rejected with a t-statistic of exactly 1.96. How confident should we be that the null hypothesis is false? A moment's reflect tells us that our answer depends on the ex ante probability that the null is true (our prior) and the power of the test. While researchers conducting RCTs frequently prepare power calculations for an alternative hypothesis with mass at a single point, allowing for a distribution of alternative values generally lowers the estimated power. Since we generally do not test hypotheses that we are confidence are false and often use techniques with much lower power than we claim, even in the absence of p-hacking our confidence in results that narrowly reject the null should be low. I estimate a structural model in which the t-statistics are drawn with some probability from true nulls that, therefore, have a t-distribution and a set of false nulls that, given the true coefficient value and test power have an exponential distribution. The model estimates imply that most narrowly rejected nulls are true and should not be rejected.