

Chapter 8

How Research Works: Funding, Ethics, Evaluation, and Openness

Framing the chapter

The previous chapter described who does research and the institutions that support it. This chapter turns to the practices that determine whether research happens at all, and how it is judged once it does: how a grant application is actually built and reviewed; the norms – and the well-documented failures of those norms – that define honest research conduct; how researchers and their publications are themselves evaluated, by peers and by numbers; how that work is presented at conferences, genuine and otherwise; and the open-science expectations that increasingly govern how the resulting data and code must be shared. None of what this course has traced across six chapters of discovery happens today without passing through this machinery.

8.1 Funding schemes: national and European

Research funding in Europe runs through two broad, overlapping layers. **National funding agencies** – the Agence nationale de la recherche (ANR) in France, the Deutsche Forschungsgemeinschaft (DFG) in Germany, UK Research and Innovation (UKRI) in Britain, among many others – fund projects based within their own country, typically through annual or twice-yearly thematic or generalist calls. **European-level schemes**, funded through the Horizon Europe framework programme, include the individually held European Research Council (ERC) grants and the mobility-focused Marie Skłodowska-Curie Actions introduced in the previous chapter, as well as large multi-partner “collaborative” projects that fund consortia of several institutions across different member states working on a shared work programme.

The ERC in particular structures its grants explicitly by career stage, which makes it a useful illustration of how funding schemes track the career pipeline described in the previous chapter: a **Starting Grant**, for researchers a small number of years after the PhD (traditionally a window of roughly two to seven years, though exact eligibility rules and extensions – for career breaks, for instance – are periodically revised), funds up to EUR 1.5 million over five years; a **Consolidator Grant**, for a somewhat later career stage (traditionally around seven to twelve years post-PhD), funds up to EUR 2 million over five years; an **Advanced Grant**, with no fixed career-stage limit but requiring a substantial track record, funds up to

EUR 2.5 million over five years; and a **Synergy Grant** funds a group of two to four principal investigators jointly, up to EUR 10 million over six years, for questions that genuinely require combining distinct expertises rather than a single investigator's project scaled up. Every scheme, national or European, PhD-, postdoctoral-, or senior-researcher-facing, is defined by a call document – sometimes called a *lettre de cadrage* or scope note in French usage – that fixes the eligible topics, career-stage or nationality restrictions, funding ceiling, and submission deadline that every application to that call must satisfy; reading this document closely, rather than reusing a proposal written for a different call, is the first and most commonly overlooked step of putting together a competitive application.

Historiographical note

The amounts, durations, and eligibility windows above are current, indicative ERC figures and change over successive Horizon Europe work programmes; a candidate should always verify the current call text rather than relying on a fixed number from any single year, including this one. The broader point this section is making is structural rather than numerical: European funding is organised as a graded set of instruments tied to career stage, from doctoral and postdoctoral fellowships through to senior single-investigator and multi-investigator grants, a structure that mirrors, and partly funds, the career pipeline described in the previous chapter.

8.2 Anatomy of a grant proposal

However much funding schemes differ in scale and eligibility, the document a successful applicant actually submits tends to contain a similar set of components, whether the call is a national ANR call or a European ERC or collaborative call. This section lists the most consistent ones; the exact required sections, page limits, and templates always come from the specific call document, not from this list.

Scientific context and state of the art situates the proposed work against existing research, states explicitly what gap or open question it addresses, and – particularly for the more competitive schemes – makes an explicit case for originality and feasibility rather than simply summarising the field.

Methodology and work plan, usually broken down into a small number of *work packages* (WPs), each with its own objectives, tasks, and expected deliverables, is normally accompanied by a timeline diagram – most often called, after the American engineer Henry Gantt who popularised the format in the early twentieth century, a **Gantt chart** – showing how the work packages are scheduled and staged relative to one another over the life of the project.




Figure 8.1: A schematic Gantt chart of the kind expected in a multi-year grant proposal's work plan: illustrative work packages and milestones only, not a real project.

