Abstract:

Not all empirical findings are published. What is the socially optimal rule for deciding whether a study should be published in scarce journal space? We show that if the goal of publication is to inform policymakers about a policy-relevant parameter, then one should publish extreme or ``surprising'' results. For specific policy objectives, the optimal rule may take the form of a one- or two-sided test comparing the point estimate to the prior mean. Dynamic considerations may additionally justify the publication of precise null results. However, if one insists on a rule for which the absence of publication is not informative, or for which frequentist inference remains valid conditional on publication, then the publication rule must not select on the study's findings (but may still select on the study's design). We also characterize publication rules that maximize alternative objectives such as learning or the plausibility of published findings