

Chapter 9

Writing and Publishing a Research Article

Framing the chapter

The previous chapter examined the machinery that decides which research gets funded, how misconduct is defined and investigated, and how peer review, citation counts, and open-science norms judge a piece of research once it exists. This chapter turns to the document itself: the artefact – the research article – through which a result actually enters the scientific record. It walks through the anatomy of a typical article section by section; the tools, principally L^AT_EX and the collaborative platform Overleaf, that most of physics and adjacent fields now use to typeset one; the surprisingly different molds the same underlying science must be poured into depending on the target journal; and the concrete, often bureaucratic steps – submission portals, conflict-of-interest forms, data-availability statements, and finally proofs – that stand between a finished manuscript and a published paper. Chapter 8’s discussion of peer review, author-contribution statements, and open data supplies the policy background this chapter assumes; here the focus shifts to the practical question a first-time author actually faces: what, concretely, has to be produced, in what order, and to whose rules?

9.1 Front matter: title, authors, and abstract

The title is, in practice, the most consequential sentence of the whole article: it is what appears in a search result, a citation list, a conference programme, or a colleague’s skim of a table of contents, and for the great majority of readers it is the only part of the paper they will ever see. A good title is specific enough to be found by the right search terms and honest enough not to overclaim – “Evidence for” and “Observation of” are not interchangeable, and a referee or a later reader will notice the difference – while staying short enough to read as a single clause rather than a run-on summary of the abstract. Some venues (several biology and biomedical journals, in particular) favour declarative, “result-stating” titles; most physics journals favour a more neutral noun-phrase title naming the system and effect studied, with the specific claim left to the abstract.

Author list and affiliations. Authorship order carries field-specific conventions that are easy to misread across disciplines: in much of biomedicine and chemistry, the first author is normally the person who did the bulk of the hands-on work and the last author is the

senior group leader or principal investigator, with a specific asterisk or footnote sometimes marking two people as “equal first authors”; in large physics collaborations (particle physics detector papers, for instance), author lists are often simply alphabetical across hundreds or thousands of names, with contribution recorded separately rather than encoded in the ordering at all; in a typical small physics group’s paper, order is frequently closer to “rough order of contribution,” with the supervising PI placed last as a matter of local custom rather than any universal rule. Each author is linked to one or more numbered affiliations (institution, department, city, country), and increasingly to a personal, persistent **ORCID iD** – a free identifier (orcid.org) that disambiguates a researcher from anyone else who happens to share their name, and that many journals now require at submission for every author, not only the corresponding one. Exactly who qualifies to be listed as an author at all is not merely a matter of local custom: the practice recommended by the International Committee of Medical Journal Editors (ICMJE), and echoed by most journals across disciplines, ties authorship to four joint conditions – a substantial contribution to the work’s conception or data, drafting or critically revising the manuscript, final approval of the submitted version, and agreement to be accountable for the work’s accuracy and integrity – a definition worth re-reading in light of the honorary-authorship problem Chapter 8 (§ ??) discussed.

The corresponding author. Exactly one author (occasionally two, for large or multi-institution collaborations) is designated, usually with an envelope symbol or asterisk next to their name, as the point of contact for the editor, the referees (indirectly, via the editor), and any reader who writes in after publication. In practice this person handles the submission portal, receives and coordinates the response to referee reports, and – critically, as § 9.7 below describes – receives and returns the proofs under a tight deadline; a corresponding author is not necessarily the person who did the most work, but is the one formally responsible for ensuring every co-author has seen, approved, and disclosed any conflicts of interest relevant to the final submitted version.

The abstract is a self-contained summary, written to be understood without reading the rest of the paper, and it is what a database such as Web of Science, Scopus, or Google Scholar actually indexes and displays in a search result – which makes its wording, not just its content, a deliberate choice about how the paper will be found. Conventions differ sharply by venue: many physics letters (Physical Review Letters among them) require a single unstructured paragraph, often startlingly short, with no citations, equations, or undefined abbreviations; many biomedical and some general-audience journals instead require a *structured* abstract with explicit sub-headings (Background, Methods, Results, Conclusions); Nature-family journals ask for an unstructured paragraph of roughly 150–200 words that states the question, the key result, and its significance in that order. Whatever the format, an abstract is normally written last, once the paper’s actual findings and framing are settled, rather than drafted first and left unrevised.

