RESEARCH ON RESEARCH INSTITUTE

Using metascience to ride the AI shockwave

Tom Stafford
t.stafford@researchonresearch.org
tomstafford.github.io

These slides:
bit.ly/tomstafford

RESEARCH ON RESEARCH INSTITUTE

AI in four moves
AI & research funding
- the present
- the future

Artificial Intelligence in four moves

1. **Calculating machines** 1956
2. **Learning algorithms** 1986
3. **Large Language models** 2022
4. **Agents** 2025
    a. Context
    b. Agency
    c. Ownership

The agentic turn in AI



For most of their short history, Large Language Models (LLMs) have been used interactively; a user poses a question and the model responds as a ‘chat bot’.

Agentic AI changes this relationship. An agent is an LLM equipped with tools such as the ability to search the web, read documents, write and execute code, or call external services.

Grant writing and reviewing sit naturally in the space of agents. The task is goal-oriented, information-rich, rule-bound and iterative. These are exactly the conditions where agents perform well.





ANTHROP\C

Agentic coding and persistent returns to expertise

Published
June 16, 2026

Authors
Zoe Hitzig, Maxim Massenkoff, Eva Lyubich, Ryan Heller, and Peter McCrory

https://www.anthropic.com/research/claude-code-expertise

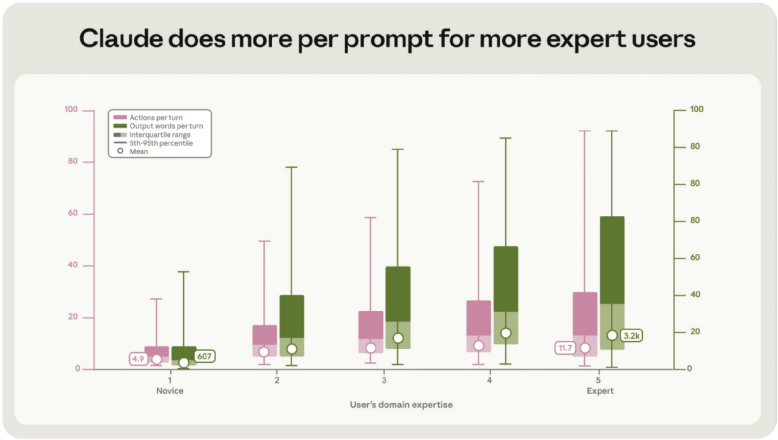


Figure 3: Claude does more per prompt for more expert users
Claude produces more actions (left bar) and text output per prompt (right bar) for more expert users. Boxes span the interquartile range (split at the median). Whiskers represent the 5th to 95th percentile. White dots are geometric means. Both upward trends are statistically significant ($p < 0.001$), as is each adjacent-level step, and they remain significant (at +9% actions and +13% output per expertise level) in a regression controlling for work mode, task value, month, occupation, and model family, with standard errors clustered by user.

September 6, 2026 Research Publication Safety

Research acceleration: The view inside OpenAI

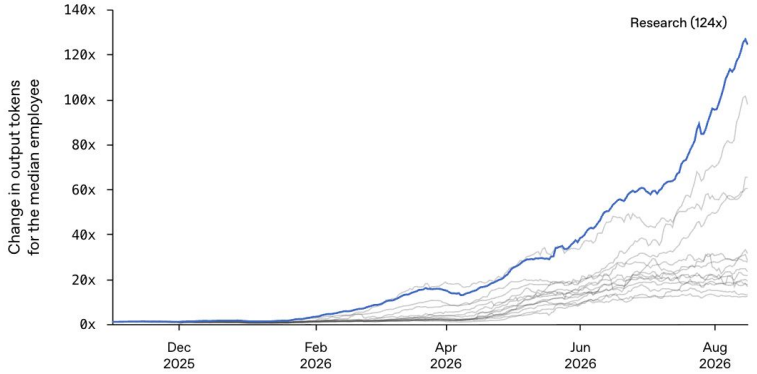
Listen to article 13:39 Share

For AGI to benefit all of humanity, we believe it must be democratically governed. This can only happen through an informed public debate about the capabilities, risks and safeguards of highly capable AI systems. People everywhere need to understand the likely future trajectory of frontier AI, so they can have a meaningful voice in how it develops.

Transparency about specific risks, incidents and safeguards is necessary, but not sufficient. We believe the public also needs to understand how the most capable systems are developing, and how they are driving research progress, inside of frontier labs.

We aim to safely build an automated AI researcher that can work under human supervision to further progress on deep learning and alignment, enabling iterative improvements. According to our measurements, we have now reached the goal, announced last fall, of having an automated research intern by September of this year. By

Usage growth is faster among researchers than in other parts of the company



View methods

Before June 2026, total agent runtime across the research organization was still below that of total human labor. That has since changed. In terms of a standard 8 hour workday, as of mid-August, in total, the research organization uses 3.1 agent-workdays of effort for every workday of human labor.

https://openai.com/index/research-acceleration-view-inside-openai/

Risks

Automation & alignment

Bias

- **Prototypicality**
- **Incumbency**
- **Offloading**
- **Hallucination**
- **Neglect**

Capital

RESEARCH ON RESEARCH INSTITUTE

AI & research funding - now

AI & funding

RoRI RESEARCH ON RESEARCH INSTITUTE

RoRI Summary note

Demand management: findings from a multi-funder workshop

Alyson Fox (Wellcome), Tom Stafford, Peter Kolarz (Research on Research Institute)

April 2026

Produced in partnership with:

wellcome

European Commission

Marie Skłodowska-Curie Actions

Developing talents, advancing research

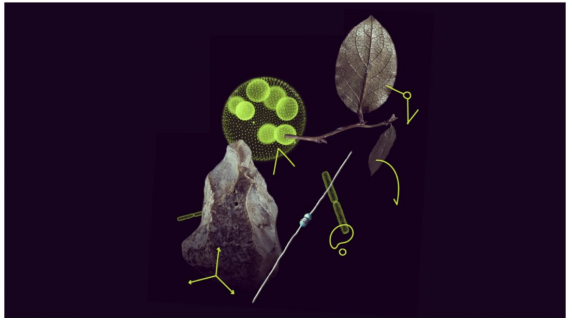
Home | About MSCA | Actions | Funding | Jobs | Resources | What's new | 30 years of MSCA

You are here: Home / What's new / News

Published: 11 Sep 2026

MSCA Postdoctoral Fellowships 2026 receives record number of applications

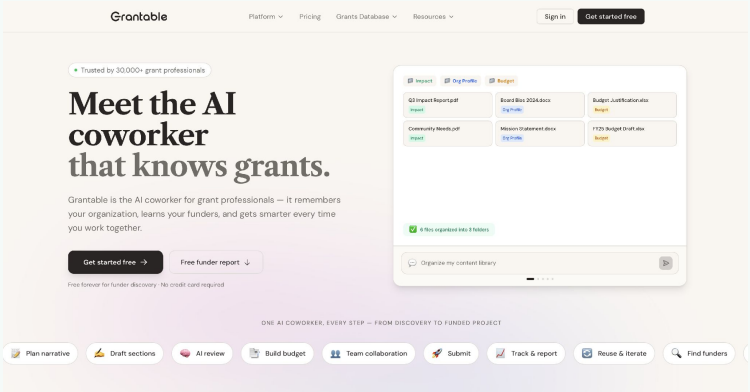
The 2026 call has attracted considerable interest from the research community with 21,627 applications submitted.



Results of the call

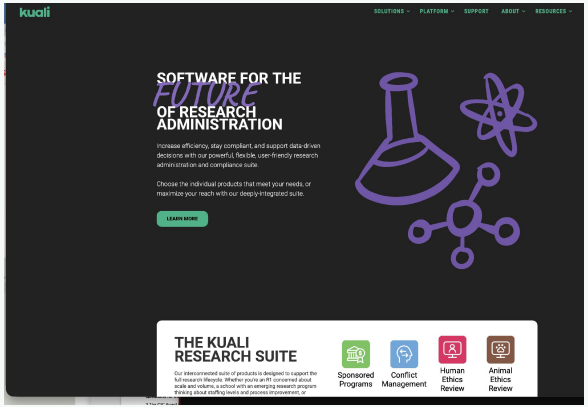
The 2026 call of the Marie Skłodowska-Curie Actions (MSCA) Postdoctoral Fellowships, part of the EU's research and innovation framework programme Horizon Europe, closed on 9 September 2026.