Most current calls, and essentially all Horizon Europe calls, additionally require a **data management and open science section**, specifying which data the project will generate, in what format, under what licence, and where it will ultimately be deposited (the repositories discussed later in this chapter are the normal answer); a **budget justification**, itemising personnel, equipment and consumables, travel, and the institution's overhead rate, with each line tied back to a specific task in the work plan rather than presented as a lump sum; and a **bibliography**. Depending on the scheme, further sections can include an ethics self-assessment (required whenever the project involves human participants, animals, or dual-use technology), a statement of the gender dimension of the research where relevant, and a dissemination and communication plan distinct from the open-science section, describing how results will reach audiences beyond the immediate scientific community.

Once submitted, a proposal is normally assessed by a panel of external peer reviewers against published criteria – for Horizon Europe schemes, typically excellence, impact, and quality of implementation – and, in some schemes, an applicant is given a rebuttal stage to respond in writing to reviewer criticisms before a final funding decision is made.

8.3 Research ethics: from questionable practices to outright fabrication

Research misconduct is not a single thing but a spectrum, and one of the more important distinctions a working researcher needs to be able to make is between practices that are merely questionable and practices that are unambiguously fraudulent. At one end of that spectrum sit what the research-integrity literature calls **questionable research practices**: selectively reporting only the results that support a hypothesis while quietly leaving out an inconvenient data point or failed trial, slicing one dataset into the minimum number of publishable papers to maximise a publication count ("salami-slicing"), or listing a colleague as

an author who made no real contribution to the work (“honorary” or “gift” authorship). These practices are common enough, and their boundaries contested enough, that they are rarely prosecuted as misconduct in a formal sense, even though they distort the scientific record in real ways. At the other end sit the three categories long recognised, in most national and institutional research-integrity policies, as outright **misconduct**: *fabrication* (inventing data or results that were never actually observed), *falsification* (manipulating research materials, equipment, processes, or data such that the research is not accurately represented), and *plagiarism* (presenting another person’s ideas, words, or results as one’s own without proper attribution). The distance between quietly omitting one inconvenient data point and inventing an entire dataset from nothing is exactly the distance this spectrum is meant to capture.

The Schön scandal remains one of the most thoroughly documented cases of outright fabrication in physics, precisely because it was investigated so publicly. Between 2000 and 2002, Jan Hendrik Schön, a young physicist working at Bell Labs (then part of Lucent Technologies) in New Jersey, published a rapid sequence of high-profile results in organic semiconductors and molecular-scale electronics – including claimed molecular transistors and superconductivity in organic crystals – in *Science*, *Nature*, and the *Physical Review* journals, at a rate that briefly made him one of the most cited young physicists in the world. The fraud was uncovered in April 2002, when colleagues reviewing one of his papers for a patent application noticed that two supposedly independent datasets, describing different devices, were in fact numerically identical; a five-member investigative committee, chaired by the Stanford physicist Malcolm Beasley and convened by Bell Labs in May 2002, went on to document that entire noise traces had been reused across experiments purportedly run at very different temperatures – a physical impossibility – and that several of the graphs presented as experimental data had in fact been generated from smooth mathematical functions rather than any real measurement. Schön kept no laboratory notebooks and had, by his own account, erased the relevant computer files and discarded the physical samples in question, leaving no way to independently verify or refute most of his claims after the fact. The committee’s September 2002 report substantiated at least sixteen of twenty-four specific allegations of scientific misconduct and found that the fabrication was Schön’s alone – his co-authors, including his immediate supervisor and frequent collaborator Bertram Batlogg, were cleared of fabricating data themselves, though the episode prompted a wider and still-unresolved debate about how much responsibility a senior co-author bears for not having independently verified a junior colleague’s results before attaching their own name to them. Twenty-nine papers were formally retracted across the five journals involved between late 2002 and mid-2003; the University of Konstanz revoked Schön’s doctorate in 2004 (a decision Germany’s Federal Constitutional Court upheld in 2014), and the Deutsche Forschungsgemeinschaft barred him for eight years from applying for funding, serving on review panels, or peer reviewing on its behalf.

