

National University of Kyiv-Mohyla Academy
English Language Department

English for Professional Development

ACADEMIC WRITING IN ECONOMICS

A Guide for MA Students, Post-graduates and Early-Career
Researchers

Prepared by:

Alla Shuhai

Oksana Kharchenko

Kyiv

2026

SUMMARY

This guide addresses a common problem in academic economics: the gap between strong analytical work and clear academic writing. Many papers are rejected not because of weak research, but due to avoidable writing issues. The guide offers practical tools to identify and improve such problems.

It is organised into eight chapters. The first chapter presents the basic structure of economic argument: claim, evidence, analysis, and implication. The following chapters cover key academic genres, including research papers, literature reviews, and case studies, treating them as forms of argument. Other chapters focus on writing about quantitative data, citation and academic integrity (including AI use), and the writing process, such as planning, revising, and using feedback. A separate chapter is devoted to writing effective abstracts.

The guide uses examples from leading economics papers to demonstrate how writing choices support clear argumentation. It is designed for flexible use, allowing readers to focus on specific sections or study it as a whole.

The central idea of the guide is that clear writing and clear thinking are closely connected and should be developed together.

CONTENTS

Introduction: Academic Writing as Disciplinary Practice

- Why Writing Remains a Challenge for Experienced Economists
- What Distinguishes Economics Writing from Other Disciplines
- The Publication Landscape in Economics
- The Structure of This Guide

Chapter One: The Architecture of Economic Argument

- 1.1 Argument vs. Description: The Foundational Distinction
- 1.2 The Four-Move Argument Structure
- 1.3 The Role of Assumptions
- 1.4 Hedging and Epistemic Responsibility
- 1.5 Paragraph Structure as Argument Structure
- 1.6 Sentence-Level Precision
- 1.7 Vocabulary and Sentence Models

Chapter Two: The Research Paper for Publication

- 2.1 Understanding What Journals Expect
- 2.2 Developing and Articulating the Research Question
- 2.3 The Introduction
- 2.4 The Literature Review within a Research Paper
- 2.5 The Methodology Section
- 2.6 The Results Section
- 2.7 The Discussion Section
- 2.8 The Conclusion
- 2.9 The Acknowledgements Section
- 2.10 Responding to Referee Reports
- 2.11 Vocabulary and Sentence Models

Chapter Three: The Literature Review as a Standalone Publication

- 3.1 The Model Text: Goldfarb and Tucker (2019)
- 3.2 The Literature Review as a Scholarly Contribution
- 3.3 Scoping and Selecting the Literature
- 3.4 Synthesising Rather Than Summarising
- 3.5 Organising the Review
- 3.6 Writing the Review: Attribution, Voice, and Hedging
- 3.7 Positioning the Paper's Contribution
- 3.8 Diagnosing Problems in Literature Reviews
- 3.9 Vocabulary and Sentence Models

Chapter Four: The Case Study as an Academic Publication

- 4.1 The Model Text: Farronato and Fradkin (2022)
- 4.2 What Makes a Case Study Publishable
- 4.3 The Architecture of the Case Study Paper
- 4.4 Constructing the Theoretical Framework
- 4.5 Data, Measurement, and the Credibility of Results

- 4.6 Presenting Results: Prose, Tables, and Figures
- 4.7 The Generalisation Problem
- 4.8 Limitations: Proportionate, Not Performative
- 4.9 Writing from a Ukrainian and Regional Context
- 4.10 Vocabulary and Sentence Models

Chapter Five: Data Interpretation and Quantitative Writing

- 5.1 The Model Text: Chetty, Hendren, Kline, and Saez (2014)
- 5.2 Measurement Decisions and Their Justification
- 5.3 Communicating Magnitudes
- 5.4 The Correlation–Causation Discipline
- 5.4b Contrast: Reinhart and Rogoff (2010)
- 5.5 Tables and Figures
- 5.6 Writing the Empirical Framework Section
- 5.7 Robustness, Heterogeneity, and the Limits of the Analysis
- 5.8 Connecting Quantitative Findings to Policy
- 5.9 Vocabulary and Sentence Models

Chapter Six: Citation, Attribution, and Academic Integrity

- 6.1 APA 7th Edition: In-Text Citations
- 6.2 APA 7th Edition: The Reference List
- 6.3 When and What to Cite
- 6.4 Quoting, Paraphrasing, and Summarising
- 6.5 Academic Integrity: The Spectrum of Failures
- 6.6 Open Science and Reproducibility
- 6.7 Vocabulary and Sentence Models

Chapter Seven: The Writing Process

- 7.1 Planning and Sequencing the Draft
- 7.2 The First Draft: Speed Over Perfection
- 7.3 Revision: Argument, Structure, and Sentence Level
- 7.4 Getting and Using Feedback
- 7.5 Writing in English as an Additional Language
- 7.6 Vocabulary and Sentence Models

Chapter Eight: Writing the Abstract

- 8.1 What the Abstract Must Do
- 8.2 The Model Abstract: Friedman (1968) Annotated
- 8.3 Structure and Word Count
- 8.4 Abstracts Across Genres
- 8.5 Common Failures in Economics Abstracts
- 8.6 Revision: From Draft to Final Abstract
- 8.7 Practical Considerations: Submission, Search, and Reuse
- 8.8 Vocabulary and Sentence Models

Conclusion: Writing and Thinking as a Single Practice

Appendix A: APA 7th Edition Quick-Reference Card

Appendix B: Pre-Submission Checklist

References

INTRODUCTION

Academic Writing as Disciplinary Practice

◆ Why Writing Remains a Challenge for Experienced Economists

There is a persistent and somewhat uncomfortable truth in academic economics: the ability to construct a rigorous model, interpret a dataset with precision, or identify a genuine gap in the empirical literature does not, by itself, produce clear and persuasive writing. The two capacities are related but distinct, and the distinction matters enormously at the point of publication. A paper whose analytical contribution is genuine but whose written exposition is unclear, poorly structured, or imprecisely argued will struggle to reach its audience — and in a competitive publication environment, it may never do so at all.

This is not a minor or peripheral concern. Desk rejection rates at leading economics journals routinely exceed sixty or seventy per cent, and while many rejections reflect genuine limitations in the research itself, a significant proportion reflect failures of presentation: introductions that do not establish a clear contribution, methodology sections that describe without justifying, discussion sections that present findings without interpreting them, and conclusions that summarise rather than conclude. These are writing problems, and they are correctable — but only if they are recognised as such.

The challenge is compounded by a disciplinary culture that has historically treated writing as secondary to analysis. In economics, the implicit assumption has long been that if the model is sound and the data are clean, the writing will take care of itself. Neugeboren and Jacobson (2005) were among the first to challenge this assumption directly, arguing that writing in economics is not merely the transcription of completed thinking but is itself a mode of thinking — that the act of putting an argument into precise written form tests, refines, and sometimes fundamentally alters the argument. This guide takes that position as its starting point and extends it to the specific demands of academic publication and postgraduate research development.

◆ What Distinguishes Economics Writing from Other Disciplines

Academic writing across all disciplines shares certain foundational commitments: clarity of argument, honesty about evidence, precision of language, and appropriate engagement with existing scholarship. But each discipline also develops its own conventions, its own tolerance for particular kinds of ambiguity, and its own hierarchy of values — and economics is no exception.

The most fundamental of these disciplinary values is precision. Economics writing prizes the exact statement over the approximate one, the qualified claim over the sweeping generalisation, and the parsimonious explanation over the elaborate one. Where a historian might develop an argument through narrative accumulation, and a sociologist through theoretical elaboration, an economist is expected to build an argument in the most direct line

possible from question to evidence to conclusion, stripping away anything that does not directly serve that progression. This is not merely a stylistic preference: it reflects the discipline's deeper epistemological commitments to falsifiability, replicability, and the careful separation of what the data show from what the analyst believes.

A second distinguishing feature is the relationship between formal modelling and written exposition. In many areas of economics, the mathematical or econometric component of a paper carries the analytical weight, and the written text is sometimes treated as little more than a wrapper around the formal content. This misconception produces some of the weakest writing in the discipline. The written sections of an economics paper — the introduction, the literature review, the discussion, the conclusion — are not decorative; they are the spaces in which the paper makes its case to a reader who may be unconvinced, sceptical, or simply unfamiliar with the specific literature. A formal model cannot argue for its own relevance. Only the writing can do that.

A third consideration is the audience. Economics writing is not addressed to a single homogeneous readership. A paper submitted to the *American Economic Review* is written for a different reader than a policy brief for the European Central Bank or a chapter in an edited volume on development economics. The register, the assumed knowledge, the density of technical language, and the appropriate level of hedging all shift according to audience, and the ability to calibrate writing for the intended reader is a skill that experienced economists, no less than novice ones, need to develop consciously.

◆ The Publication Landscape in Economics

Understanding the context in which academic economics writing is produced and evaluated is essential for anyone seeking to write more effectively within it. The publication landscape in economics has distinctive features that directly shape the writing decisions a researcher must make.

Journal publication remains the primary currency of academic reputation in economics, and the hierarchy of journals — from field-leading general interest journals such as the *American Economic Review*, the *Quarterly Journal of Economics*, and the *Review of Economic Studies*, through well-regarded field journals, to specialist and regional outlets — is both well-established and consequential for career progression. Each tier carries different expectations regarding the scope of the contribution, the rigour of the methodology, and the clarity of the writing. It is a common error to assume that writing standards are lower at less prestigious outlets: in practice, editors at journals of all tiers respond negatively to poorly written submissions, and a paper that cannot make its argument clearly is unlikely to survive the review process regardless of its analytical merits.

The review process itself deserves attention as a writing context. Referee reports are, among other things, writing feedback — often highly detailed, sometimes contradictory, and occasionally harsh in tone. Learning to read referee reports analytically, to distinguish fundamental methodological objections from matters of presentation, and to respond in a revision letter that is both comprehensive and persuasive is a distinct and important writing

skill. For many economists, the revision process is where the most significant improvements in a paper's written quality occur, yet it receives remarkably little systematic attention in discussions of academic writing development.

Working papers occupy an increasingly important role in the economics publication pipeline. Pre-publication circulation through repositories such as NBER, SSRN, and CEPR allows researchers to establish priority, receive informal feedback, and build a readership before formal publication. The writing demands of a working paper are in most respects identical to those of a published article, but the absence of editorial oversight means that the responsibility for writing quality rests entirely with the author. A poorly written working paper that circulates widely can shape a reader's impression of the underlying research in ways that are difficult to correct later.

Beyond journal articles and working papers, economists increasingly contribute to policy reports, institutional publications, book chapters, and edited volumes — each of which carries its own genre conventions and writing demands. This guide addresses the research paper and literature review as its central concerns, while also attending to case study reports and data-driven analysis writing, recognising that faculty members and postgraduate researchers are likely to work across several of these formats over the course of their academic careers.

◆ The Structure of This Guide

This guide is organised around the core writing tasks that economists face in academic publication and postgraduate research. It is intended primarily as a reference for faculty members seeking to strengthen their own writing practice and to support the writing development of colleagues and postgraduate students. It may also be used directly by MA and doctoral students in economics, particularly those preparing their first journal submissions or working on extended research writing projects.

Each chapter addresses a distinct genre or dimension of economics writing. Chapter One establishes the foundational principles of economic argument — the architecture of claim, evidence, analysis, and implication that underlies all the genres addressed in subsequent chapters. Chapters Two through Four address the principal publication genres: the research paper, the standalone literature review, and the case study. Chapter Five focuses on the specific challenge of writing about quantitative data and econometric analysis — an area in which technically strong economists frequently encounter difficulty. Chapter Six addresses citation, referencing conventions, and academic integrity, including the emerging challenges posed by AI writing tools. Chapter Seven addresses the writing process itself: planning, drafting, revision, feedback, and writing in English as an additional language.

Throughout the guide, consistent reference is made to Neugeboren and Jacobson's *Writing Economics* (2005), whose framework for understanding economics writing as a form of disciplined argumentation remains the most rigorous and practically useful treatment of the subject available. Where this guide extends or develops their framework — particularly in its attention to publication-level writing, referee response, and the specific demands of writing from a post-Soviet academic context — those extensions are clearly signalled.

The guide is designed to be used modularly. Faculty members with specific and immediate concerns — a research paper under revision, a literature review in progress, a postgraduate student struggling with data interpretation writing — will find it useful to go directly to the relevant chapter. Those approaching the guide as a more comprehensive resource for developing their writing practice are encouraged to begin with Chapter One, whose foundational principles inform everything that follows.

Clear writing and rigorous thinking are not separate achievements. They are the same achievement, approached from different directions.

CHAPTER ONE

The Architecture of Economic Argument

The chapters that follow cover specific genres of economics writing: the research paper, the literature review, the case study, and the quantitative report. Before those genres can be

discussed usefully, a more basic question must be answered: what does it mean to build an argument in economics, and how does that argument appear in writing at the level of the sentence, the paragraph, and the paper as a whole? This chapter answers that question. Everything in the chapters that follow depends on it.

1.1 Argument vs. Description: The Foundational Distinction

The most important distinction in economics writing — and the one that most reliably separates strong manuscripts from weak ones — is the distinction between argument and description. It is also the distinction that is most often ignored in submitted work, from MA dissertations to journal manuscripts.

A description tells the reader what the case is. An argument tells the reader why it is the case, what it means, and why it matters. Description is necessary: clear information about the data, the context, or the institutional setting is often necessary. However, it is not sufficient on its own. A paper that only describes an economic phenomenon, without making a claim about it, supporting that claim with evidence, and drawing out its implications, is not yet an academic argument. It is, at best, a structured account.

Neugeboren and Jacobson (2005, p. 7) put this directly: 'The goal of economics writing is not to inform but to persuade — to convince the reader that your analysis is correct, your interpretation is sound, and your conclusions follow from your evidence.' Some economists find this uncomfortable, because persuasion sounds like rhetoric, and rhetoric sounds like the opposite of empirical science. But the concern is misplaced. Persuasion in academic writing does not mean manipulation. It means building a case that a sceptical, informed reader will find convincing. Every decision a writer makes — what to emphasise, what to qualify, what to put in an appendix — is aimed at that goal.