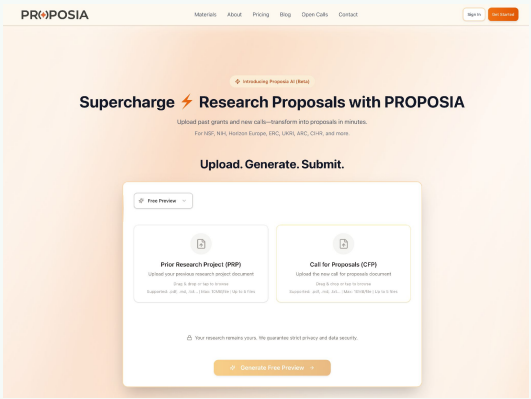
The European Research Executive Agency received **21,627 proposals**.

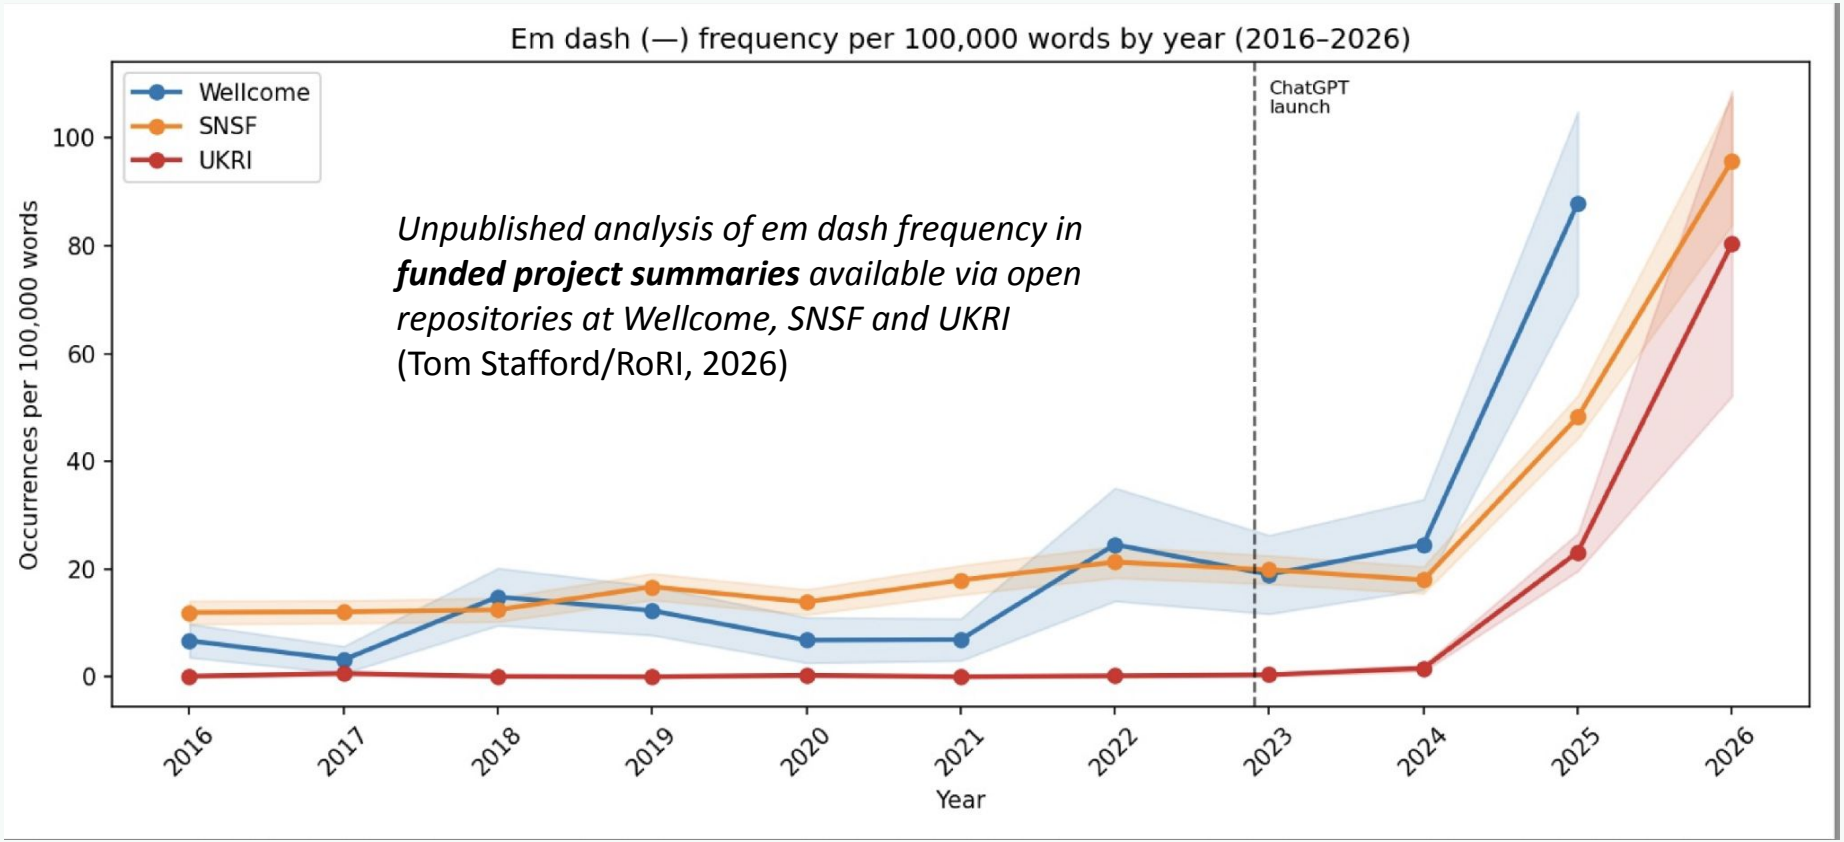




GRANTWRITING ISN'T THE PROBLEM.
EXECUTION IS.






Em dash (—) frequency per 100,000 words by year (2016–2026)
Wellcome
SNSF
UKRI
ChatGPT launch
Occurrences per 100,000 words
Year
0
20
40
60
80
100
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
*Unpublished analysis of em dash frequency in **funded project summaries** available via open repositories at Wellcome, SNSF and UKRI*
(Tom Stafford/RoRI, 2026)

Score compression

Can grant evaluation still distinguish scientific excellence?