9.2 The body: introduction, methods, results, discussion

Introduction. The conventional shape is a funnel: open with the broad scientific question and why it matters, narrow to what is already known and, specifically, what gap or open disagreement remains, and close with an explicit statement of what this particular paper contributes – a sentence often signalled by a stock phrase such as “Here we show” or “In this Letter we report.” Citing prior work fairly and specifically (not merely a block citation to “previous studies [1–12]” with no indication of what each one actually found) is both a matter of scholarly honesty and, practically, the first thing an expert referee checks, since it signals whether the authors have genuinely engaged with the existing literature or are unaware of – or worse, deliberately omitting – a directly relevant prior result.

Methods (sometimes titled Experimental Methods, Materials and Methods, or simply Methods) exists to let another competent researcher reproduce the work from the description alone – the same reproducibility standard that Chapter 8’s discussion of research integrity, and the referee exercise earlier in this course’s evaluation, both turned on. It is normally written in the past tense, describing what was actually done rather than what is generally done, and should state not only what was measured but how many times, under what conditions, and what was varied versus held fixed or left uncontrolled. Several high-profile journals now split this section in two: a short version in the main text and a longer, more technical version relegated to an online-only Methods section or Extended Data, precisely so that the main narrative stays within the journal’s length limit (§ 9.5) without sacrificing the detail a specialist reader needs to replicate the result. Any human-subjects, animal-welfare, or biosafety approval is normally stated here (or in a dedicated Ethics statement), and increasingly so is a statement of what software, code, or analysis pipeline was used and where it can be found.

Results presents the findings themselves, generally through figures and tables that are meant to carry much of the argument on their own, each with a self-contained caption a reader could understand without reading the surrounding text. Reporting uncertainty honestly – error bars, confidence intervals, sample sizes, and a stated basis for any claim of statistical significance – is not an optional refinement but the standard this course’s earlier referee exercise (paper_1 and paper_2) specifically asked students to check for: a results section that shows a clear trend without ever quantifying how confidently that trend can be distinguished from chance is exactly the kind of flaw a competent referee, and a competent reader, is trained to notice.

Discussion / Conclusion. This is where results are interpreted against the question posed in the introduction and against the existing literature, and where limitations – small sample size, an uncontrolled variable, instrument precision, an assumption the analysis depends on – are stated explicitly rather than left for a reader to infer or, worse, omitted entirely. Some journals keep Discussion and Conclusion as two distinct sections (the former interpreting results, the latter summarising the paper’s contribution and pointing to future work); shorter formats, including PRL, typically merge the two into a single closing section or even a final paragraph. A conclusion that claims more certainty than the results actually support is one of the more common – and most consequential – weaknesses a referee is asked to catch.

9.3 Back matter: acknowledgments, contributions, references, and supplementary material

Acknowledgments record funding sources (usually with the specific grant number, since funders track this to demonstrate their own impact), technical or intellectual help that falls short of the authorship threshold discussed in § 9.1 (a colleague who read a draft, a technician who ran a specific measurement, a useful conversation at a conference), and, where relevant, a facility-access statement – large research infrastructures such as synchrotrons, neutron sources, or supercomputing centres typically require, as a condition of beam-time or compute allocation, that any resulting publication cite the facility and the specific proposal or allocation number.

Author contributions. Building directly on the CRediT taxonomy introduced in Chapter 8 (§ ??), most journals now either require or strongly encourage a short, explicit statement of who did what, typically in a form such as: “A.B. and C.D. designed the study; A.B. performed the experiment; C.D. and A.B. analysed the data; A.B. wrote the manuscript with input from all authors.” This is precisely the statement students are asked to write honestly for their own article-writing assignment (Section 5 of *Article_to_write*) later in this course’s evaluation – the same structural safeguard against honorary authorship, seen this time from the writer’s side of the page rather than the referee’s.

References. Citation style is journal-specific and non-negotiable at submission: numbered, bracketed citations in order of appearance (the convention used throughout the American Physical Society and IOP Publishing journals), superscript numbers (common at Nature-family journals), or author–year “Harvard” style (common across much of biology, chemistry, and the social sciences). Reference-management software – BibTeX or its more modern successor biblatex within L^AT_EX, or standalone tools such as Zotero, EndNote, or Mendeley – exists specifically to make switching between these styles a one-line change rather than a manual rewrite, and to attach a persistent Digital Object Identifier (DOI) to each entry wherever one exists. Accuracy here is not a minor formality: this course’s referee exercise built several of its planted errors (falsified volume numbers, a reference attributed to the wrong journal, a citation to an entry that does not exist in the bibliography) specifically around bibliography mistakes, because they are simultaneously common, easy to introduce by copy-paste, and straightforward for a careful reader to catch by checking the source.