For writers, the practical implication is this: at every stage of drafting, do not ask 'what do I know about this topic?' Ask instead, 'what am I claiming, and what supports that claim?' Moving from the first question to the second is the move from description to argument. It does not happen automatically. It requires a deliberate choice about what the paper is trying to do.

◆ The thesis statement in economics

Every economic argument starts with a thesis statement: a precise, contestable claim that tells the reader what the paper will argue and why it matters. A thesis statement in economics is not a topic announcement ('This paper examines the relationship between trade openness and income inequality') and it is not a description of method ('Using panel data from 47 countries, this paper analyses...'). It is a claim that can be disputed, supported with evidence, and connected to a debate in the literature.

Consider the difference between these two formulations:

✗ Topic announcement	✓ Arguable thesis
----------------------	-------------------

<i>This paper examines the effects of monetary policy on inflation expectations in low-income countries.</i>	<i>Conventional inflation-targeting frameworks systematically underperform in low-income countries because they fail to account for the structural role of food price volatility in shaping household inflation expectations.</i>
<i>This paper analyses the relationship between institutional quality and foreign direct investment.</i>	<i>The positive effect of institutional quality on FDI inflows is significantly moderated by sectoral composition: the relationship holds for manufacturing but largely disappears in extractive industries, where resource rents dominate location decisions.</i>

The right-hand formulations are argumentative, not just topical. They take a position — one that could in principle be wrong — and they signal the analytical logic the paper will develop. A reader who encounters either of these thesis statements knows not only what the paper is about, but what it claims. They can then judge whether the evidence and analysis that follow support that claim.

The thesis statement does not have to be a single sentence. In economics, it is often spread across two or three sentences at the end of the introduction. But wherever it appears, it must be present and specific. A vague thesis produces vague paper.

1.2 The Four-Move Argument Structure

Once the thesis is in place, the argument must be developed. The most reliable way to do this — at both the paragraph level and the paper level — is through what this guide calls the four-move sequence: claim, evidence, analysis, implication. This is not a formula to be followed mechanically. It is a description of the logical steps that any well-constructed economic argument must take. Learning to recognise when one of these steps is missing is one of the most useful skills a writer can develop.

1	Claim A precise, contestable statement of what you are asserting. The claim should be specific enough to be falsifiable and consequential enough to be worth making. It can appear at the opening of a paragraph (as a topic sentence) or at the opening of a section.
----------	---

2	Evidence The empirical, statistical, or theoretical support for the claim. In economics, evidence is typically quantitative, but it may also include theoretical results, historical documentation, or comparative institutional analysis. Evidence without analysis is inert — it requires interpretation to become argumentative.
3	Analysis The intellectual work of connecting the evidence to the claim: explaining why the evidence supports the assertion, what it rules out, and what it leaves uncertain. This is the move most frequently absent or underdeveloped in submitted manuscripts.
4	Implication The articulation of what follows from the claim, for the broader literature, for policy, or for the research agenda. The implication move is what transforms a local finding into a contribution. A paper that stops analysis — that demonstrates a point without drawing out its consequences — leaves its own work unfinished.

Milton Friedman's writing is a useful example of the four-move structure at the paragraph level. In 'The Role of Monetary Policy' (Friedman, 1968), one of the most cited papers in macroeconomic history, Friedman writes:

What if the monetary authority chose the 'natural' rate — either of interest or unemployment — as its target?

One problem is that it cannot know what the natural rate is. Unfortunately, we have as yet devised no method to estimate accurately and readily the natural rate of either interest or unemployment. And the 'natural' rate will itself change from time to time. But the more fundamental problem is that even if the monetary authority knew the 'natural' rate at the moment, it would not be sensible for it to try to maintain that rate. The natural rate changes, and the monetary authority would have to change its target continually. Moreover, the monetary authority, if it attempted to peg the nominal rate of interest, would be setting out on a path that has no ending — continually accelerating or decelerating inflation.

Friedman, M. (1968). The role of monetary policy. American Economic Review, 58(1), p. 8.

What makes this passage work: the four-move structure

Claim: The monetary authority cannot and should not attempt to peg either the natural rate of interest or unemployment.

Evidence: Two distinct problems are identified: an epistemological problem (the natural rate is unobservable) and a dynamic problem (the natural rate changes over time). Friedman is careful to distinguish between these — they are not the same objection.

Analysis: The critical analytical move comes in the phrase 'more fundamental problem'. Friedman does not merely list the two difficulties; he ranks them, making clear that the second is not just an additional problem but a deeper one. The phrase 'no ending' does significant analytical work: it signals that the policy path is not merely difficult but structurally unstable.

Implication: The implication is stated directly in the final sentence — the attempt to peg the nominal rate leads to 'continually accelerating or decelerating inflation'. This is not a marginal caveat; it is the mechanism of the later-named 'inflation spiral', and it follows with logical necessity from the analysis that precedes it.

Notice also what Friedman does not do: he does not hedge excessively, he does not bury the implication in a subordinate clause, and he does not leave the reader to draw the conclusion independently. The argument arrives where it must.

The four-move structure works not only within paragraphs but across whole papers. The introduction makes the paper's central claim. The literature review provides the theoretical and evidential context. The methodology and results sections develop the evidence. The discussion carries out the analysis. The conclusion draws out the implications. When a paper feels loose or unconvincing, the problem almost always traces back to one of these four moves being absent, too short, or out of place.

The implication move deserves special attention because it is the one most often omitted or underdeveloped in submitted manuscripts. It answers the reader's most important question: so what? A paper that demonstrates a finding but does not explain what follows from it — for the literature, for policy, for future research — has not completed its task. As one frequently cited senior referee puts it: 'I can see what you found. I cannot see why I should care.' The implication move is the answer to that challenge.

1.3 The Role of Assumptions

All economic arguments depend on assumptions — about how agents behave, how markets are structured, how data are generated, and where the boundaries of the model lie. This

is not a weakness of economics as a discipline. It is what makes economics produce testable, falsifiable claims. What matters for writing is not whether assumptions are made — they always are — but whether they are stated clearly, justified, and honestly limited in scope.

The most common writing failure related to assumptions is concealment: building an argument on assumptions that are never stated, which makes the analysis appear more general than it actually is. Referees are particularly alert to this. An argument whose assumptions are hidden cannot be properly evaluated. Its scope cannot be assessed, and its implications cannot be correctly calibrated.

Friedman's handling of assumptions in the 1968 paper is a good example of this discipline. His key assumption — adaptive expectations — is central to the whole argument, and he introduces it explicitly before the analysis begins. He writes:

Let us suppose that the monetary authority tries to peg the 'market' rate of unemployment at a level below the natural rate. For definiteness, suppose that it takes 3 per cent as the target rate and that the 'natural' rate is higher than 3 per cent. Suppose also that we start in a position of equilibrium at the natural rate with a stable price level.

Friedman, M. (1968). The role of monetary policy. American Economic Review, 58(1), p. 9.

What this passage demonstrates about assumptions

Friedman does not begin the analysis and reveal its assumptions later. He states them at the outset — specific numerical values ('3 per cent'), specific initial conditions ('equilibrium at the natural rate', 'stable price level'), and specific relational constraints ('below the natural rate'). The reader knows exactly what world the argument inhabits before it begins.

The phrases 'for definiteness' and 'suppose also' are doing important rhetorical work: they signal that these are stipulated conditions, not empirical claims. This is the correct way to introduce assumptions — not apologetically, but transparently, as an explicit framing of the analytical space within which the argument will operate.

The practical lesson for faculty writing: state your assumptions before your argument requires them, not after the reader has already started to question them. An assumption introduced as a defence against a challenge is far less persuasive than one established as part of the argument's foundation.

The obligation to state assumptions clearly extends to the conclusion. A finding that holds under specific modelling assumptions should not be presented as if it holds generally. This is not excessive caution. It is intellectual honesty — and referees, editors, and readers will notice when it is absent.

1.4 Hedging and Epistemic Responsibility

Academic writing in all disciplines requires the writer to match the strength of a claim to the strength of the evidence behind it. In economics, this is especially important because the discipline's claims often have policy implications — and overstated findings can, in the worst cases, affect decisions that matter to many people. The linguistic tools for doing this are called hedges: words and phrases that adjust the degree of certainty attached to a claim.

There is, however, a tension in hedging that writers must handle carefully. Under-hedging — stating conclusions more confidently than the evidence supports — is an intellectual failure and a common reason for referee rejection. But over-hedging — qualifying every claim until the paper appears to argue nothing at all — is an equally serious failure of a different kind. It is the approach taken by writers who are not sure what they are claiming, or who are unwilling to commit to a position that might be criticised.

◆ The grammar of hedged claims

The formulations below represent a spectrum from strong assertion to heavy qualification. Effective economics writing uses the full range deliberately, reserving the strongest formulations for the claims best supported by the evidence:

Formulation	Epistemic strength
<i>The evidence demonstrates that...</i>	Strong — use only when support is unambiguous
<i>The results show that...</i>	Strong — appropriate for direct empirical findings
<i>The findings suggest that...</i>	Moderate — appropriate for interpretive claims
<i>The data are consistent with...</i>	Moderate — does not rule out alternative explanations

<i>It appears that...</i>	Cautious — appropriate where uncertainty is genuine
<i>One interpretation is that...</i>	Weak — signals that other interpretations exist
<i>This may indicate that...</i>	Weak — appropriate for speculative or preliminary claims

Friedman's (1968) paper uses hedging sparingly but deliberately. Where the theoretical logic is tight, he states conclusions directly. Where the empirical grounding is thinner, he qualifies. The sentence 'Unfortunately, we have as yet devised no method to estimate accurately and readily the natural rate of either interest or unemployment' is honest about a limitation without weakening the theoretical argument, which does not depend on the natural rate being measurable — only on its being real.

The practical rule is this: hedge the claim to the level the evidence supports, and no further in either direction. A paper that consistently over-hedges signals a writer who does not believe in their own argument. A paper that consistently under-hedges signals a writer who has not yet engaged seriously with their referee reports.

1.5 Paragraph Structure as Argument Structure

If the four-move sequence is the logic of argument, the paragraph is the unit in which that logic appears on the page. A well-constructed economics paragraph does not merely group related sentences. It advances the argument by a specific, identifiable step. Each paragraph should make one — and only one — distinct argumentative move, and the reader should be able to say clearly what that move is.

◆ The topic sentence as argumentative commitment

The first sentence of a paragraph should function as an argumentative commitment: a statement of what the paragraph will establish, not just a label for what it will discuss. This distinction matters because it determines whether the paragraph has a direction or merely a subject.

✗ Label (no argument)	✓ Argumentative commitment
-----------------------	----------------------------

<i>This section discusses the relationship between money supply and inflation.</i>	<i>The relationship between money supply and inflation is stable in the long run but subject to significant short-run variation driven by velocity shocks.</i>
<i>There are several reasons why the natural rate of unemployment changes over time.</i>	<i>The natural rate of unemployment is not a fixed structural constant but shifts in response to changes in labour market institutions, demographic composition, and the sectoral structure of the economy.</i>

The left-hand formulations tell the reader what subject is coming. The right-hand formulations tell the reader what the paragraph will argue about that subject. A paper whose topic sentences are consistently argumentative is a paper whose argument is visible at paragraph level — meaning a referee can follow the paper's logic without reading every word. This is not a concession to lazy reading. It is good academic writing.

◆ Development, evidence, and analysis within the paragraph

Once the topic sentence has made its commitment, the sentences that follow must deliver on it. The most common paragraph-level failure in economics manuscripts is a mismatch between the topic sentence's claim and what follows: a paragraph that promises to establish X but provides only background information about X, or that shifts midway to an adjacent claim without completing the first one.

Friedman's 1968 paper is a good model at the paragraph level. Immediately after his introductory discussion of the two targets of monetary policy — the interest rate and the unemployment rate — he writes:

Every central bank, since its inception, has regarded the regulation of the quantity of money as one of its primary functions. Historically, regulation of the interest rate has been the primary means through which central banks have exercised this function. More recently, there has been increasing emphasis on the level of employment or unemployment as an additional target. However, both targets share a fundamental defect that makes them unsuitable as guidelines for monetary policy: they are defined in terms of nominal rather than real magnitudes.

Friedman, M. (1968). The role of monetary policy. *American Economic Review*, 58(1), p. 5.

Paragraph structure: how it works

The final sentence — 'both targets share a fundamental defect' — is the argumentative payload of the entire paragraph. Everything before it is development: historical context, the naming of both targets, their sequential introduction to the reader. The structure is deliberate: Friedman builds towardss the claim rather than opening with it, because the claim needs its context to be intelligible.

This is a valid structural choice — sometimes the topic sentence comes last, functioning as a conclusion rather than a heading. What matters is that the paragraph has one clear argumentative move, and that the reader can identify it. Here, the move is: both standard targets of monetary policy share a defect that disqualifies them.

The phrase 'fundamental defect' is doing double work: it is both analytically precise (the defect is defined in the next sentence) and rhetorically forceful (it signals that what follows is not a minor caveat but a structural critique). This is economy of language in the service of argument.

◆ Transitions between paragraphs

The transition between paragraphs is the point where a paper's momentum is most easily lost. A transition that only announces a new topic — 'Another important factor is...' or 'Turning now to...' — interrupts the argument. It tells the reader that one thing has ended and another has begun, without explaining how they are connected. A transition that links the previous paragraph's conclusion to the following paragraph's claim maintains the argument and keeps the reader oriented in the paper's larger structure.

The practical rule is simple: the opening sentence of each paragraph should connect, explicitly or implicitly, to the closing sentence of the previous one. The connection does not always need to be stated — sometimes the logical relationship is clear from the content. But the writer should always be able to articulate it. If they cannot, the paragraphs do not yet belong together, or their order needs to change.