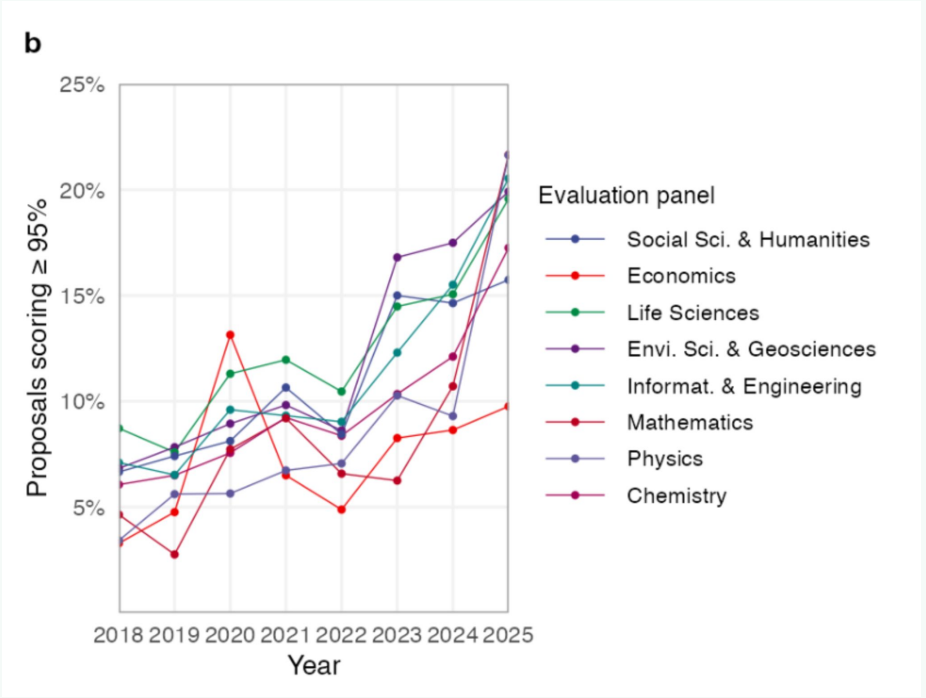
Martin Bulla[1*], James Wilsdon[2] & Peter Mikula[1]

[1]bullab, Department of Ecology, Faculty of Environmental Sciences, Czech University of Life Sciences Prague, Prague, Czechia

[2]Research on Research Institute (RoRI), Department of Science, Technology, Engineering & Public Policy, University College London; London, UK

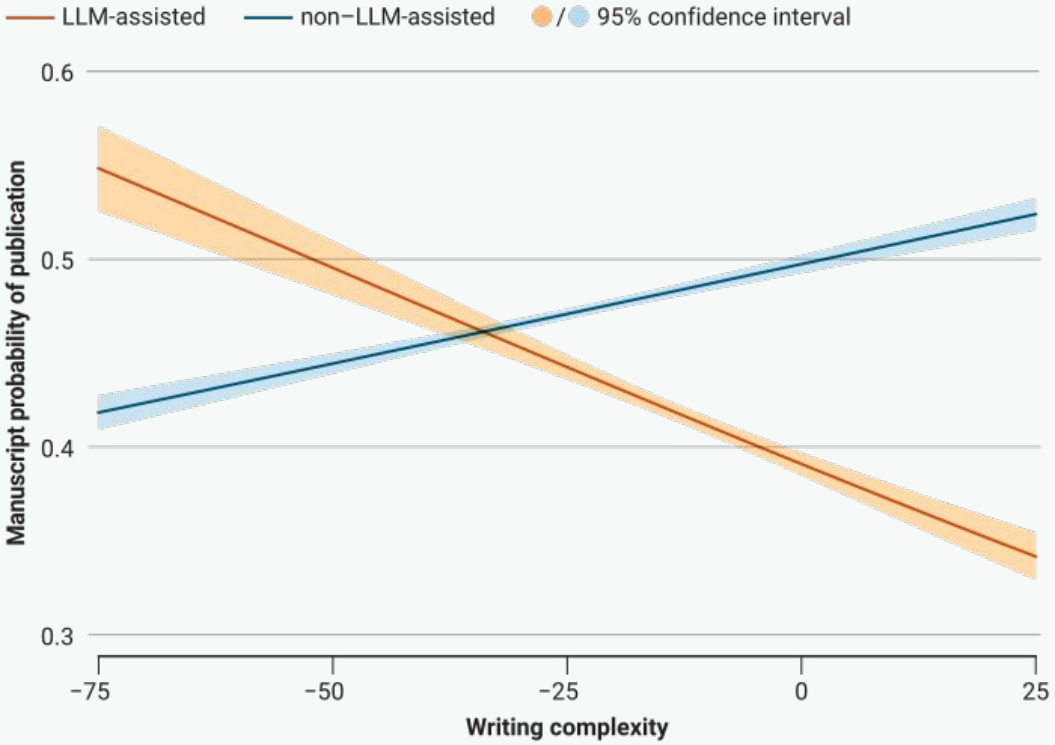
*Author for correspondence: bullam@fzp.czu.cz

Abstract. **Grant peer review increasingly depends on fine-grained numerical scores to rank many formally excellent proposals. In the era of generative and agentic AI, this architecture may be losing resolution. We analyse score-distribution data from EU *Marie Skłodowska-Curie Actions* (MSCA) Postdoctoral Fellowships and show sharp compression of evaluation scores in the 2025 call. Compared to 2024, there was a 64.6% increase in the volume of applications received, low-scoring proposals became rare, and ~20% of proposals received evaluation scores of ≥95%. The timing is consistent with AI-assisted optimization, but alternative**



Scientific production in the era of large language models

https://doi.org/10.1126/science.adw3000



Signal loss

+ Increased productivity
+ Gain greatest for non-native english speakers
+ More diverse literature cited