Supplementary material (also called Supporting Information or, at some journals, Extended Data) holds content that supports the paper’s claims without belonging in the tightly length-limited main text: additional figures, derivations, control experiments, or an extended methods description. It is worth distinguishing supplementary material – typically a PDF bundled with the article itself – from a genuinely deposited **dataset or code repository**, carrying its own persistent identifier through a service such as Zenodo, the discipline-specific repositories, or France’s HAL, exactly as Chapter 8’s discussion of FAIR data and open science (§ ??) described; the data-availability statement required at submission (§ 9.6 below) is usually where this distinction has to be made explicit.

9.4 Writing in L^AT_EX: Overleaf and journal templates

Physics, mathematics, computer science, and most of the physical sciences write their papers in L^AT_EX, a typesetting system built on Donald Knuth’s original T_EX engine, rather than in a word processor. The appeal is structural rather than cosmetic: equations, cross-references, citations, and figure numbering are all handled programmatically – a citation key or an equation label is written once and resolved automatically everywhere it is used – which keeps a long, heavily mathematical, frequently revised document internally consistent in a way that manually numbering equations and citations in a word processor does not. Biomedical and many social-science fields, by contrast, still submit predominantly in Word, in large part because their journals’ own production workflows and reviewer expectations were built around it; a physics student collaborating across fields should expect to encounter both.

Overleaf (overleaf.com) is the platform that has, over roughly the past decade, made L^AT_EX genuinely collaborative: it is a browser-based editor that compiles a shared project in real time, tracks a full revision history, supports inline comments and track-changes-style review, and removes the older, clumsier workflow of emailing `.tex` files back and forth between co-authors and manually merging conflicting edits. Many universities, including some in France, provide their students and staff with an upgraded institutional Overleaf account at no cost – worth checking your own institution’s library or IT pages for before paying for a personal subscription. For the assignment associated with this chapter (*Article_to_write*), a shared Overleaf project is exactly the collaborative workspace each pair of students will use.

Journal templates. Almost every major publisher maintains a L^AT_EX class file, freely available (often directly from Overleaf’s own template gallery) and pre-configured with the journal’s required fonts, margins, section headings, and reference style, so that an author writes content rather than reformatting it by hand. The most relevant families for this course’s students are:

- **revtex4-2** – the American Physical Society’s own class, required for *Physical Review Letters* and the entire *Physical Review* family (A, B, C, D, E, X, Applied, Research, Fluids);
- **iopart** – IOP Publishing’s class, used for the *Journal of Physics* series and similar titles (*EPL*, formerly *Europhysics Letters*, uses its own related **ep12** class);
- **elsarticle** – Elsevier’s general-purpose class, spanning a large number of that publisher’s journals across disciplines;
- **aastex** – the American Astronomical Society’s class, used for *The Astrophysical Journal* and *The Astronomical Journal*;
- **IEEEtran** – the standard class for IEEE journals and conference proceedings, common in applied and engineering physics.

Nature-family journals are a partial exception: initial submission is deliberately flexible about formatting (Nature’s own author guidance explicitly states that format and style are not strictly enforced at first submission), and the great majority of Nature-family manuscripts

arrive as Word documents; a L^AT_EX-written manuscript is nonetheless accepted at initial submission at essentially all of them, with the journal’s production team handling the final formatting conversion after acceptance rather than requiring authors to match the house style themselves in advance. The general rule worth remembering, across almost every venue: at initial submission, editors and referees are reading for the science, not the typesetting, and a clean, readable, correctly structured manuscript in a generic `article`-class template is normally accepted for review even when it does not yet match the target journal’s final template exactly.

Historiographical note

L^AT_EX support, required templates, and even which submission formats a journal accepts change from time to time as publishers update their production systems; the authoritative source is always the specific journal’s current “Instructions for Authors” page (usually linked directly from the journal’s Overleaf template, when one exists), not a fixed list memorised once and never rechecked.