1.6 Sentence-Level Precision

The qualities of argument discussed in sections 1.1 through 1.5 operate at the level of the paper's overall structure: its thesis, its internal logic, its paragraph organisation. But the force of that structure depends, in the end, on the quality of individual sentences. A sentence that is grammatically correct but imprecise, overly long, or unnecessarily passive weakens the argument even if the structure is sound.

◆ Economy of expression

Economics favours parsimony in models. The same principle applies to writing: say what needs to be said, and nothing more. This does not mean writing only short sentences or using a minimal style. It means that every word in a sentence should be doing something useful. Consider the following revision:

✗ Inflated	✓ Concise
<i>It is possible to argue that there exists a significant degree of uncertainty with respect to the relationship between the variables under consideration.</i>	<i>The relationship between these variables remains uncertain.</i>
<i>Due to the fact that the data do not provide sufficient evidence to support the claim that...</i>	<i>Since the data do not support the claim that...</i>

The inflated formulations are not wrong; they are wasteful. They ask the reader to do unnecessary work, and they hide the argumentative content behind procedural language. In a paper with many such sentences, the cumulative effect is significant: the argument becomes harder to follow, and the writing signals a writer who has not yet committed fully to what they are saying.

◆ The grammar of causation

Economics is fundamentally about causal relationships — between variables, between policies and outcomes, between institutions and economic performance. The language of causation must therefore be used with care. Different verbs and connectives carry different degrees of causal strength, and the choice should reflect what the analysis and evidence actually support.

Expression	Causal strength implied	Use when...
<i>X causes Y</i>	Strong / direct causation	Experimental or quasi-experimental identification

<i>X leads to Y</i>	Strong / directional	Robust causal mechanism established
<i>X increases Y</i>	Moderate / directional	Strong correlational evidence with plausible mechanism
<i>X is associated with Y</i>	Weak / correlational	Observational data, no causal identification
<i>X is consistent with Y</i>	Neutral	Data do not rule out the relationship
<i>X may contribute to Y</i>	Cautious / speculative	Exploratory analysis or preliminary findings

One of the most serious writing errors in empirical economics is using causal language — 'X increases Y', 'X leads to Y' — when the identification strategy supports only correlational claims. This is not just a stylistic error. It is an analytical misrepresentation. Referees in empirical fields are acutely alert to it. The choice of verb is not cosmetic; it is a claim about what the paper has established.

◆ **Technical vocabulary: precision without over-jargonising**

Economics has a precise technical vocabulary — terms such as elasticity, externality, moral hazard, rent-seeking, and time inconsistency carry specific analytical meanings that cannot be replaced by everyday equivalents without losing precision. Using this vocabulary correctly is a mark of disciplinary competence. Writers should not avoid it when writing for an economics audience.

The failure to avoid, however, is using technical vocabulary as a substitute for clear argument. Jargon used without precise meaning produces writing that sounds authoritative whilst communicating very little. The test is simple: can you explain, in plain language, what the technical term means in the specific context where you are using it? If not, the term is probably hiding a gap in the argument rather than expressing a precise idea.

1.7 Vocabulary and Sentence Models

This section provides a reference bank of key terms and sentence patterns used throughout the chapter. The vocabulary items are the technical and semi-technical words that appear frequently in the genre of economic argument. The sentence models are drawn from established

economics writing and can be adapted for your own papers. Each model is annotated to show the function it performs.

◆ Key vocabulary

Term	Meaning in economics writing	Example of the word in a sentence
thesis	The central claim of a paper — a specific, contestable statement of what the paper argues. Distinct from a topic (what the paper is about) or a question (what the paper asks).	<i>'This paper argues that persistent inflation results not from excessive money growth in isolation, but from the interaction between monetary expansion and anchored expectations.'</i> (thesis statement)
argument	A structured sequence of claims, each supported by evidence or reasoning, that leads to a conclusion the reader could not reach from the evidence alone. In economics, an argument is not a dispute; it is a chain of justified reasoning.	<i>'Friedman's argument proceeds from the definition of the natural rate, through the mechanism by which pegging the nominal rate creates disequilibrium, to the conclusion that sustained low unemployment cannot be purchased with permanent inflation.'</i>
claim	A statement that asserts something to be true and that could, in principle, be shown to be false. Claims are the building blocks of argument. A claim without evidence or reasoning to support it is an assertion.	<i>'The claim that monetary policy cannot permanently lower unemployment requires a theory of expectations, which Friedman provides in the second section of the paper.'</i>
hedge	A word or phrase that reduces the certainty of a claim — for example, 'tends to', 'is likely to', 'appears to', 'suggests that'. Hedges are used when evidence is incomplete or when the relationship is probabilistic rather than universal.	<i>'The evidence suggests that regions with high social capital tend to have higher intergenerational mobility, though the direction of causality remains uncertain.'</i>

assumption	A condition stipulated to be true for the purposes of an argument, without full empirical verification. Assumptions bound the context in which a theoretical result holds. Economists state assumptions explicitly so that readers can evaluate whether they apply.	<i>'We assume competitive labour markets, so that wages equal the marginal product of labour. Relaxing this assumption — as in models with monopsonistic employers — yields different predictions about minimum wage effects.'</i>
implication	A consequence that follows logically from an argument's claims and evidence. A paper's implications section shows what the analysis means beyond the specific results — for policy, for theory, or for future research.	<i>'The implication of this finding is that place-based policies may be more effective than previously recognised, since mobility varies substantially across regions that are administratively similar.'</i>
contestable	Capable of being disputed or shown false. A good thesis is contestable: a reasonable person, not yet persuaded, could disagree. A claim that no one could dispute is not a thesis — it is a description.	<i>'The statement that monetary policy has effects on real variables in the short run is contestable; economists have disagreed about the size, sign, and duration of those effects for decades.'</i>
parsimony	The principle of using the simplest explanation or model that is adequate for the purpose. In writing, parsimony means using no more words than necessary. In modelling, it means using no more assumptions than necessary.	<i>'The model is deliberately parsimonious: it includes only the two mechanisms that the paper's evidence can speak to, and defers the richer general equilibrium treatment to future work.'</i>

◆ Sentence models

The sentences below model key moves in economic argument writing. Each can be adapted by replacing the bracketed elements with your own content.

Function	Model sentence	Notes for the writer
----------	----------------	----------------------

Stating the thesis	<i>This paper argues that [X], and that this relationship holds specifically under [conditions Y], with implications for [Z].</i>	Place this move at the end of the introduction. The 'specifically under' clause bounds the claim; the 'implications for' clause shows why it matters.
Distinguishing argument from description	<i>Rather than describing [phenomenon], this paper explains [mechanism]: [claim].</i>	Use when the paper's contribution is analytical rather than empirical. The contrast 'rather than describing... explains' signals the argumentative register explicitly.
Stating an assumption transparently	<i>We assume [X]. This assumption holds when [conditions], and is relaxed in [section/extension].</i>	State every non-trivial assumption, give the conditions under which it is appropriate, and flag where it is tested or relaxed. This is how Friedman introduces his stipulated conditions.
Hedging a claim appropriately	<i>[X] tends to [Y], though the relationship appears to depend on [condition Z], and may not hold in [context].</i>	Use modal verbs ('tends to', 'appears to', 'may') and conditional clauses to calibrate certainty. Avoid both bare assertion and paralysing qualification.
Writing an argumentative topic sentence	<i>[Key term] plays a more important role in [outcome] than [prior accounts] have recognised.</i>	Topic sentences should commit to a position. This model makes the comparative claim ('more important than') that gives the paragraph its direction.
Making the analytical move in a paragraph	<i>What [finding/observation] reveals is not simply [surface interpretation], but [deeper mechanism]: [explanation].</i>	The phrase 'what X reveals is not... but' performs the analytical move from observation to interpretation. It is the difference between reporting and arguing.

Stating the implication of a finding	<i>If [result] is correct, then [consequence follows] — which suggests that [policy/theoretical conclusion], provided that [limiting condition].</i>	Link the empirical result to its consequence, then add the limiting condition that qualifies the implication. The 'provided that' clause prevents overreach.
Positioning the paper's contribution	<i>This analysis extends [prior work] by [specifying what is new], and contributes to [debate/literature] by showing that [claim].</i>	Contribution statements must be specific. 'Extends prior work' alone says nothing; the 'by' clause carries all the weight.

The principles in this chapter — argument over description, the four-move structure, explicit assumptions, calibrated hedging, paragraph coherence, and sentence-level precision — are not stylistic preferences. They are the conditions under which economic writing becomes persuasive: the conditions under which a reader who does not yet believe your argument might come to do so. Everything in the chapters that follow applies these conditions to specific genres and specific writing contexts.

→ **Chapter Two: The Research Paper for Publication**

CHAPTER TWO

The Research Paper for Publication

The research paper is the central form of academic economics. It is where new knowledge is developed, evaluated, and, if successful, published within the discipline. Other genres in this guide either prepare for or develop from the research paper. This chapter approaches the research paper not as a fixed format, but as an argumentative task. Each section has a clear intellectual purpose, and understanding this purpose — and the consequences of not fulfilling it — is essential for writing that can pass peer review.

The chapter uses Milton Friedman's *The Role of Monetary Policy* (1968) as a key example. Although not recent, it is a highly self-aware paper, with a clear argument that shapes every section. More than fifty years later, it remains a strong model of effective economic writing.

2.1 Understanding What Journals Expect

Before discussing the structure of a research paper, it is useful to establish the external context in which papers are submitted and evaluated. An economics manuscript arrives at an editorial office with a specific identity, a specific readership, and a specific set of expectations about what constitutes a publishable contribution. Writers who do not understand that context are at a disadvantage from the start.

◆ The anatomy of a publishable paper: IMRaD and its variations

The standard structure of an empirical economics research paper follows the IMRaD model — Introduction, Methodology, Results, and Discussion — with a literature review either inside the introduction or as a separate section, depending on the journal and the nature of the contribution. Theoretical papers may vary significantly from this structure, organising sections around the model's derivation, properties, and implications rather than around an empirical workflow. The table below summarises the main function of each standard section and the most common failure associated with each.

Section	Primary function	Most common failure
Introduction	Establish the gap, announce the contribution, roadmap the paper	Describes context without identifying a gap; announces method rather than contribution
Literature Review	Synthesise the field; position the paper's contribution	Catalogues sources without synthesis; fails to create the gap the paper fills
Methodology	Justify analytical choices; link method to question	Describes procedure without justification; omits the link to the research question

Results	Present findings with precision; distinguish data from interpretation	Interprets in the results section; omits reference to tables and figures in prose
Discussion	Interpret findings; engage the literature; acknowledge limits	Restates results; fails to engage with existing work; over-claims conclusions
Conclusion	Answer the research question; state limitations; open further enquiry	Summarises rather than concludes; introduces new material; over-generalises

Read the table from right to left as well as left to right. The failure column is not an abstract list of mistakes to avoid. It describes specific writing habits that produce specific types of referee objection. Understanding the connection between a section function and its characteristic failure is one of the most direct ways to pre-empt rejection.

◆ Matching structure to journal conventions

Different journals have different structural expectations, and reading the author guidelines carefully is not optional. Some journals require a structured abstract with fixed headings; others expect a single paragraph. Some have strict word limits that apply separately to the main text and appendices; others count everything together. Some expect robustness checks in the body of the paper; others place them in online supplements. These are not trivial details. A submission that does not follow a journal's conventions signals to the editor that the authors have not read the journal they are submitting to — which is, at best, careless, and at worst a sign of a scattergun submission strategy that editors recognise and dislike.

Beyond formal conventions, each journal has expectations about the nature and scope of the contributions it publishes. The American Economic Review publishes papers that make fundamental contributions to broad areas of the discipline. The Journal of Development Economics publishes work with direct relevance to low- and middle-income country contexts. The Review of Economic Studies treats methodological rigour as a near-absolute condition. Submitting a narrowly empirical country study to the AER, or a purely theoretical paper to a policy-oriented journal, reflects a misunderstanding of what the journal is for — and the paper will be rejected, regardless of its quality.

◆ What the covering letter must do

The covering letter is the first thing an editor reads, and it is a writing task in its own right. Its purpose is not to reproduce the abstract but to present the paper's contribution in terms that are directly relevant to the journal's readership and scope. A strong covering letter answers

three questions concisely: what is the paper's central claim, why does that claim matter for the debates this journal's readers are engaged with, and what does the paper contribute that existing work does not. A letter that says only 'Please find enclosed our paper for consideration' has not done this work.

2.2 Developing and Articulating the Research Question

The research question is the organising principle of the entire paper. Every section — the literature review, the methodology, the results, the discussion — gets its purpose and content from the question it is designed to address. A paper with a poorly defined research question is structurally incoherent from the start, because no section can do its job when the question it is answering is unclear.

A publishable research question in economics must satisfy three conditions at the same time. It must be tractable — answerable, in principle, using available data and methods. It must be non-trivial — not already answered in the existing literature, or capable of being extended, challenged, or refined. And it must be consequential — its answer must matter for the discipline, for policy, or for how an economic phenomenon is understood. A question that is tractable but trivial produces a technically competent paper with nothing to say. A question that is consequential but intractable produces an ambitious proposal with no executable research design. Editors and referees are looking for all three conditions together, and the introduction of the paper must show that they are present.

◆ From broad topic to focused question

The most common failure in research question development is insufficient specificity. A question such as 'What is the relationship between monetary policy and inflation?' is not a research question; it is a field of study. Moving from topic to question requires the writer to identify a specific aspect of the broader phenomenon that is both unresolved in the literature and addressable with a defined methodology. This process of narrowing is not a concession. It is the condition under which rigorous argument becomes possible.

Friedman's 1968 paper shows what a sharpened question looks like. The broad topic is monetary policy. The specific question is whether monetary policy can be used to fix either the interest rate or the unemployment rate at a target level. This is a question with a clear logical structure — it can be answered yes or no — and it has a direct connection to the policy debates of the time. That specificity is what allows the argument to be decisive rather than merely suggestive.

The relationship between the research question and the chosen methodology is not incidental. The question constrains the method: a question about causal mechanisms requires an identification strategy; a question about distributional effects requires disaggregated data; a question about long-run dynamics requires a sufficiently long time series. Writers who choose a method first and then build a research question around it tend to produce papers where the question and the method sit awkwardly together — a structural mismatch that referees notice quickly.