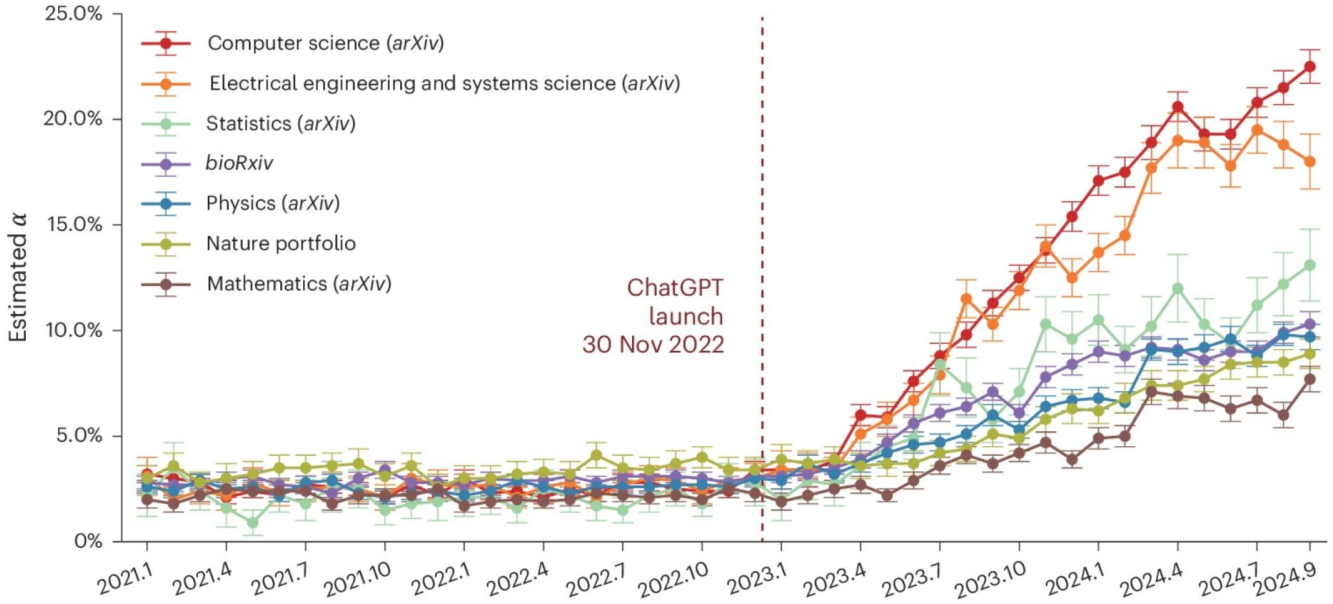
RESEARCH ON RESEARCH INSTITUTE

AI & research funding - the future

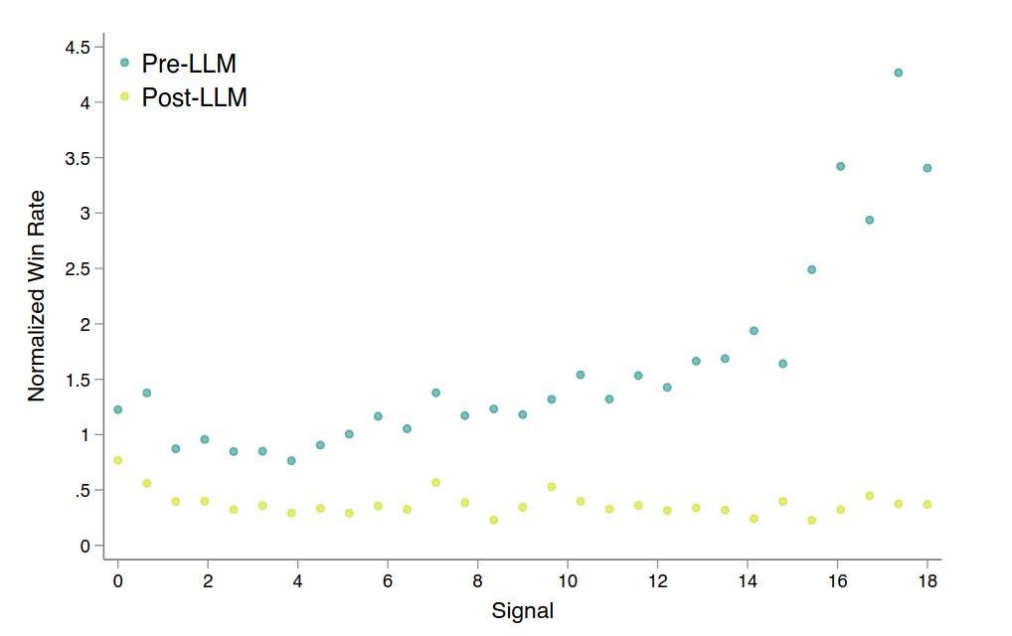
Quantifying large language model usage in scientific papers (Liang et al, 2025) https://doi.org/10.1038/s41562-025-02273-8

Fig. 1: Estimated fraction of LLM-modified sentences across research paper venues over time.

From: Quantifying large language model usage in scientific papers



Making Talk Cheap: Generative AI and Labor Market Signaling. Galdin & Silbert,(2025).arXiv preprint arXiv:2511.08785.



https://tomstafford.substack.com/p/when-trusted-signals-collapse

Diversity Loss

nature human behaviour

Article

https://doi.org/10.1038/s41562-026-02550-0

The shrinking landscape of linguistic diversity in the age of large language models

Received: 17 February 2025

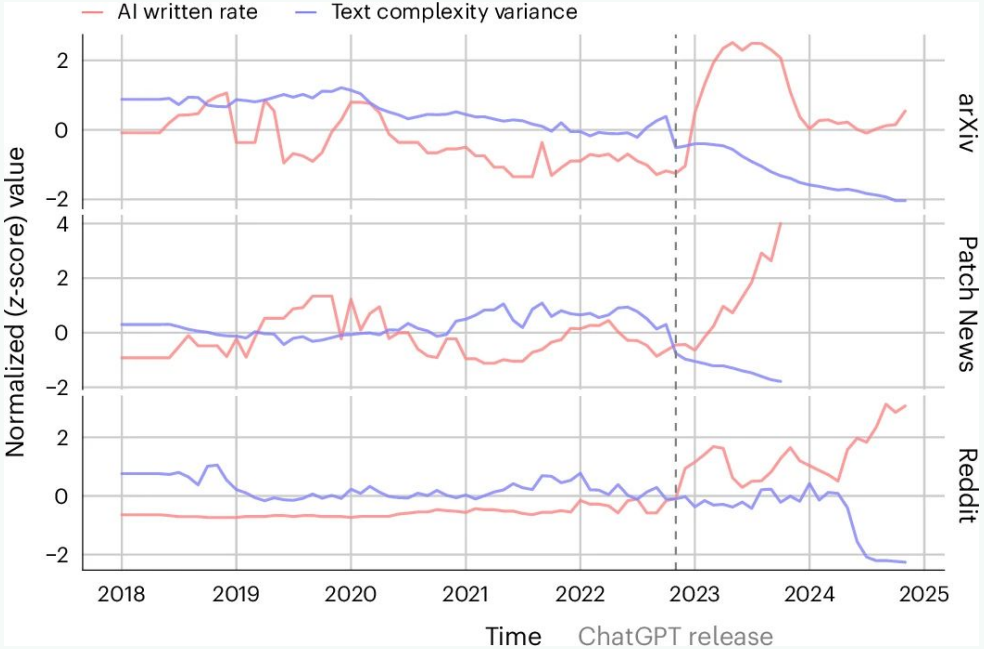
Accepted: 10 July 2026

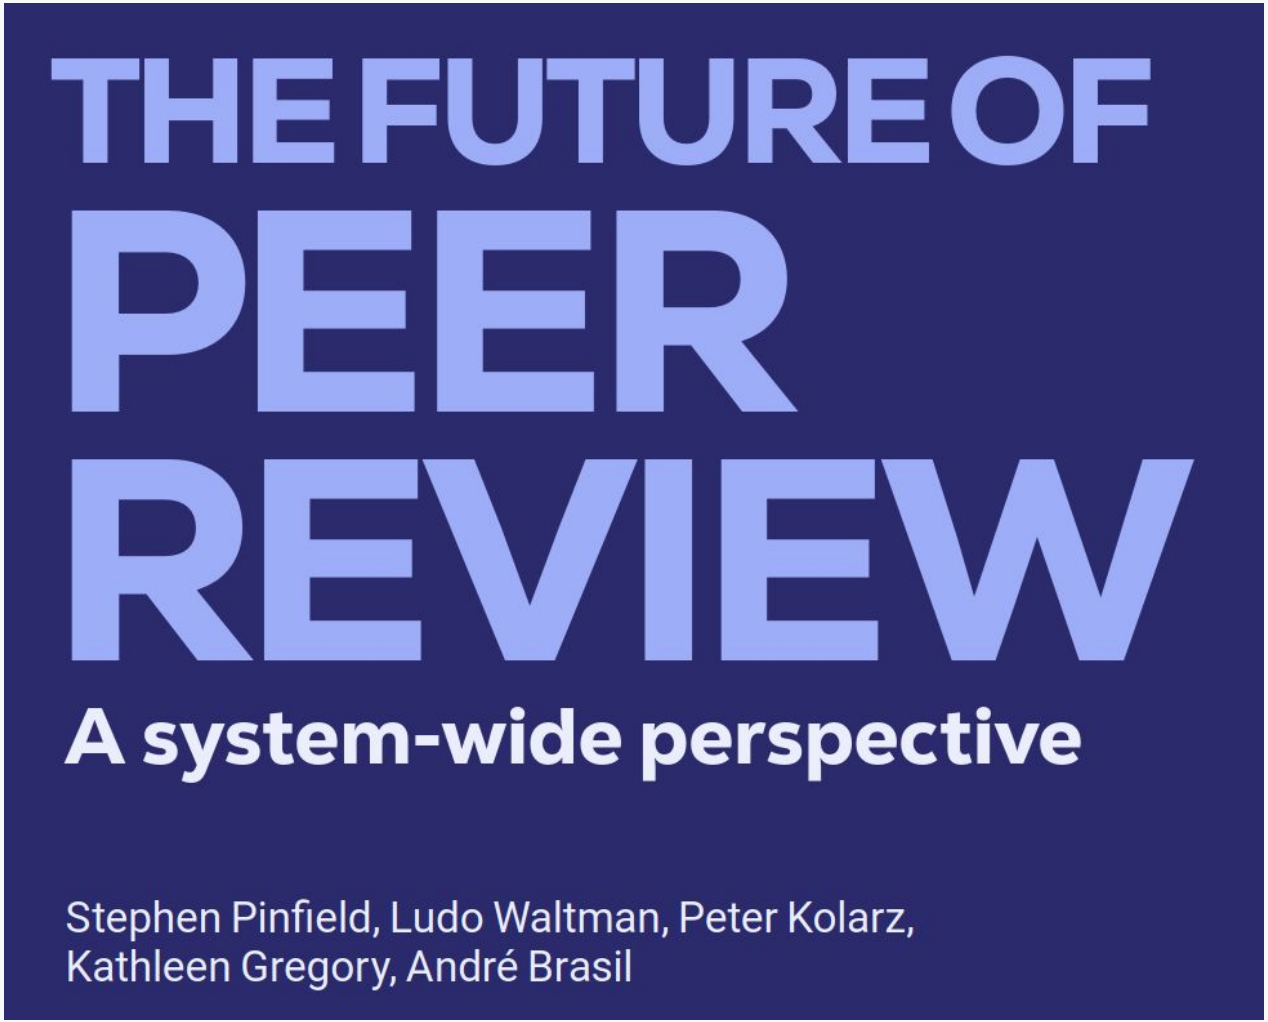
Published online: 24 August 2026

Check for updates

Zhivar Sourati[1,2], Farzan Karimi-Malekabadi[2,3,5], Meltem Ozcan[3,5], Colin McDaniel[3], Alireza Ziabari[1,2], Jackson Trager[2,3], Ala N. Tak[1], Meng Chen[3], Fred Morstatter[1,4] & Morteza Dehghani[1,2,3]