9.5 One science, many molds: comparing journal formats

The same underlying result can be written up as a four-page PRL, a several-thousand-word Nature Article, or an unrestricted-length *Physical Review B* regular article, and the differences are not merely cosmetic: a severe length limit forces a different, more selective kind of writing, pushing full technical detail out of the main text and into Methods, Extended Data, End Matter, or a companion longer paper, while an unrestricted-length venue allows – and to some extent expects – the full derivation, the complete set of control measurements, and a more exhaustive literature comparison to stay in the main text itself. Table 9.1 compares a handful of illustrative cases a physics student is likely to encounter.

Journal / type	Main-text length	Abstract	Figures/tables	References
<i>Phys. Rev. Lett.</i> (Letter)	≈4 journal pages (~3000–3750 word-equivalent, incl. figures)	≤600 characters, one paragraph	no fixed cap (typically 3–4)	no fixed cap
<i>Phys. Rev. B/D</i> etc. (regular article)	no fixed cap	unstructured paragraph	no fixed cap	no fixed cap
<i>Nature</i> (Article)	≈2500–4300 words	≈150–200 words	typically 4–6	up to ≈50
<i>Science</i> (Report)	≈2500 words	≤125 words	≈4	≈30
<i>Science</i> (Research Article)	≈4500 words	≤125 words	≈6	≈40–70

Table 9.1: Illustrative, approximate comparison of length and format constraints across a handful of high-profile and specialist physics venues. These numbers are drawn from each journal’s own current author guidelines and change periodically – see the histnote below.

Historiographical note
Every number in Table 9.1 is illustrative rather than a value to memorise: journals revise their length formulas, abstract limits, and reference guidance periodically (<i>Physical Review Letters</i> , for instance, redefined how it counts length in 2011 and has adjusted the rule since), and some limits are enforced as hard desk-rejection criteria while others are closer to a strong convention an editor can override. Before submitting to any journal, an author checks that journal’s own, currently published Instructions for Authors rather than a fixed table from a course chapter.

The most selective, broadest-audience venues (PRL, *Nature*, *Science*) impose their tight caps deliberately, less to save printed pages – most of this content is now read online – than to force a compressed, broadly legible framing aimed at readers outside the immediate sub-field, with technical completeness deferred to Methods, Supplementary Material, or a companion paper. A common and entirely legitimate practice in physics is exactly that pairing: a short, high-visibility Letter announcing a result, published alongside or shortly after a longer companion article in a specialist journal (a *Physical Review B* or *D* regular article, say) that gives the full derivation, the complete dataset, and every control measurement the Letter’s length limit could not accommodate. This differs in an important way from the “salami-slicing” practice Chapter 8 (§ ??) flagged as a questionable research practice: a Letter–companion-paper pair reports the same result once, at two levels of detail, each paper stating clearly that the other exists, rather than dividing one dataset into the minimum number of separately publishable units to inflate a publication count.

9.6 Submitting a manuscript: cover letters, systems, and disclosures

A **cover letter** is not an optional courtesy: almost every journal’s online submission system will not let a manuscript file through without one, and the handling editor typically reads it before opening the manuscript itself, so it is often the very first thing that shapes an editor’s first impression of the paper. Addressed directly to the handling editor rather than written as a generic form letter, it normally covers:

- the manuscript’s title and a one- or two-sentence statement of what was done and why it matters;
- an explicit, journal-specific case for fit and significance – why *this* journal and *this* readership, not a paragraph reused unchanged across five different submissions;
- confirmation that the manuscript is original work, is not simultaneously under review elsewhere, and that all authors have seen and approved the submitted version;
- where the journal allows it, a short list of suggested referees (with a brief reason each is qualified) and, sometimes, a short list of researchers to explicitly exclude, typically for

a stated conflict of interest.

A cover letter that merely restates the abstract, without making a specific case for the journal's readership, is a common and easily avoidable weakness – editors read dozens of these a week and notice the difference immediately. Most journals in physics have, since roughly the mid-2010s, explicitly welcomed manuscripts that have already been posted as a preprint (on arXiv, in physics' case), reversing an older and stricter norm, sometimes informally called the Ingelfinger rule after the *New England Journal of Medicine* editor who first articulated it, under which prior public disclosure of a result could disqualify it from consideration; *Nature* and *Science* both now explicitly permit preprinted submissions as well, though a specific journal's current preprint policy is always worth checking rather than assumed.