2.3 The Introduction

The introduction is the most important section of a research paper. It is where the editor decides whether to send the paper for review, and where the referee forms the first impression that will shape their reading of everything that follows. A strong introduction does not just describe what the paper will do. It builds a case for why the paper matters and implicitly promises what it will deliver. That promise must be kept throughout the paper.

The introduction has four distinct functions. Each must be performed, roughly in the order given, and each has a characteristic writing failure associated with it.

◆ Function 1 — Context: situating the problem

The opening of the introduction must establish the economic phenomenon or question at stake, explain why it matters, and place it within the landscape of existing work. This is not the literature review. It is a brief, purposeful account of the context that makes the paper's specific question comprehensible and its significance clear. The failure here is over-contextualisation: introductions that spend two or three pages on background material before reaching the research question have buried the paper's contribution under scene-setting that the reader did not need in such detail.

Friedman opens 'The Role of Monetary Policy' with an immediate statement of context — the history of professional opinion on monetary policy — that is concise and precisely calibrated to set up the paper's central claim:

There is wide agreement about the major goals of economic policy: high employment, stable prices, and rapid growth. There is less agreement that these goals are mutually compatible or, among those who regard them as incompatible, about the terms at which they can and should be substituted for one another. There is least agreement about the role that various instruments of policy can and should play in achieving the several goals. My aim in this lecture is to distinguish sharply between two roles of monetary policy, to specify the limits of monetary policy, to identify the activities it can perform, and to draw some implications for current policy.

Friedman, M. (1968). The role of monetary policy. *American Economic Review*, 58(1), p. 1.

What this opening achieves

In three sentences, Friedman establishes the intellectual landscape (competing views on policy goals), locates the area of maximum disagreement (the role of instruments), and announces precisely what the paper will do. The final sentence of the second paragraph — 'My aim in this lecture is to...' — is a model of the roadmap sentence: it is specific, it lists discrete tasks, and it implies a structure that the reader can follow.

Notice the movement from broad agreement to narrow disagreement. Friedman does not begin with a general statement about monetary policy; he begins with a specific observation about the structure of professional debate. This immediately signals that the paper has a precise intellectual target. The opening does not waste the reader's attention — every sentence is doing work.

◆ Function 2 — Gap: identifying what is unresolved

The gap statement is the intellectual heart of the introduction. It identifies what the existing literature has not established, has established incorrectly, or has established only in part — and explains why that gap matters. The gap must be genuine. A paper that claims novelty in a space that existing work has already occupied — either through ignorance of that work or through misleading framing — will be rejected on that basis alone.

The gap statement must also be specific. 'This area has not been fully explored' is not a gap statement; it is an evasion. A proper gap statement identifies a specific claim, relationship, or mechanism that the literature has not addressed and explains why the paper can address it. In Friedman's case, the gap is stated directly: the existing literature on monetary policy has confused two targets — the interest rate and the unemployment rate — with the instruments available to achieve them, and has therefore reached incorrect conclusions about what monetary policy can accomplish.

◆ Function 3 — Contribution: what this paper does

The contribution statement explains what the paper adds to knowledge. It should be specific, honest about scope, and clearly distinguished from a description of the paper's content. There is a significant difference between 'This paper examines X' and 'This paper shows that X'. Only the second is a contribution statement.

The contribution statement is also where writers most often over-claim. A paper that says it will 'resolve' a long-standing debate, 'transform' understanding of a phenomenon, or 'definitively establish' a relationship that the data can support only tentatively is setting itself up for a hostile referee response. The contribution must be proportionate to what the paper

actually delivers — significant enough to be worth publishing, honest enough to survive scrutiny.

✗ Over-claimed contribution	✓ Proportionate contribution
<i>This paper resolves the long-standing debate about the relationship between trade openness and inequality by providing definitive evidence that openness increases inequality in all developing countries.</i>	<i>This paper provides evidence that trade openness is positively associated with within-country income inequality in a panel of 38 developing economies over 1990–2015, with the relationship strongest in countries with relatively underdeveloped financial sectors.</i>
<i>This paper fundamentally revises our understanding of the transmission mechanism of monetary policy.</i>	<i>This paper identifies a previously undocumented channel through which changes in the policy rate affect small-firm investment in credit-constrained environments, using administrative data from Poland's 2012 monetary tightening cycle.</i>

◆ Function 4 — Roadmap: signposting the paper's structure

The roadmap paragraph is the final element of the introduction. Its purpose is functional, not argumentative: it tells the reader how the paper is organised so that they can navigate it. The failure in roadmap paragraphs is mechanical enumeration — a sentence-by-sentence account of each section's content that reads like a list rather than an orientation. A well-written roadmap identifies the logical sequence of the paper's argument and briefly explains why each section is necessary for the argument to proceed.

The roadmap should be short. Two or three sentences are usually enough. A roadmap paragraph that runs to half a page has gone beyond its intended purpose and is taking up space that should be devoted to the argument.

2.4 The Literature Review within a Research Paper

The literature review within a research paper is not a survey of everything written on the topic. It is a selective, purposeful account of the scholarship that directly informs the paper's research question and positions its contribution. Its function is to show that the writer knows the field, that the gap identified in the introduction is genuine, and that the paper's approach is grounded in existing methodology and theory.

The embedded literature review in an empirical paper is typically built around three moves: establishing what is already known with confidence; identifying what remains contested or unresolved; and explaining how the present paper addresses what is unresolved. This three-move structure connects the literature review directly to the gap statement in the introduction and prevents the review from becoming a free-standing catalogue of related work.

◆ **Synthesis rather than summary**

The most common failure of embedded literature reviews at publication level is substituting summary for synthesis. Summary reports what individual sources say. Synthesis identifies what the literature as a whole reveals — its areas of agreement, its methodological divisions, its unresolved debates, its empirical inconsistencies. A literature review that moves from source to source without integrating them has not yet performed its function.

Friedman's handling of the existing literature on monetary policy targets is characteristically synthetic. Rather than listing positions chronologically, he organises the literature around a structural argument: that the history of central banking reveals a pattern of targets that are either impossible to achieve or self-defeating in the long run. The literature is used not as background but as evidence for a claim about the systematic errors that the paper will correct.

◆ **Calibrating reference density**

Reference density — the number of citations per unit of text — varies across economics subfields and journal conventions, and calibrating it correctly is a matter of judgement rather than formula. Under-citation signals insufficient engagement with the literature. Over-citation produces writing that is more footnote than argument, and may suggest a writer who is using citations defensively rather than engaging critically with the scholarship.

As a working principle: cite the most important, most rigorous, and most directly relevant work; do not cite every paper that has touched the topic. A referee who is a specialist in the field will notice both significant omissions and unnecessary padding, and will respond negatively to both.

2.5 The Methodology Section

The methodology section has a specific and often misunderstood function. It is not a description of what the researcher did. It is a justification of why those choices — of data source, model specification, identification strategy, and estimation technique — are the right ones for the research question. The difference between description and justification is the difference between a methods section and an argument.

◆ **Describing data with precision**

Data description should be precise, complete, and organised to help the reader understand the analysis that follows. The reader needs to know the source, coverage, and unit of observation of the data; the time period; any cleaning or construction decisions that affect the sample; and any limitations that will be relevant when evaluating the results. These are not procedural details. They are part of the evidential basis of the paper, and gaps in data description create vulnerabilities that referees will identify.

The failure to avoid is what might be called the 'black box' description: a brief statement that data were obtained from a named source, with no account of how they were processed, which observations were included or excluded, or what the key variables actually measure. This is not adequate at publication level, and it typically indicates that the writer has not yet distinguished between their own research notes and the account a reader needs to evaluate the analysis.

◆ Justifying the identification strategy

In empirical economics, the identification strategy — the approach by which the paper establishes that the relationships it observes are causal rather than merely correlational — is the methodological core of the paper. The methodology section must not only describe the identification approach but justify it: explain why this strategy is appropriate for the research question, what assumptions it requires, and why those assumptions are plausible given the data and context.

Friedman's (1968) paper is theoretical rather than empirical, so it has no identification strategy in the econometric sense. But it has an analytical identification strategy: the use of the distinction between real and nominal magnitudes as the key that unlocks the paper's central argument. That distinction is introduced in the methodology-equivalent section of the paper — where Friedman sets out the framework of the argument — with explicit acknowledgement of the assumptions it requires:

To state this conclusion differently, there is always a temporary trade-off between inflation and unemployment; there is no permanent trade-off. The temporary trade-off comes not from inflation per se, but from unanticipated inflation, which generally means, from a rising rate of inflation. The widespread belief in the existence of a permanent trade-off is a sophisticated version of the confusion between 'high' and 'rising' that we all recognise in simpler forms. A rising rate of inflation may reduce unemployment, a high rate will not.

Friedman, M. (1968). The role of monetary policy. *American Economic Review*, 58(1), p. 11.

The analytical framework as a writing achievement

This passage is the methodological heart of the paper. Friedman is not merely making a claim; he is identifying the conceptual distinction — between levels and rates of change, between anticipated and unanticipated inflation — that makes the claim possible. Without this distinction, the paper's central argument cannot be stated.

The final two sentences are a model of analytical compression: 'A rising rate of inflation may reduce unemployment, a high rate will not.' This is the paper's key theoretical result, stated in plain English, without a single equation. The methodological achievement is not computational but conceptual — and the writing is equal to it.

The lesson for methodology sections: the framework or identification strategy should be introduced not merely as a procedure but as an intellectual move — one that makes the analysis that follows both possible and justified. The reader should be able to see why the analytical framework was necessary, not just what it is.

◆ Linking method to question

The methodology section must close the connection between the paper's research question and its analytical approach. A reader who finishes the methodology section should be able to say clearly: given this question, these data, and this method, the paper is in a position to produce findings that bear directly on the question. If that connection is not explicit, the paper has a structural gap at its methodological core — and referees will identify it.

2.6 The Results Section

The results section presents the paper's empirical or analytical findings with precision and clarity. Its function is to show what the data reveal, what the model produces, or what the analysis demonstrates — and to do so in a way that is unambiguous, complete, and directly connected to the research question. What it must not do is interpret: interpretation belongs in the discussion section, and mixing the two creates a structural problem that makes both sections harder to evaluate.

◆ Prose, tables, and figures: who leads

A results section in empirical economics typically combines prose, tables, and figures. The prose leads: it directs the reader's attention, identifies the key findings, and explains what the tables and figures show without simply repeating their content. A table or figure that is

never discussed in the prose has not been integrated into the paper. It is merely attached to it, and referees will note the omission.

The key principle for writing about tables and figures is selectivity. The prose should not walk through every cell or data point. It should identify the findings that matter most for the paper's argument and explain their relevance. The reader can consult the table for full details. The prose should tell them where to look and why it matters.

✗ Narrating the table	✓ Directing the reader
<i>Table 2 shows that in 1967, the unemployment rate was 3.8 per cent. In 1968 it was 3.6 per cent. In 1969 it was 3.5 per cent. In 1970 it rose to 4.9 per cent.</i>	<i>Table 2 shows a sustained decline in the unemployment rate through the late 1960s, followed by a sharp reversal in 1970 — a pattern consistent with the short-run trade-off predicted by the model, and with its subsequent breakdown.</i>
<i>The coefficient on the trade openness variable is 0.34. The standard error is 0.12. The t-statistic is 2.83. The p-value is 0.005.</i>	<i>The coefficient on trade openness is positive and statistically significant at the 1 per cent level ($\beta = 0.34$, $SE = 0.12$), suggesting a meaningful association with the dependent variable in the baseline specification.</i>

◆ Maintaining the results/discussion distinction

The boundary between the results section and the discussion section is one of the most frequently crossed structural conventions in economics manuscripts. Results sections that interpret, contextualise, or draw implications from findings are doing the discussion work — and in doing so, they deprive the discussion of its content and blur the paper's analytical structure.

The practical test is simple: a sentence that begins 'This suggests that...' or 'This finding implies...' belongs in the discussion, not in the results. A sentence that begins with 'The coefficient is...' or 'The model produces...' belongs in the results. This distinction is not pedantic. It reflects the difference between reporting what was found and arguing about what it means.

2.7 The Discussion Section

The discussion section is the intellectual core of the research paper. It is where the paper's findings are turned from reported results into argued contributions: where the evidence is connected to the claim, where the literature is engaged rather than merely cited, and where the paper answers the question it set out to address. If the introduction is the paper's promise and the results section is its evidence, the discussion is its argument — and it is the section that most directly determines whether the paper makes a genuine contribution.

It is also the section most often underdeveloped in submitted manuscripts. Writers who have worked intensively on data collection, model specification, and results presentation often produce discussion sections that are thin, formulaic, or simply restatements of the results in different words. This is one of the most reliable predictors of rejection at the revision stage: a paper with strong empirics and a weak discussion will be sent back for precisely that reason.

◆ From findings to interpretation to implication

The discussion must do three things in sequence. First, it must interpret the findings: explain what they mean, why the results are what they are, and how they connect to the analytical framework set out in the methodology. Second, it must engage the existing literature: show how the findings confirm, contradict, or extend what prior work has established, and explain why discrepancies arise. Third, it must state the implications: what the findings mean for the broader literature, for policy, and for the research agenda.

Friedman's discussion of the implications of the natural rate hypothesis is a good model of this three-move structure. Having established the theoretical result — that there is no permanent trade-off between inflation and unemployment — he does not simply restate the finding. He extends it into its policy implications with explicit logical steps:

A second limitation of monetary policy is that it can force a change in nominal quantities — the price level, the exchange rate, the nominal level of national income, the quantity of money — but it cannot, except for transitory periods, force a change in real quantities — the real rate of interest, the real wage rate, the level of employment, the real level of national income. The reason is that any attempt to force a change in real quantities will set in train reactions that alter nominal quantities in the same direction.

This is the fundamental limitation that makes it impossible for monetary policy to peg real quantities for any extended period.

