

Statistics in a social context

Opening remarks

Maximilian Kasy

May 10, 2019

Introduction

- Current debates across the social and life sciences:
 - Publication bias and p-hacking,
 - replicability and replications,
 - pre-analysis plans and other reform proposals, ...
- Motivation of this conference:
 - These debates raise a number of foundational questions,
 - which, I believe, are not well addressed using textbook frameworks,
 - and which require input from several disciplines.

Roadmap for these opening remarks

1. Where I am coming from:
 - a Research with Isaiah Andrews on “which findings get published.”
 - b Research with Alex Frankel on “which findings should be published.”
2. Three alternative perspectives on statistics:
 - a Decision problems,
 - b (optimal) communication,
 - c research as a social process.
3. Brief preview of conference.

Which findings get published?

Andrews, I. and Kasy, M. (2018). [Identification of and correction for publication bias](#)

1. **Published research is selected**

in various ways (significance at different levels, sign, ...)


- Lab experiments in economics and psychology: Statistical significance
- Effect of minimum wages on employment: Statistical significance, sign.
- Deworming: Inconclusive.

2. How do we know?

Form and magnitude of **selection** are **nonparametrically identified**.




- Using systematic replication studies.
- Using meta-studies.

Evidence on selective publication



- Data from systematic replication study of Camerer et al. (2016).
- Z-statistic W , replication W^r , estimate X , standard error Σ .
- Absent selection:
 1. z-statistics should be continuously distributed.
 2. Original and replication estimates should be symmetrically distributed.
 3. Estimates from studies with larger standard errors should be more dispersed, but not shifted.

Selection implies publication bias



- Suppose only findings with z-stats > 1.96 are published.
- The figures plot, as a function of the true mean θ ,
 1. The bias of Z as estimator of θ ,
 2. the coverage of $Z \pm 1.96$ as confidence interval for θ ,
 3. the naive and the correct Bayesian posterior density, for a normal prior, when no finding is published.

Which findings should be published?


Frankel, A. and Kasy, M. (2018). *Which findings should be published?*

- Publication bias motivates calls for reform:
Publication should not select on findings.
- But: Is eliminating bias the right objective?
How does it relate to informing decision makers?
- We characterize **optimal publication rules from an instrumental perspective**:
 - Study might inform the public about some state of the world.
 - Then the public chooses a policy action.
 - Take as given that not all findings get published (prominently).

Key findings


1. **Optimal** rules selectively **publish surprising findings**.
In leading examples: Similar to two-sided or one sided tests.
2. But: Selective publication **always distorts inference**.
There is a trade-off policy relevance vs. statistical credibility.

Example of optimal publication region



- Optimal publication region (shaded). Axes:
 - left Estimate X , standard error S .
 - right "t-statistic" $t = (X - \mu_0)/S$, standard error S .
- Note:
 - Given S , publish outside symmetric interval around μ_0 .
 - Critical value for t-statistic is non-monotonic in S .

A standard foundation of statistics: Decision theory



Questions to ask in this framework:

- Objective function?
- Set of feasible actions?
- Prior information?

Is this an appropriate description of empirical research?

- Some questions:
 - Why do we not just print all the data?
 - Why do we need researchers?
 - What is the purpose of pre-committing to a research design?
 - Does commitment make sense without conflicts of interest?
 - How do we cumulatively learn from published research?
- Can we make sense of publication bias, pre-analysis plans, etc., using textbook foundations of statistics?
- Or do we need alternative foundations, taking into account the social dimension of research?

Alternative foundations

Different ways of thinking about statistics / econometrics:

2. Statistics as (optimal) communication.

- Not just “you and the data.”
- What do we communicate to whom?
- Subject to what costs and benefits?
Why not publish everything? Attention?

3. Statistics / research as a social process.

- Researchers, editors and referees, policymakers.
- Incentives, information, strategic behavior.
- Social learning, paradigm changes.

Proposed agenda

- Derive **optimal methodological recommendations**,
- assuming the **goal** is to promote some notion of collective **learning through communication** of summaries of empirical findings,
- taking into account the **constraints** of human psychology and the social organization of research.

To better understand these constraints, draw on

1. psychology,
2. sociology and history of science,
3. microeconomic theory and information economics.

Conference outline

- **Applied perspectives**

- Katherine Casey (development economics)
Comments on pre-specification and analysis plans
- Simine Vazire (psychology) *The Credibility Revolution in Psychological Science*
- Ben Olken (development economics) *Promises and Perils of Pre-Analysis Plans*
- Daniel Mellow (meta studies)

- **Microeconomic models**

- Jann Spiess (econometrics)
Optimal Estimation when Researcher and Social Preferences are Misaligned
- Alex Frankel (economic theory) *Which findings should be published*
- Isaiah Andrews (econometrics) *Statistical Reports for Remote Agents*
- Marco Ottaviani (economic theory) *Strategic Sample Selection*

- **Philosophical and historical perspectives**

- Theodore Porter (history of statistics) *Statistics, a Tool of Science?*
- Deborah Mayo (philosophy of statistics)
3D Statistics: 7 Responses to Challenges for Statistical Testers
- Zoe Hitzig (economic theory, philosophy)
The Problem of New Evidence: P-hacking and Pre-analysis Plans