The manuscript itself is uploaded through an online **submission system**, several of which a student is likely to encounter across different journals: **Editorial Manager** (built by Aries Systems and used across much of Elsevier's, Springer's, and many independent societies' journal portfolios), **ScholarOne Manuscripts** (also marketed as Manuscript Central, built by Clarivate and used widely by Wiley and IOP Publishing among others), the American Physical Society's own portal at authors.aps.org, and Nature-family journals' own manuscript-tracking system. Whichever system is used, the submission form itself typically requires, beyond the manuscript file:

- the full author list, affiliations, and – increasingly mandatory rather than optional – each author's ORCID iD;
- subject classification or keywords (once standardised across physics through the Physics and Astronomy Classification Scheme, PACS, retired by the American Physical Society during the 2010s in favour of free-text terms and arXiv's own subject categories);
- a **data-availability statement**, specifying where and how the data underlying the results can be accessed – directly connected to the FAIR-data expectations Chapter 8 (§ ??) described, and, for most funders operating under an open-science mandate (the European **cOAlition S / Plan S** initiative among them, directly relevant to ANR- and Horizon-Europe-funded French researchers), no longer optional;
- a **conflict-of-interest (or competing-interests) disclosure**, covering financial relationships (employment, consulting, patents, equity) and relevant non-financial relationships for every author, and, separately, used by the handling editor to screen potential referees for the same kinds of conflict;
- an **open-access and licensing choice**, since most journals today offer both a traditional subscription route and a “gold” open-access route funded by an article-processing charge (APC) under a Creative Commons licence (typically CC-BY), a choice increasingly constrained by funder mandates rather than left entirely to the author; and
- , where relevant, a statement of **ethics approval** (an institutional review board or ethics-committee reference number for human-subjects or animal research) – a near-universal requirement outside physics, and an occasional one within it, for instance in biophysics or medical-physics work.

9.7 After submission: peer review, decisions, and proofs

A submitted manuscript first passes an **editorial screening**: an editor checks it against basic scope and quality criteria and, at journals with a hard length limit, against that limit itself – a manuscript that exceeds *Physical Review Letters*’ page limit, for instance, is normally returned to the authors without ever being sent to a referee. A manuscript that clears this stage is assigned to a handling or associate editor, who selects typically two or three independent referees (drawing on the field expertise, and avoiding the conflicts of interest, discussed in § 9.6); their written reports, together with the editor’s own reading, lead to one of a small number of standard decisions – accept, minor revision, major revision, or reject, sometimes with encouragement to resubmit a substantially reworked version as a new submission. A revision is normally returned together with a point-by-point rebuttal letter, addressing every referee comment individually and stating exactly what, if anything, changed in the manuscript in response – itself a piece of technical writing with its own conventions, and often the single most consequential document an author writes during the entire review process.

Once a paper is finally accepted, it moves to **production**: a copy-editing team formats the accepted manuscript into the journal’s house style and typesets it, generating a **proof** – historically called a galley proof – that the corresponding author (and, at many journals, every co-author) is asked to check within a short, firmly enforced window, typically somewhere between 24 and 72 hours. Proof-stage corrections are meant to be narrow: fixing a typesetting error, a missing funding or grant number, a misspelled name, or a genuine factual slip, not rewriting sentences, restructuring the argument, or adding new data or analysis, which most journals either refuse outright at this stage or route back through editorial approval, sometimes at the author’s expense. Once proofs are returned and any last corrections applied, the paper is typically published “**online first**” with its own DOI well before it is assigned to a specific printed volume and issue, becoming, at that point, the citable **Version of Record**.

9.8 Bibliometrics in practice: measuring a researcher, measuring a journal

Once a paper reaches that Version of Record, both it and its authors start accumulating the citation-based numbers Chapter 8 (§ ??) introduced in principle – the journal impact factor and the h-index, and the DORA pushback against over-relying on either. This section works through those numbers as they actually appear on screen: the databases that compute them, the fuller family of metrics beyond the h-index and the impact factor, and current, concrete figures for a handful of journals this course has already discussed.