Friedman, M. (1968). The role of monetary policy. *American Economic Review*, 58(1), pp. 10–11.

The discussion as logical extension, not restatement

Friedman does not restate his earlier results; he extends them. The distinction between nominal and real quantities — introduced in the methodology section of the paper — is now deployed to generate a new implication: the fundamental limitation of monetary policy. This is what a discussion section should do: take the findings established in the results and develop their consequences.

The phrase 'This is the fundamental limitation' is the kind of declarative sentence that many writers avoid, fearing it will seem too assertive. But this is precisely the function of the discussion: to state clearly what the analysis has established. Hedging this conclusion — 'This may perhaps suggest a possible limitation...' — would be a failure of intellectual nerve, not a mark of scholarly caution.

Notice also the structural clarity of the two-paragraph construction: the first paragraph presents the mechanism; the second states the conclusion. The reader is given the argument before they are given the verdict.

◆ Acknowledging limitations without undermining the argument

Every empirical paper has limitations — in its data, its identification strategy, its scope, or its generalisability. Acknowledging those limitations is a mark of intellectual honesty, and referees will notice their absence. The failure to avoid, however, is the limitations section that is so extensive or so fundamental that it raises questions about whether the paper's conclusions can be trusted at all.

The principle is proportionality: acknowledge the limitations that are most significant and most likely to be raised by referees, explain why they do not invalidate the paper's central findings, and move on. A limitation that genuinely undermines the paper's conclusions should be addressed through additional analysis — robustness checks, alternative specifications, bounding exercises — rather than simply acknowledged and left to stand.

2.8 The Conclusion

The conclusion of a research paper is not a summary of what has gone before. A conclusion that simply repeats, in condensed form, the introduction and the discussion has not done its job. The conclusion must answer the research question directly and explicitly — tell the reader, in clear terms, what has been established — and explain what follows from that answer: for the literature, for policy, or for the next generation of research.

Neugeboren and Jacobson (2005, Ch. 6) identify the habit of 'retreating to the introduction' as the most common conclusion failure: the writer, unsure how to close, returns to the broad contextual framing of the opening paragraphs rather than delivering the specific, grounded conclusion that the analysis has produced. The result is a paper that ends where it began, creating the impression that the intervening work was circular rather than progressive.

◆ Answering the research question directly

The opening of the conclusion should state, clearly and without excessive qualification, what the paper has found and why it matters. This is not the place for hedging that belongs in the discussion. By the time the reader reaches the conclusion, the analysis has been performed and the evidence assessed. The conclusion's opening is where the writer must be willing to state plainly what the paper has established.

Friedman's conclusion to the 1968 paper is a good model of this directness. Having traced the argument from the limits of monetary policy through the natural rate hypothesis to the dangers of activist stabilisation policy, he ends with a clear statement of what good monetary policy should look like:

The first and most important lesson that history teaches about what monetary policy can do — and it is a lesson of the most profound importance — is that monetary policy can prevent money itself from being a major source of economic disturbance. The Great Contraction might not have occurred at all, and if it had, it would have been far less severe, if the monetary policy that has prevailed from 1929 to 1933 had been replaced by the policy of steady monetary growth.

As it happens, preventing money from being a source of economic disturbance does not call for any drastic or spectacular action. It calls for a steady course.

Friedman, M. (1968). The role of monetary policy. *American Economic Review*, 58(1), p. 17.

The conclusion as earned verdict

This conclusion works because it is proportionate to the argument that precedes it. Friedman does not over-claim: he does not say that monetary policy is irrelevant or that it can solve all economic problems. He says precisely what the analysis supports — that the most important function of monetary policy is negative (preventing money from being a source of disturbance) and that this function requires not activism but steadiness.

The phrase 'it calls for a steady course' is the payoff of seventeen pages of careful argument. It is direct, it is specific, and it is proportionate. It is also, characteristically, expressed in plain English — a reminder that the conclusion is the moment when the writer owes the reader maximum clarity, not maximum complexity.

◆ Policy implications: handling with care

Many economics papers, particularly in applied and development economics, end with a discussion of policy implications. These are valuable — they are part of what makes research consequential — but they require care in how they are framed. Policy implications should follow logically from the paper's findings, be calibrated to the scope of the evidence, and acknowledge the gap between what a paper can establish and what policy decisions require.

A paper that finds a positive association between trade openness and growth in a panel of middle-income countries cannot legitimately conclude that all countries should liberalise trade. The scope of the finding — specific countries, specific period, specific measure — constrains the scope of the implication. Over-reaching in the policy implications section is one of the more reliable triggers for a request for major revision.

◆ Directions for further research

The section on further research directions, where included, should be substantive rather than formulaic. The standard closing sentence — 'Further research is needed in this area' — is not a conclusion; it is a placeholder. Genuine directions for further research identify specific questions that arise from the paper's findings, specific data or methodological developments that would allow those questions to be addressed, and specific aspects of the phenomenon that the paper could not examine but that its results make newly relevant.

2.9 Writing the Abstract

The abstract is the most widely read part of any paper. Editors use it to decide whether to send a paper out for review; referees use it to orient themselves before reading; researchers use it to decide whether the full paper is worth their time. Despite this, abstracts are among the most consistently underwritten components of economics manuscripts. Many writers treat the abstract as a summary of what the paper does rather than what it finds — describing the research design rather than communicating the results. An abstract of this kind fails at its primary function.

The practical discipline is this: write the abstract last, once the paper is complete, and treat it as a compressed version of the paper's argument rather than a table of contents. The abstract should be self-contained. A reader who reads nothing else should understand what question was asked, how it was answered, what was found, and why it matters.

◆ The four components of an economics abstract

A well-constructed abstract for an economics paper contains four components, in order: the question, the method, the finding, and the implication. Each component is typically one to two sentences. The total length for most journals is 100–200 words; check the target journal's author guidelines, as some impose strict word limits that differ from this range.

The question states what the paper investigates and positions it in the relevant debate. It is not a topic announcement ('This paper studies housing markets') but a specific question with content ('This paper asks whether short-term rental platforms reduce the welfare of long-term residents by tightening the housing supply'). The difference is that a specific question tells the reader immediately what the paper can and cannot answer and what kind of evidence it will use.

The method states, in compressed form, the data, the identification strategy, and the context. For empirical papers, the method sentence is typically two elements: what data are used and what variation identifies the effect of interest. For theoretical papers, it states the model and its key assumptions. The method sentence must be specific enough to allow the reader to evaluate the paper's credibility — 'we use panel data' is insufficient; 'we use a difference-in-differences design exploiting staggered entry of Airbnb listings across US cities' is informative.

The finding states the paper's main result, with its magnitude. This is the most frequently underwritten component of economics abstracts. A finding stated as 'we find that X affects Y' without a magnitude, a direction, or a scope condition is not a finding that a reader can evaluate. Chetty et al. (2014) provide the appropriate model: the abstract of their paper gives the full range of the mobility measure (35.8 to 46.2 percentile points), names the endpoints of the distribution (Charlotte and Salt Lake City), and provides a comparator (larger than the US–Denmark gap in cross-country comparisons). In two sentences the reader knows not just that mobility varies but how much it varies and what that variation means relative to a familiar benchmark.

The implication states what follows from the finding — for the literature, for policy, or for future research. The implication should be calibrated to the evidence: a correlational paper

should not imply policy recommendations that require a causal effect; a single-country study should not imply universal generalisability. The discipline of the implication sentence is the discipline of the whole paper in miniature: state what the evidence supports and no more.

◆ **Structured and unstructured abstracts**

Some journals — particularly those in health economics, development economics, and interdisciplinary outlets — require a structured abstract with explicit headings such as Background, Objective, Methods, Results, and Conclusion. Structured abstracts are easier to write correctly because the required components are labelled, but they are harder to write well: the label 'Results' does not help a writer who does not know what result to state or at what level of specificity. The four-component logic above applies regardless of whether the abstract is structured or unstructured; the structure simply makes the components visible.

◆ **Common abstract failures in economics manuscripts**

Four failures recur across submitted manuscripts. The first is the method-heavy abstract that describes the data and estimation approach at length while stating the finding in a single vague sentence. This abstract tells the reader how the paper works but not what it shows. The second is the contribution abstract that describes the paper's place in the literature ('this paper extends the model of X by incorporating Y') without stating any empirical or theoretical result. The reader learns that the paper exists but not what it finds.

The third is the over-hedged abstract that qualifies every finding so heavily that nothing is asserted: 'results may suggest that, under certain conditions, there could be a possible relationship between X and Y.' A reader who encounters this abstract learns that the authors are uncertain about their own findings, which is not a basis for proceeding to the full paper. The fourth is the abstract written before the paper is finished — which produces a description of what the authors planned to do rather than what they actually found. The discipline of writing the abstract last prevents this failure.

2.10 The Acknowledgements Section

The acknowledgements section appears in the published version of a paper immediately after the conclusion and before the reference list, though in manuscript submissions it is typically placed on the title page or in the paper's front matter. It is short — rarely more than a paragraph — and it performs a specific set of functions that are distinct from both the body of the paper and the reference list. Understanding what belongs in acknowledgements, and what does not, prevents both the under-acknowledgement that looks careless and the over-acknowledgement that looks unprofessional.

◆ **What acknowledgements must contain**

Four categories of acknowledgement are required by most journals and expected by the research community even where not formally required.

Funding sources. Any grant, fellowship, research programme, or institutional funding that supported the research must be named, with the grant number where applicable. Journals increasingly enforce this requirement as a condition of publication, and funding bodies require it for audit purposes. The standard form is: 'This research was supported by [Funder Name] under grant number [N].' Where no external funding was received, some journals require an explicit statement to that effect: 'This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.'

Data and institutional access. Where the analysis depends on access to proprietary, administrative, or restricted data — as in Chetty et al. (2014), whose use of federal tax records required institutional approval — the acknowledgements section names the institution or body that provided access and states any relevant conditions. If the data provider reviewed the paper for accuracy, that review role should be stated and bounded explicitly, as Farronato and Fradkin (2022) did for Airbnb: acknowledging that the company verified the paper for confidential information accuracy, while making clear that editorial control remained with the authors.

Conflicts of interest. Any financial relationship between the authors and the subject matter of the paper — employment, equity holding, advisory roles, or paid consulting — must be disclosed. Journals differ in how they collect this information (some use separate declaration forms), but where a conflict exists, disclosure in the acknowledgements is the minimum standard. Farronato and Fradkin's disclosure that one author held stock in Airbnb is the appropriate model: brief, specific, and placed where readers expect to find it. Undisclosed conflicts of interest are a serious breach of academic integrity (see Chapter Six).

AI writing tool use. Where AI tools have been used in the preparation of the paper — for language editing, initial drafting, literature summarisation, or any other function — this use should be disclosed in the acknowledgements with a brief description of the tool and the function it performed. This norm is now stated in the author guidelines of a growing number of journals and reflects the broader principle of scholarly transparency described in Chapter Six.

◆ **What acknowledgements should contain**

Beyond the required categories, the acknowledgements section conventionally thanks individuals who contributed to the paper without qualifying as authors. The standard is that acknowledgements are owed to people who provided substantive intellectual input — comments on drafts, advice on methodology, data or code assistance, or research assistance — but whose contribution did not rise to the level of authorship. The distinction is not always sharp, but the authorship criteria used by most journals (see Chapter Six, section 6.5) provide the relevant threshold.

Specific things to acknowledge: named seminar audiences and conference presentations where the paper received substantive feedback; colleagues who provided detailed comments on drafts; research assistants who contributed to data collection, cleaning, or analysis; and the editor and referees, whose comments improved the paper, where this is accurate. The last

convention — thanking referees — is standard in acknowledgements accompanying final accepted manuscripts, even though referees are anonymous. It is a professional courtesy, not an empty formality.

◆ **What acknowledgements should not contain**

Acknowledgements should not include: people who provided only general encouragement or moral support (family members, supervisors who did not read the paper, departmental colleagues who attended a seminar but offered no substantive comments); promotional language about the research programme or institution; or lengthy explanations of the paper's development history that belong, if anywhere, in the covering letter.

Length is a discipline. A well-written acknowledgements section for a single-authored paper is typically two to four sentences. For a co-authored paper with multiple funding sources, it may run to a short paragraph. Writing more than this usually indicates that material belonging in the covering letter or the paper's methodological notes has been placed in the acknowledgements instead.

◆ **Acknowledgements in the submission process**

In double-blind peer review — the standard at most leading economics journals — the acknowledgements section must be removed from the manuscript submitted for review, because it contains identifying information (funding sources, named colleagues, institutional affiliations). The standard practice is to include a blinded manuscript with all acknowledgements removed, and a separate title page with the full author details and acknowledgements, submitted alongside but not embedded in the main text. Check the target journal's submission requirements before preparing the manuscript: requirements vary, and non-compliance can lead to desk rejection.

2.11 Responding to Referee Reports

The referee report is not the end of the publication process. For most papers that eventually appear in print, it is the beginning of the most intensive phase of writing. Knowing how to read referee reports, how to prioritise the required revisions, and how to write the revision letter are skills as important as writing the original paper — and they receive remarkably little systematic attention in discussions of academic writing development.

◆ **Reading referee reports analytically**

The first task on receiving a referee report — particularly one requesting major revision — is to resist the natural impulse to respond immediately, either defensively or with wholesale agreement. Referee reports should be read analytically. The writer must distinguish between objections that identify genuine weaknesses in the paper, objections that reflect the referee's

own methodological preferences or theoretical commitments, and requests for additional work that are not strictly necessary to the paper's central argument.

Not all referee comments carry the same weight, and not all of them need to be accommodated. The table below offers a framework for categorising referee comments and choosing the appropriate response.