Language is far more than a communication tool; it encodes a wealth of information about a person's identity, psychological state and social context, providing valuable insights for diverse fields including psychology, marketing and healthcare. Across three studies spanning seven datasets in different domains and over 880,000 texts, we show that the widespread adoption of large language models (LLMs) as writing assistants is linked to declines in linguistic diversity, interfering with the societal and psychological insights language provides. While core content is retained when LLMs polish and rewrite texts, LLMs also homogenize writing styles, reducing writing-complexity variance by a statistically significant 21–50% across datasets and models ($P \leq 0.05$), and amplify patterns associated with dominant characteristics while suppressing others, emphasizing conformity over individuality. These trends hold across different LLMs, prompts and contexts, with potential implications for diagnostic processes, personalization efforts, hiring assessments and cultural preservation.





The Future of Peer Review:
A system-wide perspective
https://doi.org/10.6084/m9.figshare.32685351

Trends

Arms race in AI use

Demand management

Validation / verification

Trust markers of research output

Researcher communities

Pull rather than push models of funding

Absorptive Capacity

"the ability of a firm to recognize the value of new, external information, assimilate it, and apply it to commercial ends is critical to its innovative capabilities. We label this capability a firm's absorptive capacity and suggest that it is largely a function of the firm's level of prior related knowledge"

Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. Administrative Science Quarterly, 35(1), 128–152.

"That is, the performance of basic research may be thought of as a ticket of admission to an information network."

Rosenberg, N. (1990). Why do firms do basic research (with their own money)? Research Policy, 19(2), 165–174.

RESEARCH ON RESEARCH INSTITUTE

7th ANNUAL RoRI CONSORTIUM MEETING

Monday 2 November—Wednesday 4 November 2026

CosmoCaixa Museum of Science,

Carrer d'Isaac Newton, 26, 08022 Barcelona, Spain



MON 2 NOV **Agentic AI and the future of research funding (Invitation only for INFLECTION project partners[1])**

Evidence is mounting of an acceleration in AI adoption and augmentation across the research enterprise, including in grant applications and peer review. In April 2026, RoRI/UCL analysis in *Nature* of data from 12 funders found an average increase of 57% in application volumes from 2022 to 2025.[2] RoRI's Partnership Board and the UK Metascience Unit have since approved **INFLECTION**: a 15-month, intensive study with three goals:


Comment

Could agentic AI topple grant-funding systems?

- **To build a comparative dataset for real-time monitoring of the uptake and diffusion of generative and agentic AI** across diverse research funding and review contexts, in terms of application volumes, score compression, topic convergence, novelty and other indicators;
- **To use robust metascientific methods to test opportunities and limits to AI augmentation** of processes of ideation, drafting, development, peer review and evaluation of proposals;
- **To coproduce with project partners a structured menu of policy, process and strategic responses**, and consider which of these, in what combination, are most feasible.

A few places still available at the 2026 RoRI Consortium meeting – for more on agentic AI, the Meta-moment and comparative funder-engaged metascience (email: jw@researchonresearch.org)

RESEARCH ON RESEARCH INSTITUTE

Responses

RESEARCH ON RESEARCH INSTITUTE

Thank you

These slides:
bit.ly/tomstafford

Tom Stafford
t.stafford@researchonresearch.org

Join our Substack to get the latest news and views from our researchers



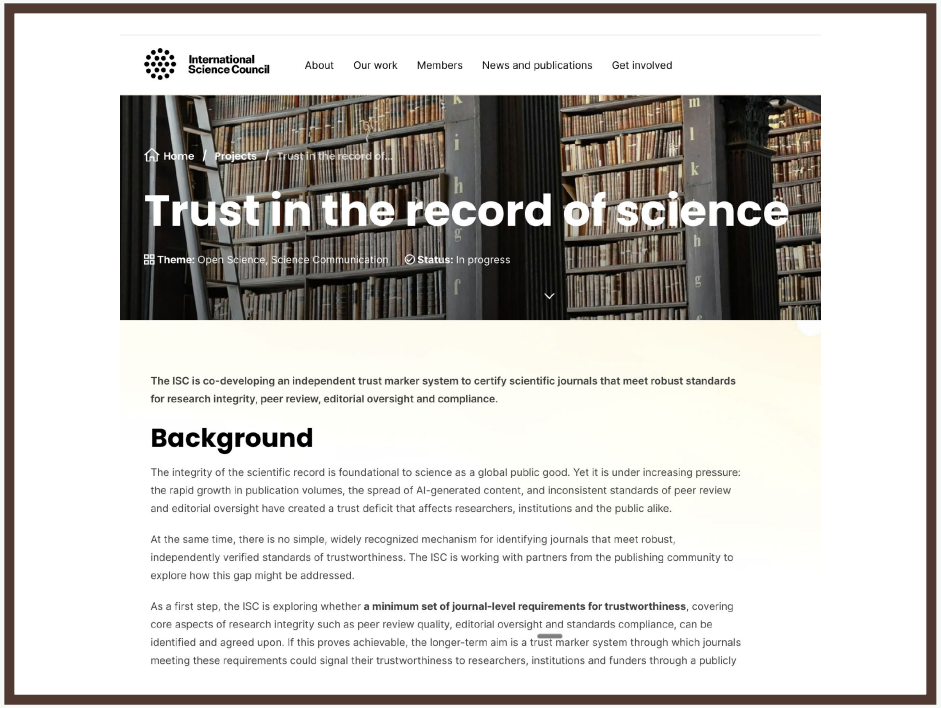
RoRI abstract for Schmidt Sciences meeting on 'Validating AI-Driven Science' (linked to RoRI's new *Inflection* project)

Conjectures are cheap, decisions are not: from push to pull in research funding

AI can now generate hypotheses faster than science can validate them. But validation capacity is finite, so someone must decide which AI-augmented ideas get research time, data and funding. That decision sits largely with research funders, whose allocation systems are already under strain. In a recent Nature article, Geraint Rees & I reported that applications to twelve global funders rose by an average of 57% between 2022 and 2025 (https://doi.org/10.1038/d41586-026-01297-). The 2026 Marie Skłodowska-Curie postdoctoral fellowships received 21,627 proposals, three times the 2022 figure, and in the previous round roughly one in five proposals scored 95% or higher. Funders are responding with disclosure rules, demand management, interviews, lotteries and AI-assisted screening. Drawing on RoRI's work with a consortium of 25 funders, I will argue that these responses all try to salvage an application-led "push" model. The more promising frontier is "pull": funders using AI to seek out whom and what to fund, and to allocate validation capacity itself, while designing out the risks of Matthew effects, model monoculture and novelty collapse.