Three databases, three different counts. The same researcher’s citation count and h-index can differ substantially depending on where it is measured, because the three major sources differ in what they index. **Google Scholar** is free, requires no subscription, and casts the widest net – indexing preprints, theses, conference proceedings, books, and other grey literature alongside peer-reviewed articles; a researcher’s public Google Scholar profile computes and displays an h-index and an i10-index automatically, updating as new citations

appear. **Web of Science** (Clarivate), through its curated “Core Collection,” indexes a more selective set of peer-reviewed venues and is the authoritative source of the Journal Impact Factor via its annual Journal Citation Reports (JCR); it is subscription-based, and a researcher’s citation record there is tied to a personal identifier (ResearcherID, now integrated with ORCID). **Scopus** (Elsevier) sits between the two in coverage – broader than Web of Science, more selective than Google Scholar – and is the source of the CiteScore, SJR, and SNIP journal metrics described below, feeding into Elsevier’s SciVal analytics platform; it too is subscription-based. Because grey literature and conference proceedings inflate a count that a curated index excludes, Google Scholar figures for the same researcher are typically the highest of the three, Web of Science typically the most conservative, and Scopus in between.

Database	Access	Coverage philosophy	What it computes
Google Scholar	Free	Broadest: journals, preprints, theses, books, conference proceedings	Author profile: h-index, i10-index (auto-updated); Google Scholar Metrics: journal h5-index
Web of Science	Subscription (Clarivate)	Curated “Core Collection” of selective, indexed venues	Journal Impact Factor (JCR, annual); author h-index via ResearcherID
Scopus	Subscription (Elsevier)	Broader than WoS, narrower than Scholar	CiteScore, SJR, SNIP; author h-index; Field-Weighted Citation Impact (via SciVal)

Table 9.2: The three bibliometric sources a researcher is most likely to encounter, and what each one actually reports.

A researcher’s profile, unpacked. Beyond the raw **number of publications** and **total citation count** – both sensitive to career length, field size, and co-authorship practices (a large collaboration paper with hundreds of co-authors credits every one of them with its full citation count) – several further metrics refine the h-index Chapter 8 introduced. The **i10-index**, specific to Google Scholar and introduced in 2011, is simply the number of a researcher’s papers that have received at least ten citations each: cruder than the h-index but easy to read at a glance. The **g-index**, a less widely used refinement, gives more weight to a researcher’s most highly cited papers than the h-index does, so that a handful of exceptional results can move it even when the rest of the record is modest. The **m-quotient** (h-index divided by the number of years since a researcher’s first publication) is a deliberate attempt to make the h-index comparable across career stages, since raw h-index otherwise mechanically favours whoever has simply been publishing longest. Scopus’s **Field-Weighted Citation Impact (FWCI)**, computed via SciVal, normalises a researcher’s (or a single paper’s) citation count against the average for papers of the same field, publication year, and document type: a

value of exactly 1.0 means precisely average impact for that context, and values above or below 1 read directly as a multiple of the field average – a way of comparing, say, a mathematician and a biologist without the raw citation counts being misleading in either direction.

Historiographical note

Any concrete h-index or citation-count range offered as an illustration – a beginning PhD student typically has a handful of papers and a single-digit h-index; an established senior researcher in an active subfield is often, but by no means always, somewhere in the range of roughly 20 to 40; the most highly cited figures in a field can exceed 60 or 100 – is an order-of-magnitude illustration, not a target, and depends enormously on field size and citation culture, career stage, and co-authorship norms (particle-physics collaboration papers, in particular, mechanically inflate every listed author’s count). Chapter 8’s account of DORA (§ ??) applies with full force here: these numbers describe a record, not a verdict on the quality of the work behind it.

A journal’s numbers, unpacked. The **Journal Impact Factor (JIF)**, published annually by Clarivate’s Journal Citation Reports (usually released around the middle of the year), divides the citations a journal receives in a given year by the number of citable items it published in the two preceding years – a two-year citation window later supplemented, for journals where this is judged more informative, by a five-year variant. **CiteScore** is Scopus’s own alternative: citations in a given year to a journal’s documents published across a several-year window (currently four years, including the current year), divided by the number of such documents, counting a broader range of document types than the JIF and recalculated monthly rather than once a year. **SJR (SCImago Journal Rank)**, freely available at scimagojr.com, weights each citation by the citing journal’s own prestige – a citation from a highly cited journal counts for more than one from an obscure title, an idea related to Google’s original PageRank algorithm. **SNIP (Source Normalised Impact per Paper)** normalises for field-specific citation practices, making it one of the few metrics that supports a reasonably fair comparison between, say, a mathematics journal and a biomedical one, whose citation cultures otherwise differ enormously. Finally, **Google Scholar Metrics** publishes its own free, journal-level **h5-index** and **h5-median** – the h-index and its associated median, computed only over each journal’s articles from the last five complete calendar years.