Referee comment type	What it usually signals	Appropriate response
Methodological objection	A genuine weakness in the identification strategy, data, or model specification	Address directly; provide additional analysis, robustness checks, or revised specification as required
Request for additional literature	Referee believes important prior work has been omitted	Check whether the omitted work is genuinely relevant; cite and engage with it if so; explain politely if not
Scope expansion request	Referee wants more than the paper claims to deliver	Push back diplomatically; explain that the paper's scope is deliberate and that expansion would require a different paper
Stylistic or presentational comment	Prose, structure, or formatting does not meet the referee's preferences	Accommodate where it improves the paper; note where conventions differ
Contradictory referee comments	Two referees disagree about what the paper should do	Explain the contradiction transparently in the revision letter; take a reasoned position and justify it
Request for data or code	Referee wants to verify results or test robustness	Provide all requested materials; treat this as a mark of seriousness, not a threat

◆ The revision letter as a scholarly document

The revision letter — the document submitted alongside a revised manuscript — is one of the most important pieces of writing associated with academic publication, and one of the most often handled badly. Its function is to show the editor and referees that the authors have read the reports carefully, have taken the substantive objections seriously, and have revised the paper in ways that address them.

A strong revision letter is organised around the referee's comments in sequence. It quotes each comment directly before responding to it. It explains clearly what change has been made to the paper and where. And where a requested change has not been made, it provides a clear explanation of why the paper is stronger without it. This last element is where many writers fail: the impulse to accommodate every comment, regardless of merit, produces a paper that has been revised incoherently in response to contradictory instructions. Measured, respectful pushback — based on argument rather than ego — is a legitimate and often necessary part of the revision process.

⚠ Five structural problems that reliably generate rejection or major revision

1. Introduction that contextualises without establishing a gap — the paper never explains what question the existing literature has left unanswered.
2. Literature review that lists rather than synthesises — sources are summarised individually without integration or argumentative direction.
3. Methodology section that describes without justifying — the paper explains what was done but not why those choices were appropriate for the research question.
4. Discussion section that restates rather than interprets — the paper reports what was found a second time without explaining what it means or why it matters.
5. Conclusion that retreats to the introduction — the paper ends with broad contextual framing rather than a direct answer to its own research question.

2.12 Vocabulary and Sentence Models

The terms below are used throughout the research paper genre. Learning them as active writing tools — not just as concepts to recognise — allows writers to move between sections confidently and to meet the genre expectations that editors and referees apply. The sentence models that follow cover the key moves in a research paper from introduction to conclusion.

◆ Key vocabulary

Term	Meaning in economics writing	Example of the word in a sentence
gap	A specific question that the existing literature has not answered, or a problem it has not adequately addressed. The gap is what justifies the paper's existence. It must be stated precisely — not as 'little is known about X' but as 'prior work has not established whether X holds when condition Y is present'.	<i>'The literature on peer-to-peer markets has established that entry affects incumbent prices, but has not determined whether this effect operates through quality competition or price competition alone — the gap this paper addresses.'</i>
contribution	What the paper adds to existing knowledge. The contribution must be specific (not 'this paper studies X') and genuinely new. It is stated in the introduction and restated, with supporting evidence, in the conclusion.	<i>'The paper's contribution is to show that the welfare effects of platform entry depend critically on the supply elasticity of incumbents — a mechanism not identified in prior work on platform competition.'</i>
research question	The specific question the paper answers, stated in terms precise enough to be answered by the paper's method and data. A research question is not a topic ('this paper examines X') and not a hypothesis ('this paper tests whether X causes Y') but a question that the analysis is designed to resolve.	<i>'The central question is: do platforms that lower barriers to market entry improve consumer welfare when incumbents can adjust quality as well as price?'</i>
identification	The strategy by which a paper establishes that its estimated relationship reflects a causal (or otherwise meaningful) effect rather than a spurious correlation. Identification requires a source of variation in the data that is plausibly uncontaminated by selection or confounding.	<i>'Identification relies on the staggered rollout of Airbnb across US cities, which creates variation in platform exposure that is not driven by pre-existing trends in hotel occupancy.'</i>

specification	The mathematical and statistical form of the model estimated in the paper: which variables are included, how they are measured, what functional form is used. A paper must justify its specification, not merely report it.	<i>'The baseline specification regresses log hotel revenue on a dummy for Airbnb entry, controlling for city fixed effects, time trends, and pre-entry city characteristics.'</i>
robustness	The degree to which results hold under variations in the specification, sample, or measurement choices. Reporting robustness checks is a way of showing that the main findings do not depend on a specific, potentially arbitrary analytical choice.	<i>'The results are robust to alternative measures of Airbnb penetration, to restricting the sample to cities above the median hotel market size, and to including state-by-year fixed effects.'</i>
peer review	The process by which submitted manuscripts are evaluated by expert referees before publication. Understanding peer review — what referees look for, how referee reports are structured, and how revision responses should be written — is part of the practical knowledge needed to publish.	<i>'The referee's central concern was that the identification strategy did not adequately address the possibility that Airbnb entered markets that were already declining — a concern addressed in the revision by including pre-trend tests.'</i>
abstract	A short summary of the paper (typically 100–200 words) that states the research question, the method, the main finding, and the implication. The abstract is the most widely read part of any paper and must be self-contained.	<i>'The abstract must stand alone: a reader who reads nothing else should understand what question was asked, how it was answered, what was found, and why it matters.'</i>

◆ Sentence models

The sentences below model the key moves across the sections of a research paper. Parenthetical cues show the section where each move typically appears.

Function (section)	Model sentence	Notes for the writer
Establishing context (Introduction)	<i>[Topic/phenomenon] has attracted substantial attention in recent years, in part because [reason why it matters economically].</i>	Open with context, not with the paper's findings. Give the reader a reason to care before making claims.
Identifying the gap (Introduction)	<i>Despite this body of work, [prior studies] have not established [specific unanswered question], because [reason — data limitation, methodological constraint, or theoretical gap].</i>	Name the prior work, name the gap, explain why it exists. Vague gaps ('little is known') are not adequate.
Stating the contribution (Introduction)	<i>This paper addresses [gap] by [method/approach], using [data], and finds that [main result].</i>	One sentence that contains all four elements: gap, method, data, result. This sentence is the paper in miniature.
Reviewing prior work (Literature review within paper)	<i>[Author(s), year] [showed/argued/estimated] that [finding]. However, [limitation of that work], which [our paper addresses/extends].</i>	Cite prior work precisely — give the specific finding, not just the topic. Then state the limitation that creates space for your contribution.
Justifying the method (Methodology)	<i>[Method] is appropriate for this research question because [reason relating to the identifying variation/data structure/research design].</i>	Every methodological choice requires a justification. The justification links the method to the specific question being asked.
Stating a result (Results)	<i>[Variable X] is associated with [variable Y] [direction and magnitude] (Table [N], Column [N]). [Explanation of what this means].</i>	State the finding in the prose — don't just point to the table. The table provides evidence; the prose states the conclusion.

Interpreting a result (Discussion)	<i>This finding is consistent with [theoretical mechanism], and suggests that [implication]. An alternative explanation — [alternative] — is addressed in [section/robustness check].</i>	The discussion section interprets; it does not re-report. Anticipate the alternative explanation and show why the preferred interpretation is more supported.
Stating the conclusion (Conclusion)	<i>This paper has shown that [main finding], using [identification strategy/data]. The finding implies that [implication for policy or theory]. Future work should address [limitation/extension].</i>	The conclusion answers the research question directly. It should not introduce new information — only the earned answer and its implications.
Responding to a referee (Revision letter)	<i>We thank the referee for this observation. We have [action taken: added/revised/extended] in [section], where we now show [what was added]. We believe this addresses the concern that [referee's specific concern].</i>	Be specific in revision responses. Name the section changed, describe what was added, and connect it explicitly to the referee's stated concern.

The research paper is the discipline's primary instrument of argument. Each of its sections performs a specific intellectual function, and the paper's strength as a whole depends on each section performing that function rather than defaulting to description, restatement, or evasion. Friedman's 1968 paper remains instructive not because its conclusions are unchallengeable — they have been extensively contested — but because its writing is equal to its argument at every level: from the opening context to the earned conclusion. That is the standard the research paper sets for itself, and the standard that editors and referees are, in effect, measuring submitted work against.

→ **Chapter Three: The Literature Review as a Standalone Publication**

CHAPTER THREE

The Literature Review as a Standalone Publication

A note on the model text for this chapter. Chapters One and Two drew on Milton Friedman's 'The Role of Monetary Policy' (Friedman, 1968) as their primary example — a paper chosen for the precision of its argument architecture and the clarity of its prose. From this chapter onwards, the guide's primary model text shifts to a more recent work: Goldfarb and Tucker's 'Digital Economics' (Goldfarb & Tucker, 2019), published in the *Journal of Economic Literature*. The reasons for this choice are discussed below, but the most important is straightforward: 'Digital Economics' is itself a standalone literature review, and no more direct model for this chapter's subject matter is available.

3.1 The Model Text: Goldfarb and Tucker (2019)

Goldfarb, A., and Tucker, C. (2019). Digital Economics. *Journal of Economic Literature*, 57(1), 3–43.

'Digital Economics' reviews a large and fast-growing literature on how digital technology has changed economic activity. Its central contribution is a five-cost framework: the authors argue that the economic importance of digitalisation is best understood through the reductions it has produced in five distinct cost categories — search costs, replication costs, transportation costs, tracking costs, and verification costs. Each cost reduction draws on a different body of economic theory, generates a different set of empirical questions, and has produced a different research trajectory. By organising the literature around this framework — rather than around chronology, subfield, or individual paper — Goldfarb and Tucker make a contribution that is genuinely analytical: the framework itself is the intellectual added value of the review.

The paper is a good primary model for this chapter for several reasons. It is recent enough to reflect contemporary standards of academic writing in economics. It is published in one of the discipline's most prestigious outlets for review articles. It addresses a field — digital markets and platforms — that is directly relevant to contemporary economics research and to the marketing and management contexts in which many faculty members work. And it is written with a structural discipline and clarity of purpose that make its writing choices easy to analyse and learn from. The five-cost framework, the handling of contested evidence, the management of a large and varied literature, the positioning of the contribution — all are handled in ways that illustrate the principles this chapter addresses.

Throughout this chapter, passages from Goldfarb and Tucker (2019) are quoted directly and analysed. Readers are strongly encouraged to obtain and read the full paper alongside this chapter; it is openly available via the NBER working paper series (Working Paper 23684) and through institutional access to the *Journal of Economic Literature*.

3.2 The Literature Review as a Scholarly Contribution

A standalone literature review is publishable when it makes a contribution that individual readers cannot make for themselves by consulting the primary sources. That contribution may take several forms: it may identify patterns of consensus and dissent that are not visible from within any single sub-debate; it may trace the methodological development of a field; it may identify a persistent gap or blind spot that individual papers have not been able to see; or it may provide an organising framework that makes a dispersed literature coherent and navigable. What a literature review cannot do — at publication level — is simply describe what a set of papers found. A review organised as a sequence of summaries, however accurate and comprehensive, has not synthesised; it has catalogued.

This distinction between catalogue and synthesis is exactly what separates a publishable literature review from a useful PhD literature chapter. Goldfarb and Tucker make the distinction explicit in their opening pages. They do not announce that they will survey the literature on digital economics; they announce that they will organise it around a specific analytical framework. The five-cost structure is stated in the abstract and restated at the opening of the paper's first substantive section:

Digital technology is the representation of information in bits. This technology has reduced the cost of storage, computation, and transmission of data. Research on digital economics examines whether and how digital technology changes economic activity. In this review, we emphasise the reduction in five distinct economic costs associated with digital economic activity: search costs, replication costs, transportation costs, tracking costs, and verification costs.

Goldfarb, A. and Tucker, C. (2019). Digital Economics. *Journal of Economic Literature*, 57(1), p. 3.

The organising framework as the review's primary contribution

Three features of this opening warrant attention. First, the definition — 'Digital technology is the representation of information in bits' — is not contextual padding; it is an analytical move. By defining digital technology in terms of its fundamental property (bitwise representation), Goldfarb and Tucker anchor their framework in a characteristic that is theoretically tractable: it is because bits can be stored, copied, and transmitted at near-zero marginal cost that the five cost reductions follow. The definition does intellectual work.

Second, the five costs are announced immediately and specifically. The reader does not have to wait to discover the organising principle of the review; it is presented in the abstract and confirmed in the first paragraph of the text. This is a deliberate structural choice that signals, from the outset, that the review has an argument — not merely a coverage plan.

Third, the phrase 'we emphasise' is doing important honest work. Goldfarb and Tucker are not claiming that these five costs are the only way to organise the literature, or that their framework is exhaustive. They are stating that it is the framework they have chosen, and implicitly inviting the reader to evaluate it. This is intellectual transparency about scope — a virtue that reviewers and editors recognise and value.

◆ The distinction between systematic and narrative reviews

In economics and adjacent social sciences, standalone literature reviews typically take one of two forms. A systematic review follows an explicit, pre-specified protocol for identifying, selecting, and assessing studies — databases searched, search terms used, inclusion and exclusion criteria — so that the review is, in principle, replicable. Systematic reviews are most common in economics subfields with close connections to public health and policy evaluation, and they are the dominant form in meta-analytic work.

A narrative review does not follow a pre-specified protocol. It is selective in ways that reflect the author's scholarly judgement about what is most important and relevant. The best narrative reviews are more intellectually powerful than systematic ones precisely because they are organised around an argument rather than a procedure. 'Digital Economics' is a narrative review of the highest quality: its coverage is extensive, but the selection and organisation of material is governed throughout by the five-cost framework rather than by a bibliographic procedure. For most faculty writing in economics, the narrative review is the appropriate form — provided that selectivity does not shade into cherry-picking, and that the framework chosen is genuinely useful rather than merely convenient.

3.3 Scoping and Selecting the Literature

Before a literature review can be written, two prior decisions must be made: what falls within its scope, and how sources will be selected and evaluated. These are not just practical decisions. They are intellectual commitments that shape the review's argument and its credibility.