Historiographical note

The Schön case is unusually well documented because Bell Labs ran, and fully published, a formal investigation – most suspected misconduct is resolved far less visibly, if it is resolved at all, and detected fraud is understood in the research-integrity literature to be a substantial undercount of

the true rate, precisely because a well-executed fabrication, by definition, is not straightforward to distinguish from a genuine and merely surprising result. The case is also sometimes taught alongside other large, publicly investigated fabrication cases from outside physics – for instance the Dutch social psychologist Diederik Stapel, found in 2011 to have fabricated data in dozens of published studies over more than a decade – specifically because the cross-disciplinary recurrence of the same underlying pattern (isolated access to raw data, high publication pressure, and insufficiently sceptical co-authors and reviewers) suggests the vulnerability is structural rather than confined to any one field.

A far more recent case, from the same broad field as Schön's, shows the same pattern recurring two decades later under intensified scrutiny. Between 2020 and 2023, the physicist **Ranga Dias**, at the University of Rochester, published a sequence of high-profile claims of room-temperature superconductivity: first in a carbonaceous sulfur hydride under extreme pressure (*Nature*, October 2020), and then, after that first paper's retraction, in a nitrogen-doped lutetium hydride he reported as displaying near-ambient-pressure superconductivity (*Nature*, March 2023), a claim that briefly drew the same order of international attention Schön's had two decades earlier. Both *Nature* papers were retracted – the first in 2022, after outside physicists identified irregularities in the reported resistance data that could not be independently reconciled with the raw measurements; the second in November 2023, at the request of most of Dias's own coauthors. A separate 2021 paper in *Physical Review Letters* was also retracted, in mid-2023, after an independent review found what the journal described as convincing evidence of data fabrication. A University of Rochester investigation, informed by a review commissioned by the US National Science Foundation, concluded in March 2024 that Dias had engaged in research misconduct, including fabrication and falsification of data and, separately, plagiarised text and figures in his own grant applications; the university dismissed him in November 2024, by which point five of his papers had been retracted in total.

Historiographical note

Not every headline-grabbing superconductivity claim of the 2020s was fraudulent, and the contrast is instructive. In July 2023, a South Korean team posted two preprints – not submitted to, or reviewed by, any journal – describing a room-temperature, ambient-pressure superconductor they called LK-99. Unlike Schön's or Dias's papers, LK-99 was not fabricated data dressed up as a discovery: it was a genuine, if seriously incomplete, experimental result, released directly onto a preprint server and amplified almost immediately by social media, well beyond the specialist community actually equipped to evaluate it. What followed was peer review's underlying logic operating by other means: dozens of laboratories worldwide attempted to replicate the result within weeks, entirely outside any formal editorial process, and converged within a few months on a mundane explanation, made definitive in a detailed replication study published that November – a ferromagnetic impurity phase, largely copper sulfide, present in the samples produced a sharp drop in electrical resistance near body temperature that superficially resembled a superconducting transition, but was, on closer examination, an ordinary structural phase change with no connection to superconductivity at all. The episode is a caution about how far an unreviewed claim can now travel, via preprint and social media, before anyone competent has checked it; it is equally a demonstration that open, rapid, decentralised scrutiny by the wider research community – exactly

the function formal peer review is meant to serve, more slowly and more narrowly – can still catch an erroneous claim within months rather than years, and did so here without any journal’s peer review ever having been invoked at all.

Institutional responses to this spectrum of misconduct have themselves professionalised considerably since the Schön case: most research institutions and funders now maintain a formal research-integrity office and a defined procedure for investigating an allegation, journals increasingly follow shared guidelines from bodies such as the Committee on Publication Ethics (COPE) for handling suspected misconduct in a submitted or published paper, and independent projects such as Retraction Watch maintain public, searchable databases of retracted papers specifically to make the historical record of detected fraud more visible and harder to quietly forget.

The question of whether fabrication actually occurred is not always settled quickly, and it is not confined to the cases modern research-integrity offices formally investigate. Gregor Mendel’s foundational 1866 paper on inheritance in pea plants became the subject of exactly this question seventy years after his death. In 1936 the statistician Ronald Fisher re-examined Mendel’s published ratios and found them to fit the theoretical Mendelian expectations implausibly well – too well, Fisher argued, to be the product of ordinary biological variation and sampling, concluding that the data had, somewhere in the process, been adjusted toward the expected result, whether by Mendel himself or by an assistant tending the plants. Fisher’s own statistical standing was such that this conclusion was repeated in genetics textbooks for the better part of a century essentially as settled fact.