New international initiatives to maintain trust and credibility in the scientific system





Where this goes over next 2 years, and 3 questions for funders

Five expectations

- **The written proposal loses most of its remaining discriminating power** (high confidence). Diffusion is past the median in publications; polish sways human and machine judges; a competent bid costs almost nothing.
- **AI enters review at scale as scout, not judge** (high). Expect at least one visible false-negative failure of the Nagarajappa type before 2029, and regulation shaped by it.
- **Allocable scarcity migrates from proposals to validation** (medium-high). New instruments: validation vouchers, user-facility access models, milestone-gated funding, AI data stocktakes.
- **Demand management and interviews spread, then plateau** (medium). Lotteries above threshold become routine for a minority of schemes.
- **Disclosure regimes converge on the DFG pattern** (medium), generating the first longitudinal data on AI use in review.

Questions for funders

1. Do we know our own curve? What share of our applications and reviews, involved AI agents in 2026, and how would we find out?

2. Which parts of our assessment can survive proxy collapse? Track record, interviews, verified feasibility, or none of the above?

3. Where do we redeploy assessment effort: validation capacity, shared AI tooling and bias audits with RoRI funder partners, or defending the proposal stage?

Reading about agents is not the same as studying them at work. Commission a live or shadow agentic run against a funding call and put the result in front of the Board.

RESEARCH ON RESEARCH INSTITUTE

RESERVE SLIDES FOLLOW

(not part of talk)

Labour market responses

Arms race

- **Applicants submit more applications**
- **Employeers use AI triage**

Journal responses

NATURE INDEX | 13 October 2026

Inundated AI conferences impose submissions caps

Researchers question whether these will work to stem the deluge of manuscipts.

By Dalmeet Singh Chawla

Save article



https://www.nature.com/articles/d41586-026-03153-5

Section: *News / UK / Publishing*

Bookmark Email

Warning system cracks down on peer reviewers' AI use

By Sophie Hogan

28th July 2026



https://www.researchprofessional.com/news-articles/article/1419829

NEWS / **PUBLISHING**

Scientists decry conference's use of hidden prompts to snare AI peer reviews

The invisible messages, which instruct large language models to use telltale phrases in a peer-review report, are effective in catching artificial-intelligence misuse but also erode trust, some say.

BY **DALMEET SINGH CHAWLA**
1 JULY 2026 | 4 MIN READ

HTTPS://DOI.ORG/10.53053/FCUI5388
CITE THIS ARTICLE (1)

https://www.thetransmitter.org/publishing/scientists-decry-conferences-use-of-hidden-prompts-to-snare-ai-peer-reviews/

NEWS | 19 May 2026 | Correction 19 May 2026

Researchers who use hallucinated references to face arXiv ban

The preprint server is the latest to impose stiff penalties on authors who contribute to AI 'slop' – but not everyone is convinced it's the right approach.

https://www.nature.com/articles/d41586-026-01595-5

- ACL has adopted a policy on reviewer incentives (https://aclrollingreview.org/incentives2025). Great reviewers may receive awards and prizes (such as a free conference registration). Highly irresponsible reviewers and ACs may become ineligible to submit, resubmit, or commit their work to ACL venues for the duration of the next review cycle. For reviewers, this means that reviews need to be submitted on time and in accordance with the guidelines. Note also that while reviewers are anonymous to the authors, their identity is known to the senior researchers serving as area chairs, senior area chairs, and program chairs.

https://arxiv.org/pdf/2609.19420

Use and Effects of LLMs in Peer Review: A Randomized Experiment and Survey at ICML 2026

SUNNIE S. Y. KIM*, Microsoft Research, USA
WESLEY HANWEN DENG*, Microsoft Research, USA
JENNIFER WORTMAN VAUGHAN, Microsoft Research, USA
BUXIN SU, University of Pennsylvania, USA
WEIJIE SU, University of Pennsylvania, USA
ALEKH AGARWAL, Google Research, USA
SHARON LI, University of Wisconsin-Madison, USA
MARTIN JAGGI, EPFL, Switzerland
DANIEL G. GOLDSTEIN, Microsoft Research, USA
NIHAR B. SHAH, Carnegie Mellon University, USA
MIROSLAV DUDÍK, Microsoft Research, USA

LLMs are rapidly reshaping peer review, making it important to understand how reviewers use them in practice and how different LLM-use policies affect review outcomes. We investigate these questions through a randomized experiment and an anonymous post-survey at ICML 2026, a major machine learning conference involving over 24,000 papers and 17,000 reviewers. Reviewers were assigned to either a conservative policy prohibiting all LLM use or a permissive policy allowing limited assistance, with randomization among a subset of main-track papers and reviewers. Policy assignment had near-zero effects on final paper decisions, paper scores, and reviewer confidence, although reviews under the permissive policy were 5.5-7% longer. Post-survey responses (N=1,486) revealed diverse attitudes toward LLMs and substantial noncompliance: 22.5% of conservative-policy reviewers reported using an LLM despite the prohibition, and 36.5% of permissive-policy reviewers reported at least one explicitly disallowed use. We discuss implications for future peer-review policy and tool design.

https://www.aclweb.org/portal/content/acl-sustainable-reviewing-policy

Association for Computational Linguistics

ACL Sustainable Reviewing Policy

September 11, 2026 | BY Sarvnaz Karimi

Event Notification Type: Other
Abbreviated Title:
Location:

Attachment
ACL sustainable reviewing policy_2026.pdf

- ACL has adopted a policy on reviewer incentives (https://aclrollingreview.org/incentives2025). Great reviewers may receive awards and prizes (such as a free conference registration). Highly irresponsible reviewers and ACs may become ineligible to submit, resubmit, or commit their work to ACL venues for the duration of the next review cycle. For reviewers, this means that reviews need to be submitted on time and in accordance with the guidelines. Note also that while reviewers are anonymous to the authors, their identity is known to the senior researchers serving as area chairs, senior area chairs, and program chairs.

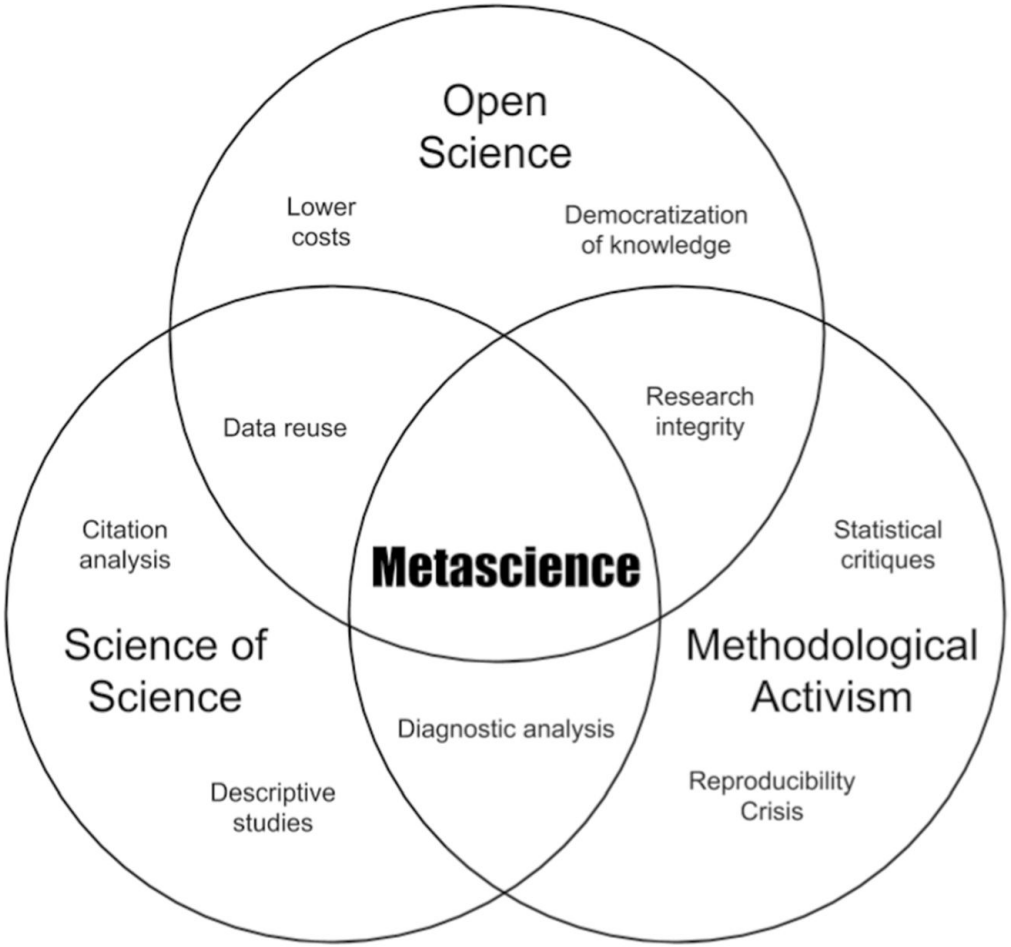


Fig. 1 The strands of metascience

Peterson, D., & Panofsky, A. (2023). Metascience as a Scientific Social Movement: D. Peterson, A. Panofsky. *Minerva*, *61*(2), 147-174. https://doi.org/10.1007/s11024-023-09490-3