◆ Defining the scope through a question, not a topic

A scope defined by topic — 'the economics of digital markets' — provides no principle for deciding what to include or exclude; every paper that mentions digital markets is potentially in scope. A scope defined by question — 'how do reductions in five specific cost categories change economic activity in digital markets?' — provides a clear selection criterion: a paper is in scope if it bears on one of the five cost reductions and their consequences. This is precisely what Goldfarb and Tucker do, and the result is a review that is simultaneously broad (covering search, pricing, copyright, privacy, reputation, and platform economics) and coherent (because every section addresses the same underlying question through a different cost lens).

The most common scoping failure is treating topic and question as the same thing. Writers who scope a review by topic tend to produce reviews that are extensive but unfocused — they include everything that might be relevant and exclude nothing — and the result is a catalogue rather than an argument. The discipline required is to state the question the review is answering and to use that question as the selection criterion throughout.

◆ Evaluating sources: epistemic standing

Not all sources carry the same evidential weight in economics, and a literature review that treats a working paper, a policy report, and a published article in the *American Economic Review* as interchangeable misrepresents the evidential landscape of the field. The table below sets out the principal source types encountered in economics literature reviews, their epistemic standing, and the appropriate use of each.

Source type	Epistemic standing in economics	Appropriate use in a literature review
Peer-reviewed journal article	Highest — externally refereed, fully published	Primary source for establishing the state of knowledge; cite for empirical findings, theoretical results, and established debates
Published working paper (NBER, CEPR, SSRN)	High — widely circulated, often pre-publication versions of eventual journal articles	Cite for recent contributions not yet in print; note working paper status; check for a published version before submission

Unpublished working paper	Moderate — not externally refereed; may change significantly	Use selectively for very recent or specifically relevant work; flag unpublished status explicitly
Policy report (IMF, World Bank, OECD)	Moderate — professionally reviewed but not peer-refereed	Appropriate for contextual data, policy framing, and applied evidence; do not treat as equivalent to refereed research
Edited volume chapter	Variable — quality depends on volume and editors	Use where the chapter represents a scholar's definitive statement on a topic; note that editorial standards vary considerably
Textbook	Established consensus — not original research	Appropriate for definitional or pedagogical claims; do not cite as evidence for contested empirical or theoretical positions
Newspaper / magazine article	Low — not subject to academic review	Appropriate only for documenting policy events, public debates, or context; never as evidence for academic claims

Goldfarb and Tucker illustrate appropriate source calibration throughout 'Digital Economics'. Their primary evidence base is peer-reviewed journal articles; they draw on working papers for recent contributions not yet in print; and they reference policy documents and industry reports only for contextual data, never as evidence for theoretical or causal claims. The distinction is maintained consistently and signals scholarly judgement about the hierarchy of evidence.

3.4 Synthesising Rather Than Summarising

The instruction to synthesise rather than summarise is the most frequently stated principle of literature review writing — and also the most frequently ignored in practice. It is worth being precise about what the distinction actually requires, because the work of synthesis is considerably harder than the word suggests.

Summarising requires the writer to represent accurately what each source says. Synthesising requires the writer to identify what the sources, taken together, reveal: where they converge and diverge, what patterns emerge across findings, how methodological choices have produced different results, and what cumulative picture the field has — and has not — drawn of the phenomenon under study. Synthesis is not the aggregation of summaries. It is an independent analytical act performed on the body of literature as a whole.

Goldfarb and Tucker's treatment of the search costs literature illustrates this distinction with clarity. Rather than summarising individual papers on online price dispersion, they identify an intellectual puzzle that cuts across the literature — a puzzle that individual papers could not have identified from their own perspective:

Early research on search costs tested straightforward models of lower costs. For example, the search literature of the late 1990s and early 2000s built directly on earlier models by Diamond (1971) and Varian (1980). As we detail below, empirical work emerged that found some inconsistencies with the simple models, and so richer models and empirical analysis of the cost reductions developed.

The early empirical literature found that price dispersion online was often greater than price dispersion offline. This was puzzling given the prediction of lower search costs that online comparison shopping would reduce price dispersion.

Goldfarb, A. and Tucker, C. (2019). Digital Economics. *Journal of Economic Literature*, 57(1), p. 4.

Synthesis as the identification of a cross-paper puzzle

The phrase 'this was puzzling' is the pivot of a synthesis move. Individual papers on online price dispersion each reported their own findings; Goldfarb and Tucker, looking across the literature, identify that those findings, taken together, contradict the theoretical prediction that motivated the research programme. No individual paper could make this observation about its own place in the literature; it requires the vantage point of the reviewer.

The structure of the passage is also worth noting: simple model → early empirical test → inconsistency → richer model. This is a chronological account, but it is not organised chronologically for its own sake; it is organised to trace the intellectual logic of the field's development in response to a specific anomaly. The sequence has an argument — the field learned something that forced it to revise its models — and the chronology serves that argument.

This is what it means to treat the literature as a resource for argument rather than a record to be transcribed. The writer is not reporting what papers found; they are tracing the intellectual consequences of what papers found — which is a fundamentally different, and more valuable, activity.

◆ The five synthesis moves

The table below identifies five core synthesis moves that a well-constructed literature review uses in combination. They are not a sequence but a description of the analytical work that synthesis requires. A writer reviewing their draft should be able to identify which moves are present and which are absent.

Synthesis move	What it does	Useful signal phrases
Establishing consensus	Identifies what the literature has collectively established with confidence	<i>'There is broad agreement that...'; 'A consistent finding across studies is...'; 'The literature has established...'</i>

Identifying methodological division	Notes that different methods or data sources produce different findings	<i>'Studies using X tend to find...; whilst those using Y report...'; 'The divergence in findings reflects, in part, methodological differences in...'</i>
Mapping a contested claim	Acknowledges genuine scholarly disagreement without taking sides prematurely	<i>'Views on X remain divided...'; 'The evidence is mixed: [Author A] finds..., whilst [Author B] argues...'; 'This debate has not been resolved...'</i>
Locating the gap	Identifies what the literature has not yet addressed, or has addressed inadequately	<i>'What remains unexamined is...'; 'No study has yet addressed...'; 'Existing work has not distinguished between...'</i>
Positioning the paper	States how the present paper addresses the identified gap	<i>'This paper addresses this gap by...'; 'The present analysis extends [Author]'s framework to...'; 'We contribute to this debate by...'</i>

All five moves appear in 'Digital Economics', used at different scales. The consensus move appears repeatedly within individual sections: Goldfarb and Tucker establish, for instance, that lower search costs have generally reduced price dispersion — though with important qualifications. The methodological division move is used to explain the price dispersion puzzle: studies finding greater online dispersion tend to use different product categories and measurement approaches than those finding lower dispersion. The contested claim move appears in the tracking costs section, where the economic consequences of personalisation and privacy remain genuinely unresolved. The gap move closes each section. And the positioning move is the purpose of the five-cost framework itself — it positions the entire review as a contribution to making a dispersed and rapidly growing field analytically coherent.

3.5 Organising the Review

Once the scope has been defined and the synthesis moves identified, the writer faces an organisational decision that will shape the review's argumentative structure. There is no universally correct answer, but there are better and worse choices depending on the subject matter and the argumentative purpose.

◆ **Framework-led organisation**

Goldfarb and Tucker's choice of framework-led organisation — each of the five cost reductions provides the heading and subject matter of a major section — is the most direct route from an organising framework to an organised text. The framework does the work that a thematic or chronological organising principle would otherwise have to do. It also makes the review's intellectual contribution explicit in its very structure: the reader encounters the five-cost argument not only in the introduction and conclusion but in the shape of the paper itself.

Framework-led organisation is appropriate when the framework is the contribution. If the review's added value is that it proposes a new way of organising a dispersed literature, then organising the review around that framework demonstrates the framework's usefulness at the same time as it makes the scholarly argument. This is an efficient rhetorical move, and it is one of the reasons 'Digital Economics' is as effective as it is.

◆ **Thematic, chronological, and methodological alternatives**

Not every review has an original organising framework to offer. For those that do not, three alternative principles are available: thematic organisation (grouping the literature by the substantive questions or debates it addresses), chronological organisation (tracing the development of a field over time), and methodological organisation (grouping by the analytical approaches used). Each is appropriate in different circumstances.

Thematic organisation is the most common and usually the most effective for a review whose argument concerns the current state of knowledge on a question. Chronological organisation is most appropriate when the review's argument is about intellectual change — about how a field has developed and what it has learned. Methodological organisation is appropriate when methodological divisions are a primary source of disagreement in the field and the review's contribution is to map those divisions clearly.

The failure common to all three is the same: the organisational principle becomes an end in itself rather than a tool for argument. A thematic review that moves from theme to theme without integrating them; a chronological review that narrates a history without interpreting it; a methodological review that catalogues approaches without assessing their relative contributions — all are forms of the same failure to keep the central question in view.

◆ **Transitions that carry the argument**

Whatever the organisational principle, transitions between sections must carry the review's argument forward rather than just announce a change of topic. A transition that says 'Having examined search costs, we now turn to replication costs' has described a sequence

without explaining a logic. An argumentative transition explains why the second section follows from the first — what new question the first section raises, what limitation it reveals, or what the two sections together establish that neither establishes alone.

✗ Transitional announcement	✓ Argumentative transition
<i>Having reviewed the evidence on search costs, this section now examines replication costs.</i>	<i>The reduction in search costs has changed how consumers find and compare products; the reduction in replication costs changes the economics of the products themselves. Where search costs affect market structure through their influence on competition, replication costs affect it through their influence on the nature of the goods being traded — and the two effects interact in ways the following section begins to trace.</i>
<i>Section 3 discussed the early literature on price dispersion. Section 4 now turns to more recent work on digital platforms.</i>	<i>The price dispersion literature established that simple models of lower search costs were insufficient — the data did not behave as predicted. The platform literature, reviewed below, represents in part the field's response to that puzzle: if lower search costs do not straightforwardly reduce prices, how do digital markets actually organise competition, and what role do platforms play in that organisation?</i>

3.6 Writing the Review: Attribution, Voice, and Hedging

The prose of a literature review must do two things at once: report accurately what others have claimed, and maintain an independent analytical voice that shapes the review's argument. Writers who focus only on the first function produce reviews that are accurate but voiceless. Writers who focus only on the second produce reviews whose independence shades into distortion. The discipline required is to do both simultaneously: represent the literature faithfully whilst using it to support an argument.

◆ Attribution registers

The language used to introduce citations is not just a stylistic choice. It signals the writer's relationship to the cited claim. Different registers convey different degrees of endorsement, critical distance, and evaluative judgement. The table below sets out the principal attribution registers used in economics literature reviews, with examples drawn from the digital economics literature.

Attribution register	Example	When to use
Neutral reporting	<i>'Smith (2018) finds that...' / 'According to Jones (2020)...</i>	Default register for summarising others' findings without evaluation
Evaluative endorsement	<i>'In a particularly influential study, Goldfarb and Tucker (2019) demonstrate...'</i>	When the cited work is a field-defining contribution that warrants explicit recognition
Critical distance	<i>'Smith (2018) claims that...' / 'Smith (2018) argues, controversially, that...'</i>	When the cited position is disputed or when you will subsequently challenge it
Methodological note	<i>'Using a panel of US internet adoption data, Smith (2018) finds...'</i>	When the method is directly relevant to evaluating the applicability of the finding
Explicit disagreement	<i>'Smith (2018) finds X; however, the present analysis produces a different result, for reasons addressed in Section 4.'</i>	When your paper's findings contradict an established result; must be handled with precision and evidence

Goldfarb and Tucker use attribution registers with precision throughout 'Digital Economics'. When citing well-established theoretical results — Diamond (1971) on search, Varian (1980) on price dispersion — they use neutral reporting and allow the date and journal to signal the established standing of the work. When introducing findings that the discussion will later complicate or qualify, they use registers that prepare the reader to evaluate the finding in context. The register is never arbitrary. It always reflects the writers' assessment of the cited work's evidential status relative to the review's argument.

◆ Maintaining analytical voice

The risk of a heavily cited review is that the writer's voice disappears behind a succession of attributed claims. 'Digital Economics' avoids this through a consistent analytical stance: every section opens with a statement of what the cost reduction in question is and why it matters economically, before any individual paper is cited. The framework provides the voice; the citations provide the evidence. The relationship is clear, and the reader is never in doubt about whose argument they are following.

Goldfarb and Tucker's handling of the tracking costs section — one of the most contested areas of the literature — illustrates how to maintain voice whilst representing genuine disagreement:

Reductions in tracking costs enable personalisation and the creation of one-to-one markets, leading to renewed interest in established economic models with asymmetric information and differentiated products such as price discrimination, auctions, and advertising models. The economics literature on privacy, both offline and online, grapples with the question of how privacy should be treated in terms of the consumers' utility function.

Privacy regulation puts a cost on tracking information flows. As such, the precise nature of privacy protection can be expected to matter a lot for the direction of innovation; it is not a matter of a simple binary choice to have privacy protection or not.

Goldfarb, A. and Tucker, C. (2019). Digital Economics. Journal of Economic Literature, 57(1), p. 22.

Maintaining voice in contested territory

The phrase 'grapples with the question' is an attribution register that signals genuine unresolved debate without abandoning analytical direction. Goldfarb and Tucker do not pretend the privacy question is settled, but they also do not simply list the competing views and leave the reader to weigh them. The second paragraph performs the analytical move: the important variable is not whether privacy protection exists, but what form it takes — 'it is not a matter of a simple binary choice.' This is the writers' own claim, grounded in the literature but not reducible to any single cited paper.

Notice the sentence construction: 'can be expected to matter a lot'. The hedge ('can be expected') reflects genuine uncertainty; the emphasis ('matter a lot') reflects genuine analytical conviction. This calibration — uncertain about the direction of effects, confident about their significance — is precisely the kind of differentiated hedging that marks a sophisticated literature review. It is neither over-cautious nor over-assertive; it is proportionate to the state of the evidence.

◆ Hedging in proportion to the evidence

The hedging principles established in Chapter One apply with particular force in literature reviews, because the review's claims are collective claims about what multiple studies have jointly established. A claim attributed to 'the literature' carries more weight than a claim attributed to a single study, and must therefore be made with care. The practical rule is proportionality: the strength of the claim should match the robustness of the consensus it describes.