Historiographical note

A 2019–2020 statistical re-audit, led by the geneticist Noel Ellis, went back to Fisher’s own chi-squared calculation rather than to Mendel’s original data, and found the accusation itself did not hold up: Fisher had pooled results from several of Mendel’s separate experiments – conducted under different conditions and, in at least one case, involving traits Mendel could not have known were genetically linked rather than independent – into a single aggregate test, in a way that manufactured much of the apparent statistical implausibility. Properly disaggregated, Mendel’s individual experimental results are not unusually close to expectation at all. This does not prove Mendel’s data were entirely free of the ordinary, low-level classification bias any nineteenth-century plant-counting experiment might carry; it shows something narrower, and for this course’s purposes more important: an accusation of scientific misconduct is itself an empirical claim, resting on someone’s stated method and assumptions, and is no more entitled to a free pass from scrutiny than the original result it targets – precisely the caution Chapter 6 raised, in the opposite direction, about the risk of overcorrecting a genuine pattern of historical erasure.

8.4 Publishing and being reviewed

A result generally only counts, in the ordinary working sense a researcher uses the word, once it has passed peer review and appeared in a journal: independent researchers in the same field – referees, typically two or three per submission and, for many physics and related journals, coordinated by a handling editor often called an **Associate Editor** – assess a submitted

manuscript for soundness, originality, and significance before an editor decides whether to publish it, reject it, or request revisions. Journals themselves range from broad, general-audience venues such as *Nature* or *Science* to narrowly specialised, discipline-specific titles whose readership and citation practices differ enormously, a distinction that matters directly for the bibliometric measures discussed in the next section, since raw citation counts are not comparable across venues of such different scope. Open-science norms increasingly shape this process directly: most funders and many journals now expect or require a data-availability statement specifying where the underlying data can be found, and posting a preprint – an unreviewed draft, publicly available before or during peer review, on a server such as arXiv in physics – has become standard practice in several fields specifically to make results available for scrutiny well before the (often lengthy) formal review process concludes.

Alongside genuine peer review of the science, a growing number of journals now also require an **author contribution statement**, itemising exactly which of a set of standard roles – conceiving the study, developing its methodology, carrying out the experiments, performing the formal analysis, writing the original draft, supervising the work, securing its funding, among others – each named author actually performed. The most widely adopted version, CRediT (Contributor Roles Taxonomy), defines fourteen such roles and grew out of a 2012 workshop convened by Harvard University and the Wellcome Trust; it is now built directly into the submission systems of many major publishers. The requirement is a direct, structural response to the honorary- and gift-authorship problem raised earlier in this chapter: a statement that lists a co-author only under, say, “Writing – Review & Editing,” with no substantive research role attached, makes an unearned byline considerably harder to sustain quietly, and gives a journal, an institution, or a later investigator a documented basis for asking what a given co-author actually did on a paper under scrutiny.

Historiographical note

Not every venue that looks like a peer-reviewed journal actually functions as one. A **predatory journal** charges publication fees – often only revealed after a manuscript has already been “accepted” – while providing little or no genuine peer review, sometimes fabricating an editorial board, an impact factor, or an indexing status it does not actually hold. The librarian Jeffrey Beall compiled an influential, if also criticised, public blacklist of such publishers from around 2010 until discontinuing it in 2017 under legal pressure; the phenomenon itself has not gone away, and identifying it now falls, in practice, on individual researchers, journal indexing services, and institutions checking a candidate’s publication list, rather than on any single maintained authority.