More metascience

Different funding mechanisms: Distributed Peer Review, partial randomisation, golden tickets, narrative CVs,

System level effects: metrics, national research evaluation, the meta-science moment, the AI shock

Platforms: funder data platform, new publication forms (PRC)

Experimental Funders Group

https://forms.gle/TV9DdsAHHP5aUDHN7

These slides:
bit.ly/tomstafford

RESEARCH ON RESEARCH INSTITUTE



The Research on Research Institute Impact Report (2023-2025)

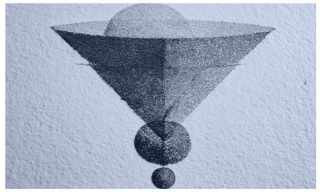
https://researchonresearch.org/roris-impact-what-weve-achieved-and-where-were-going/

Our mission is to accelerate transformative research on research systems, cultures and decision-making through metascientific methods, data and evidence.



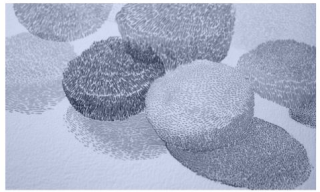
CONVENING COMMUNITY

We foster collaboration and codesign by creating spaces for funders and other stakeholders to share learning and collectively shape the future of research funding



DEVELOPING RESEARCH

We generate metascientific insights through transdisciplinary, collaborative, applied, and mixed-method research on research, responding to the strategic priorities of funding organisations



CATALYSING IMPACT

We translate evidence into action by accelerating the uptake of metascientific insights into funders' policies, strategies, and practices

- Research for Public Good
- Collaboration and Co-Design
- Innovation and Experimentation
- Evidence with Integrity

Australian Government
Australian Research Council

Claire Forsyth
Ke Zhang

FWF Austrian Science Fund

Falk Reckling

CIHR IRSC Canadian Institutes of Health Research / Instituts de recherche en santé du Canada

Alison Bourgon
Sarah Passmore

Michael Smith Health Research BC

Ellie Griffith
Kevin Sauve

NRF National Research Foundation

Gugu Moche
Nosisa Dube

novo nordisk foundation
Benefitting people and society

Gert Vilhelm Balling

CWTS Meaningful metrics

Ludo Waltman
Alex Rushforth

DIGITAL science

Jürgen Wastl
Daniel Hook

NWO

Jurgen Ganzevles
Puck Wildschut

Research England

Steven Hill

Social Sciences and Humanities Research Council of Canada
Canada

Matthew Lucas
Dominique Roche

Swiss National Science Foundation

Katrin Milzow
Srdjan Perko

The British Academy

Molly Morgan Jones
Alex Lewis

UKRI Research England

UCL

Sarah Chaytor

GORDON AND BETTY MOORE FOUNDATION

Ted Hodapp

King Baudouin Foundation
Working together for a better society

Laetitia Aerts
Sofie Bekaert

Fundació "la Caixa"