Journal	JIF	Note
<i>Nature</i>	56.1 (2025)	up from 50.5 just three years earlier (2022) – see the histnote below
<i>Science</i>	45.8 (2024)	
<i>Physical Review Letters</i>	9.4 (2025)	
<i>Physical Review B</i>	3.7 (2024)	typical of a large, specialised, high-volume physics journal

Table 9.3: Illustrative current Journal Impact Factors for four journals already discussed in this chapter, each as reported in its own most recently published Journal Citation Reports year.

Historiographical note

Table 9.3's figures are a snapshot, not a constant: Clarivate republishes the JCR annually, and a journal's JIF can move substantially from one year to the next – *Nature*'s own figure above rose by more than ten percent in three years – for reasons that include genuine shifts in citation behaviour, methodology changes at Clarivate, and a journal's own editorial choices about what counts as a “citable item.”

The DORA pushback. Every metric this section has described – the JIF and CiteScore above all, but the h-index and its variants just as much – is exactly what the **San Francisco Declaration on Research Assessment (DORA)** was written to push back against. Drafted at the December 2012 annual meeting of the American Society for Cell Biology in San Francisco, DORA's central recommendation is blunt: funders, institutions, and journals should 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 should instead assess the scientific content of a candidate's actual work directly. Its more specific recommendations bear directly on the numbers this section has just walked through: journals should stop actively marketing their JIF, and any institution that does use a quantitative metric at all should report it alongside a range of complementary indicators – an early, 2012 statement of the same multi-database, multi-metric caution this section has taken toward Google Scholar, Web of Science, and Scopus individually. Thousands of institutions and individual researchers worldwide have signed the declaration since 2012, including, in France, the CNRS in 2018; how consistently a given hiring or funding committee's actual practice matches that commitment is, as Chapter 8 (§ ??) already noted, an uneven and still-debated question rather than a settled one. None of the metrics in this section – for a journal or for a researcher – should be read as a direct measure of the quality of any single paper; they are citation-based proxies, useful for some purposes and, as the h-index gaming and citation-cartel practices mentioned in Chapter 8 illustrate, not immune to being gamed.

Looking ahead

This chapter has deliberately mirrored, from the writer's side, the referee exercise and the evaluation criteria used earlier in this course: the same categories a good referee is trained to check for – a properly sourced introduction, a reproducible methods section, honestly quantified uncertainty in the results, an appropriately modest discussion, a truthful author-contributions statement, and an accurate bibliography – are exactly the categories this chapter has just walked through as things to *produce* rather than to detect. From the blank Overleaf project this chapter opened with, through the anatomy of the article, the tools used to typeset it, the very different molds it might fit, and the submission and review process, to the citation-based numbers a published paper eventually accumulates, this chapter has traced the complete path a manuscript follows into the scientific record. Together with Chapter 8's account of the funding, integrity, evaluation, and openness norms that surround it, it completes this course's account of the practical machinery of contemporary research: not only how discoveries happened historically, the subject of the course's earlier chapters, but how a researcher today

actually turns a result into a public, citable, peer-reviewed contribution to that same long record.

Further reading

- APS Journals, *Length Limits and Guidelines for Physical Review Article Types*, journals.aps.org/authors/length-guide.
- *Physical Review Letters, Information for Contributors and Style Basics*, journals.aps.org/prl.
- Springer Nature, *A Brief Guide to Initial and Revised Submissions and Initial Submission* author guidance, nature.com/nature/for-authors.
- *Science/AAAS, Information for Authors*, science.org.
- International Committee of Medical Journal Editors (ICMJE), *Defining the Role of Authors and Contributors*, icmje.org.
- CRediT (Contributor Roles Taxonomy), credit.niso.org.
- ORCID, orcid.org.
- cOAlition S, *Plan S: Making Full and Immediate Open Access a Reality*, coalition-s.org.
- Committee on Publication Ethics (COPE), *Core Practices*, publicationethics.org.
- Overleaf template gallery, overleaf.com/latex/templates.
- Clarivate, *Journal Citation Reports* (annual), jcr.clarivate.com.
- Google Scholar, *Google Scholar Metrics Help* (h5-index, h5-median), scholar.google.com.
- SCImago Research Group, *SJR – SCImago Journal & Country Rank*, scimagojr.com.
- Elsevier, *CiteScore Metrics*, scopus.com/sources.