8.5 Bibliometrics: counting citations, and their limits

Beyond judging individual papers, institutions, funders, and hiring and promotion committees also need some way of comparing researchers’ overall track records, and have increasingly reached for quantitative citation-based measures to do it. The **journal impact factor**, introduced by Eugene Garfield and still calculated and published (via Clarivate’s Journal Citation Reports) as the average number of citations articles in a given journal receive within a defined window, was designed to compare journals, not individual articles or the researchers who

wrote them, though it is very widely – and, as the next paragraph describes, controversially – used as exactly that kind of proxy in practice. The **h-index**, proposed by the physicist Jorge Hirsch in 2005, instead measures an individual researcher directly: a researcher has an h-index of h if h of their papers have each received at least h citations, a single number meant to reward both productivity and impact simultaneously while resisting distortion by one or two exceptionally highly cited papers. Both measures share the same basic limitation: citation practices differ enormously between fields (a highly cited mathematics paper and a highly cited biomedical paper are not on the same numerical scale at all), both measures reward career length more than they reward any specific discovery, and both can be, and are, gamed – through excessive self-citation, citation cartels among cooperating authors, or, in the specific case of the impact factor, editorial pressure on authors to cite the journal’s own recent articles before publication.

These well-documented limitations produced a formal, organised pushback: the **San Francisco Declaration on Research Assessment (DORA)**, drafted at the December 2012 annual meeting of the American Society for Cell Biology in San Francisco, explicitly recommends that funders, institutions, and journals stop using the journal impact factor as a proxy for the quality of an individual article, or of the researcher who wrote it, in hiring, promotion, or funding decisions, and instead assess the scientific content of a candidate’s actual work. Thousands of institutions and individual researchers worldwide have since signed the declaration, including, in France, the CNRS in 2018; implementing that recommendation in practice – replacing a single familiar number with a more qualitative, and inevitably more time-consuming, reading of a candidate’s actual publications – has proceeded considerably more unevenly than the number of signatures alone might suggest.

Historiographical note

DORA is the most widely cited reform initiative in this area, but it is a declaration of intent and a set of recommendations, not an enforceable standard, and the gap between an institution’s having signed it and an institution’s evaluation committees having actually changed how they read a CV in practice is itself a live, ongoing subject of research-policy debate rather than a settled outcome. This course states the declaration’s aims plainly, and states just as plainly that adoption in practice remains partial and uneven across the institutions that have formally signed on.

8.6 Conferences: formats, prestige, and predatory imitations

Alongside journal publication, the conference remains the other principal venue where research results are presented, and conference contributions are structured in a fairly consistent hierarchy across disciplines. A **contributed talk** or **poster**, the most common form of participation, is submitted as an abstract, reviewed by the organisers (often lightly, by scope and relevance rather than full peer review), and slotted into one of many simultaneous **parallel sessions** grouped by sub-topic, typically for ten to fifteen minutes. An **invited talk** is different in kind rather than degree: the organisers themselves select and directly invite a specific speaker, usually for a longer slot, on the strength of that person’s existing standing in the sub-field, rather than the speaker having submitted anything for review. A **keynote**

or **plenary** talk sits at the top of that hierarchy: a single session, with no competing parallel sessions at all, typically given to a field’s most prominent figures or to that year’s major prize recipients; some larger conferences additionally schedule **semi-plenary** sessions, drawing a large but not conference-wide audience, as an intermediate format between a plenary and an ordinary parallel session.

The **APS March Meeting** of the American Physical Society, now generally run jointly with the (northern-hemisphere) spring April Meeting as a single “Global Physics Summit,” is one of the largest single physics conferences in the world and a useful illustration of this structure at scale: thousands of contributed talks are organised into many dozens of simultaneous parallel and focus sessions running across the meeting’s several days, alongside separately curated invited-talk sessions and a small number of plenary and prize sessions – including, most years, the address by that year’s major APS prize or medal recipients – that draw the entire meeting into one room.

Historiographical note

The conference landscape has its own version of the predatory-journal problem: a **predatory conference** solicits attendance and paper submissions indiscriminately, often merging several unrelated scientific fields into a single nominal “conference” purely to maximise registration-fee revenue, provides little or no genuine peer review of submitted abstracts, and frequently misrepresents its own organising committee or academic sponsorship. Organisations repeatedly identified in the research-integrity literature and by regulators as running events of this kind include WASET and OMICS International; the latter was specifically fined by the United States Federal Trade Commission in 2019 for deceptive practices in both its predatory journals and its predatory conferences. As with predatory journals, an unsolicited invitation that flatters a researcher’s expertise in oddly generic terms, an event covering an implausibly broad range of unrelated topics, and registration fees that are large relative to the event’s actual scientific programme are the most reliable warning signs.