Carla Carbonell Cortés
Patricia Nadal

VolkswagenStiftung

Johanna Brumberg
Antje Tepperwien

wellcome

Alyson Fox
Kirsten Bound

ZonMw

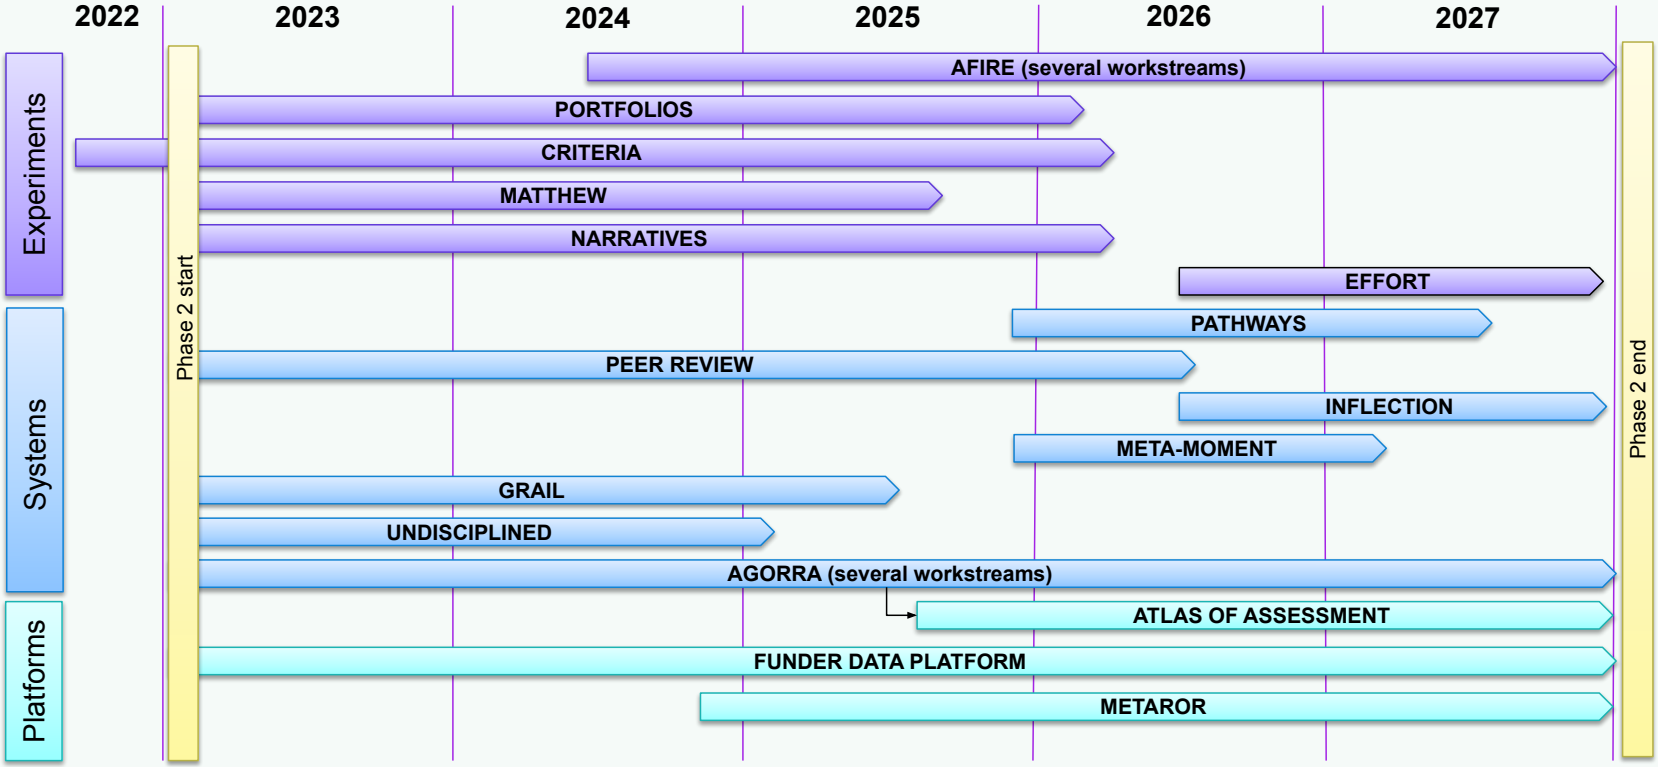
Pauline Goense
Leonie van Drooge

Our three priority areas

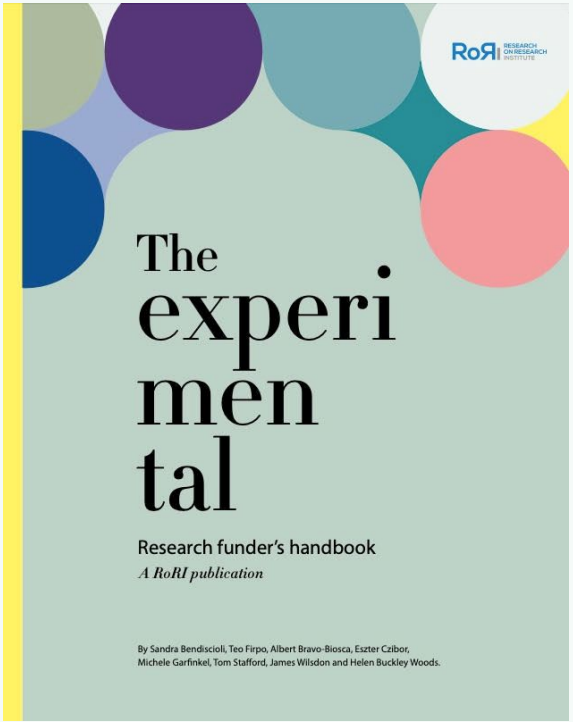
Experiments	Systems	Platforms
Controlled experiments and/or data-driven (occasionally qualitative) studies of specific elements of funding mechanisms to test effectiveness/ efficiency/ understand pros and cons.	Wider studies of the past, present or future operation of particular elements of research systems. Often more qualitative/ case study based.	Creating/ maintaining/ curating tools or facilities that can collect or hold data relevant to metascientific enquiry.

These slides:
bit.ly/tomstafford

RoRI's phase 2 projects – a timeline



We can plan/run experiments



AFIRE:

t.stafford@researchonresearch.org

josie.coburn@ucl.ac.uk

The experimental research funder's handbook (Revised edition, June 2022, ISBN 978-1-7397102-0-0). https://doi.org/10.6084/m9.figshare.19459328.v2

These slides:

bit.ly/tomstafford

What makes a RoRI project?

1. It has to involve a transdisciplinary/collaborative dimension between researchers and funder, e.g.
   - Analysis using internal funder data
   - Experiments on funders' internal systems
   - Analysis of funders' processes, strategies or tools

2. It has to involve multiple funders so that there is a comparative/ international dimension
   - Beyond single-funder consultancy projects
   - 'strength in numbers'
   - More robust & widely relevant findings

Accelerator For Innovation & Research Funding Experimentation (AFIRE)

1. Awareness
2. Motivation
3. Capacity
4. Scale and reach
5. The culture

https://researchonresearch.org/project/a-f-i-r-e/

These slides:
bit.ly/tomstafford

OUR DEFINITION OF EXPERIMENT

Principled

A research design that allows inference about what causes what (before/after, shadow experiments, true experiment/RCT).

Planned

Primary outcome measure and analysis plan declared in advance.

Public

A commitment to sharing the results regardless of outcome.

These slides:
bit.ly/tomstafford

EXPERIMENTS

Distributed Peer Review, Partial Randomisation, Desk Rejection, and robust quantitative pilots.

CAPACITY BUILDING

Sprints on AI/ML in reviewer selection. Equipping funding institutions with execution frameworks.

FORUM

Sharing work by funders, for funders. Building global awareness and initial motivation to experiment.

https://researchonresearch.org/project/a-f-i-r-e/

These slides:
bit.ly/tomstafford

AFIRE | 2024-25 Funder Forum Series

Villum Experiment – a grant scheme with double-blind evaluation
Anders Smith, Villum Foundation

Predicting Progress: A Pilot of Expected Utility Forecasting in Science Funding
Alice Wu, Federation of American Scientists & Jordan Dworkin, Open Philanthropy

Enhancing EDI in application processes in the CES Transformation Fund
Rachel Kendal, Principal Investigator CES Transformation Fund

Partial Randomisation at The British Academy
Ken Emond, The British Academy & Adrian Barnett, Queensland University of Technology

Do Grant Proposal Texts Matter for Funding Decisions? A Field Experiment at NWO
Müge Şimşek, University of Amsterdam

An experiment with Distributed Peer Review
Hanna Denecke, The Volkswagen Foundation

Two-stage versus continuous application processing
Jan-Ole Hesselberg, Stiftelsen Dam

Recordings

https://www.youtube.com/@researchonresearchinstitut6175

FORUM | Experimental Funders Group, 2026

Shorter talks, more time for discussion

13 May 16:00 CEST
AI in Funding

17 June 16:00 CEST
Distributed Peer Review

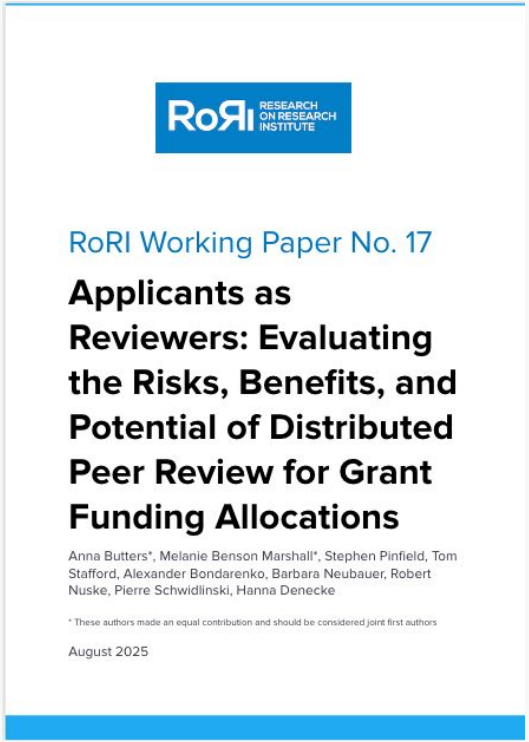
16 December 09:00 CEST
‘Lottery First’ methods

24 February 16:00 CEST
Desk Rejection

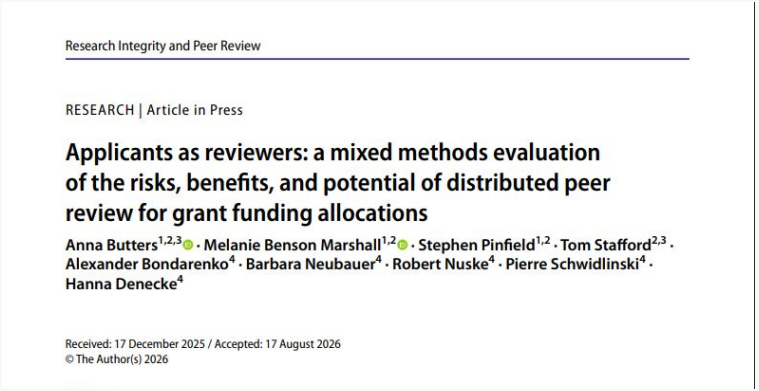
https://forms.gle/TV9DdsAHHP5aUDHN7



https://doi.org/10.6084/m9.figshare.29270534.v1



https://doi.org/10.6084/m9.figshare.29994841



https://doi.org/10.1186/s41073-026-00246-7

New work from AFIRE



https://researchonresearchinstitute.substack.com/



https://rori.figshare.com/articles/report/_b_Demand_management_findings_from_a_multi-funder_workshop_b_/32044422?file=64292685

New projects!

- Desk rejection
- EFFORT
- Researcher attitudes to funder experimentation
- Understanding DPR
- Agentic AI and funding