8.7 Open science: FAIR data, persistent identifiers, and repositories

The open-science expectations referenced earlier in this chapter, in the discussion of both grant proposals and journal publishing, rest on a shared technical and normative foundation usually summarised by the **FAIR** principles: that research data and the software used to produce it should be **F**indable, **A**ccessible, **I**nteroperable, and **R**eusable. The principles were formulated at a 2014 workshop at Leiden University’s Lorentz Center and published in 2016, by Mark Wilkinson and a large group of co-authors, in the journal *Scientific Data*; “open” and “FAIR” are related but distinct in an important way that the acronym itself does not make obvious – FAIR data must be well enough described and identified to be found and correctly reused even when it is not, for legal, ethical, or commercial reasons, fully open to the public.

Findability and reusability in the FAIR sense depend heavily on **persistent identifiers**: just as a published article is now routinely assigned a Digital Object Identifier (DOI) through

Crossref, a dataset, and increasingly a specific version of a piece of analysis software, can be assigned its own citable DOI through the parallel DataCite system, which lets a dataset be formally cited, and its use tracked, in exactly the way a journal article already is. That identifier only does useful work, in turn, if the object it points to is actually deposited somewhere durable: general-purpose repositories such as Zenodo (hosted by CERN and linked to the European OpenAIRE initiative, and free to use for research outputs of essentially any discipline) sit alongside discipline-specific repositories built around one particular kind of data – the Protein Data Bank for macromolecular structures, the Crystallography Open Database for crystal structures, GenBank for nucleotide sequences – and national infrastructures such as HAL in France, an open repository for preprints and author manuscripts across all disciplines, or Software Heritage, an initiative hosted by Inria that aims to archive, indefinitely and independently of any single hosting platform, the full history of publicly available source code. A data management plan, describing which of these repositories a given project will actually use and under what licence, is – as the earlier section on grant proposals noted – now a standard, and often mandatory, component of the funding application itself, which closes the loop this chapter has traced from how research is funded to how, once funded and completed, its results are made available for anyone else to find, verify, and build on.

Looking ahead

Together, this chapter and the one before it have described the practical machinery this course's first six chapters left implicit: the careers and institutions research runs on, and the funding, integrity, evaluation, and openness norms that determine which research gets done, which of it can be trusted, and how it reaches the rest of the scientific community and the public once it is finished. None of these norms is static – funding instruments, integrity procedures, and open-science requirements have all changed substantially within the past two decades, and are continuing to change – and situating any individual scientific result this course has discussed within the institutional conditions of its own time, exactly as earlier chapters have done for archives, journals, and academies, applies just as much to the research being done today as to any of the historical cases this course has examined.

Further reading

- European Research Council, *ERC Work Programme* (Horizon Europe, current call year), erc.europa.eu.
- Agence nationale de la recherche, *Guide au montage de dossier* (current call year), anr.fr.
- Committee on Publication Ethics (COPE), *Core Practices* and case-report database, publicationethics.org.
- Eugenie Samuel Reich, *Plastic Fantastic: How the Biggest Fraud in Physics Shook the*

Scientific World (New York: Palgrave Macmillan, 2009) – a full account of the Schön case.

- Jorge E. Hirsch, “An Index to Quantify an Individual’s Scientific Research Output,” *Proceedings of the National Academy of Sciences* 102, no. 46 (2005): 16569–16572.
- San Francisco Declaration on Research Assessment, full text and signatory list, sf-dora.org.
- Mark D. Wilkinson et al., “The FAIR Guiding Principles for Scientific Data Management and Stewardship,” *Scientific Data* 3 (2016): 160018.
- Ivan Oransky and Adam Marcus, Retraction Watch (blog and database), retraction-watch.com.