Goldfarb and Tucker apply this proportionality consistently. In areas where the evidence is consistent across methods and data sources — the positive relationship between digitalisation and product variety, for instance — they state findings directly without excessive qualification. In areas where the evidence is mixed or contested — the welfare effects of price discrimination enabled by tracking costs — they explicitly note the disagreement and explain what drives it. The level of hedging is not uniform; it varies by section according to the strength of the evidence, and this variation is itself analytically informative.

3.7 Positioning the Paper's Contribution

Whether it appears as a standalone publication or as an embedded section within a research paper, the literature review must eventually position the contribution of the present work. This is the moment at which the review's intellectual work pays off: the gap,

contradiction, or unresolved debate identified through the synthesis is connected to what the paper or review itself offers.

◆ Making the gap specific

The gap must be stated in language that is specific enough to be falsifiable. 'This area has not been fully explored' is not a gap statement; it is an observation that the literature is finite. A genuine gap statement identifies a specific claim, mechanism, or relationship that the existing literature has not established — and explains why: because the relevant data were not available, because the methods used were not suited to the question, or because the literature has been built on assumptions that prevented the relevant question from being asked.

Goldfarb and Tucker's gap statement is structural rather than empirical: they argue that the existing literature on digital economics, whilst extensive and technically sophisticated, lacks an organising framework that makes its contributions mutually intelligible and cumulative. The gap is not a missing empirical study but a missing conceptual structure. This is a more ambitious kind of gap — and a more challenging one to justify — but it is precisely what makes the five-cost framework a genuine contribution rather than a convenient taxonomy.

The paper states this gap with directness and without excessive qualification in the introduction:

Each of these cost changes draws on a different set of well-established economic models, primarily search, non-rival goods, transportation cost, price discrimination, and reputation models. Early research tested straightforward models of lower costs. As we detail below, empirical work emerged that found some inconsistencies with the simple models, and so richer models and empirical analysis of the cost reductions developed.

Goldfarb, A. and Tucker, C. (2019). Digital Economics. *Journal of Economic Literature*, 57(1), p. 4.

The contribution positioned through the framework

This passage positions the review's contribution not by identifying what has not been studied but by showing how existing studies relate to each other. The five-cost framework is justified not by the absence of prior work but by the absence of a coherent structure in which that work can be compared, cumulated, and used to identify next steps.

The phrase 'each of these cost changes draws on a different set of well-established economic models' is doing two things at once: it demonstrates that the five-cost framework is analytically grounded (each cost maps onto an existing theoretical tradition), and it explains why the review is necessary (without the mapping, the connections between sub-literatures are invisible to most researchers working within any single one of them).

For writers of standalone literature reviews, this passage suggests a model for positioning a framework-based contribution: show that the framework is grounded (it connects to established theory), show that it is necessary (without it, something important is invisible), and show that it is productive (it generates questions and research directions that the un-organised literature cannot).

Avoiding the two failure modes

Two opposing ways of failing to position a contribution appear regularly enough to deserve explicit attention.

The first is uncritical acceptance: the review describes a settled consensus, identifies no genuine gaps or tensions, and then claims — without adequate foundation — that the present paper makes a contribution to a field that the review itself has presented as complete. A review that presents a settled consensus cannot justify a new paper by reference to that consensus. The positioning move must be grounded in a genuine gap, tension, or limitation in the existing literature.

The second is unfounded dismissal: the review overstates the inadequacy of the existing literature in order to create space for the present paper's contribution. This is intellectually dishonest, and experienced referees — who have often contributed to the literature being dismissed — will respond accordingly. Goldfarb and Tucker's treatment of the existing literature on digital economics is the appropriate model: they acknowledge what the literature has established, are specific about what it has not organised well, and claim only the contribution that the five-cost framework can actually deliver — that is, a clearer and more cumulative way of organising an existing body of knowledge.

3.8 Diagnosing Problems in Literature Reviews

Four structural failures appear regularly enough in submitted literature reviews to deserve explicit identification. Each reflects a specific misconception about what a literature review is supposed to accomplish.

◆ The cataloguing problem

The most widespread failure is substituting catalogue for synthesis: organising the review as a sequence of source summaries, each reporting what the relevant paper found, none integrated into a larger analytical argument. The cataloguing review shows that the writer has read the literature. It does not show that they have understood what the literature collectively means.

The diagnostic sign of cataloguing is a series of topic sentences that begin with author names: 'Smith (2018) finds that...', 'Jones (2020) argues that...', 'Brown and Green (2021) show that...'. A synthesis move would begin differently: 'The evidence on X consistently points towards Y, with the partial exception of Brown and Green (2021), whose different result appears to reflect their use of a different identification strategy.' The contrast between these two constructions is the contrast between a list and an argument.

✗ Catalogue (author-led)	✓ Synthesis (argument-led)
<i>Brown (2015) finds that online platforms reduce search costs for consumers. Jones (2017) also finds reduced search costs. Smith (2019) finds that price dispersion is higher online than offline. Williams (2021) finds that platforms have increased price competition.</i>	<i>The balance of evidence supports the prediction that digital platforms reduce consumer search costs (Brown, 2015; Jones, 2017; Williams, 2021), though this consensus sits in tension with the finding of greater price dispersion online than offline (Smith, 2019) — a puzzle Goldfarb and Tucker (2019) attribute to the heterogeneity of goods and the endogeneity of which products move online first.</i>

◆ The false consensus problem

A literature review that presents a field as more settled than it actually is misrepresents the state of knowledge. This problem arises when writers, seeking to establish a clear foundation for their own contribution, smooth over genuine disagreements in the literature. Goldfarb and Tucker's handling of the privacy literature is instructive here: they do not

minimise the genuine uncertainty about the welfare effects of tracking costs, but they also do not present that uncertainty as a reason not to have a position. They state what is known, acknowledge what is contested, and explain what drives the disagreement. This is the appropriate model: honest about tension, analytical about its sources.

◆ **The missing gap problem**

A review that describes and synthesises the literature without identifying a specific gap that the present paper or a future paper could address has completed the descriptive work without doing the positioning work. The gap does not have to be an empirical lacuna; it may be a conceptual one, as in Goldfarb and Tucker's case — an absence of organising structure that makes the literature less cumulative than it could be. But it must be specific, falsifiable, and connected to the contribution being claimed.

◆ **The missing contribution problem**

A review that identifies the gap but does not state clearly how the present work addresses it has identified a problem without offering a solution. At publication level, the review must close with a clear statement of what follows from the synthesis: what the paper contributes, why that contribution is now possible, and why it matters for the field. This closure is not optional. It is the final move of the review's argument, and without it the review has not completed its intellectual task.

⚠ Four failure modes in literature reviews — a pre-submission checklist

1. Catalogue not synthesis — the review moves from source to source without integrating findings into a larger argument. Fix: restructure around themes, debates, or framework categories, not around individual papers.
2. False consensus — genuine disagreements have been smoothed over. Fix: identify and represent the most significant points of methodological or empirical contention explicitly, and explain what drives them.
3. Missing gap — the review does not identify a specific, falsifiable gap that the paper or a future paper could address. Fix: the concluding section must state what the literature has not established, and why.
4. Missing contribution — the review ends without positioning a specific scholarly contribution. Fix: the final paragraph must state, in clear terms, what follows from the synthesis — what the paper contributes and why it is the right vehicle for that contribution.

3.9 Vocabulary and Sentence Models

The vocabulary below covers terms specific to the literature review genre in economics. The sentence models that follow demonstrate the key moves from scoping and synthesis to gap identification and contribution positioning. All models can be adapted by replacing the bracketed elements with your own content and sources.

◆ Key vocabulary

Term	Meaning in literature review writing	Example of the word in a sentence
synthesis	The act of combining findings from multiple sources into a coherent argument that shows how the pieces relate. Synthesis is the opposite of summary: a summary says what each source found; a synthesis shows what those findings, taken together, establish or contest.	<i>'Synthesising the evidence on digital market pricing, Goldfarb and Tucker (2019) show that the cost structures of digital goods — near-zero replication cost, high fixed cost — are more important in explaining observed price patterns than the competitive dynamics that most earlier work emphasised.'</i>
thematic organisation	Structuring a literature review around topics, debates, or framework categories rather than around individual papers or chronological order. Thematic organisation makes the review's argument visible and helps the reader see how different strands of the literature relate.	<i>'The review is organised thematically: Section 2 addresses measurement, Section 3 addresses identification, and Section 4 synthesises findings on welfare effects — a structure that reflects the logical dependencies in the literature rather than the order in which papers appeared.'</i>
scholarly conversation	The ongoing exchange of arguments, findings, and methods among researchers working on related questions. A literature review positions the paper within this conversation: it shows what others have said and where the paper adds something new.	<i>'This paper enters a scholarly conversation about the welfare effects of platform entry that has been developing since at least Rochet and Tirole (2003), and engages directly with the recent empirical work surveyed in Goldfarb and Tucker (2019).'</i>

methodological debate	A point of disagreement among researchers about how a question should be studied — what data should be used, what method is appropriate, what constitutes adequate identification. Methodological debates are as important to review as empirical debates.	<i>'A key methodological debate in this literature concerns the appropriate unit of analysis: papers using city-level data find positive welfare effects, while those using individual-level data find null or negative effects, suggesting that aggregation may mask important heterogeneity.'</i>
scope	The boundaries of a literature review: the questions it covers, the disciplines it draws from, the time period it addresses. Scope decisions must be justified — a review that is too narrow misses important work; one that is too broad loses analytical focus.	<i>'The review is deliberately scoped to empirical studies of digital platform markets published since 2000, when the data and methods needed to study these markets became available. Theoretical work is discussed only where it directly motivates the empirical questions addressed.'</i>
seminal work	A paper or book that established the foundations of a research area, introduced a key concept or method, or set the terms of a debate that subsequent work has built on. Seminal works require more than a passing citation; they require explanation of their contribution and limitations.	<i>'The seminal work in this area is Diamond (1971), which showed that small search costs can sustain prices far above the competitive equilibrium — a result that digital markets were initially expected to overturn, and which Goldfarb and Tucker (2019) show has proved more durable than predicted.'</i>
contested finding	An empirical result that is disputed in the literature — where different papers, using different data, methods, or sample periods, reach different conclusions. Contested findings must be represented accurately and the sources of disagreement explained.	<i>'The finding that social capital is associated with higher intergenerational mobility is contested: Chetty et al. (2014) report a positive correlation, but subsequent work using panel data shows that the relationship is unstable across cohorts and sensitive to the measure of social capital used.'</i>

positioning	The rhetorical and analytical work of explaining how the present paper relates to the literature reviewed — what it builds on, what it departs from, and what specific gap it fills. Positioning is the final move of a literature review.	<i>'This paper builds on Goldfarb and Tucker's (2019) framework for digital goods but departs from their aggregated analysis by estimating welfare effects at the individual transaction level, which allows the identification of distributional effects that aggregate estimates obscure.'</i>
--------------------	--	--

◆ Sentence models

The models below cover the main moves across a standalone literature review, from the framing of scope to the final positioning of the paper's contribution.

Function	Model sentence	Notes for the writer
Stating the scope	<i>This review covers [topic area] as studied in [disciplines/journals] from [time period]. Studies that [out-of-scope criterion] are not included, because [reason].</i>	State scope decisions explicitly and justify them. Unjustified scope restrictions invite the inference that important inconvenient work was excluded.
Introducing a theme	<i>A central theme in this literature is [theme]: the question of whether [specific debate]. This debate has proceeded along two main lines.</i>	Use thematic headers or signalling sentences to make the review's structure visible. Readers should be able to follow the argument without reading every citation.
Synthesising across sources	<i>[Author A, year] and [Author B, year] both find [shared finding], though through different mechanisms: [A's mechanism] versus [B's mechanism]. The implication is [synthesis claim].</i>	Synthesis means showing relationships between findings. The 'both find... though through different mechanisms' structure is a typical synthesis move.

Representing a contested finding	<i>[Finding X] is contested. [Author A, year] find [result]; [Author B, year], using [different data/method], find [different result]. The disagreement is most plausibly explained by [methodological or contextual difference].</i>	Name both sides of the dispute. Explain the disagreement rather than picking a side without evidence. Readers need to know that you have engaged with the dispute, not avoided it.
Identifying a methodological limit	<i>A limitation of [most/early] work in this area is that [methodological problem — identification issue, data gap, measurement error]. This limitation means that [what prior results cannot establish].</i>	Be precise about limitations. 'Results may not generalise' is not a methodological critique; 'identification relies on cross-sectional variation that cannot rule out selection' is.
Stating the gap	<i>What the literature has not yet established is [specific question]: whether [condition A] changes the sign/magnitude/direction of [finding], particularly in [context or population not yet studied].</i>	The gap must be specific and falsifiable. A falsifiable gap names a condition whose variation would change the result — not just an area that has received less attention.
Introducing a seminal paper	<i>[Author(s), year] established the framework that most subsequent work builds on: [key contribution]. The paper showed that [finding], under the assumption that [key condition]. Subsequent work has extended this by [extension] and challenged it by [challenge].</i>	Seminal papers deserve explanation, not just citation. Show what they established and how the literature has developed in response.
Positioning the contribution	<i>Against this background, [this paper / the paper reviewed here] contributes by [specific contribution]: [brief description]. This is made possible by [data advantage / methodological innovation / context] that prior work lacked.</i>	The positioning move must be specific. The 'made possible by' clause explains why this paper can answer the question when prior work could not.

What makes 'Digital Economics' an instructive model is not that it reviews a large literature exhaustively — many reviews do that — but that its organising framework is itself the contribution. Goldfarb and Tucker do not simply describe what the digital economics literature has found; they reorganise it in a way that makes its internal logic visible, its debates tractable, and its next questions clear. That ambition — to make a contribution through the act of synthesis, not merely alongside it — is the standard against which any serious standalone literature review should be measured.

→ **Chapter Four: The Case Study as an Academic Publication**

CHAPTER FOUR

The Case Study as an Academic Publication

The case study has an ambiguous position in the hierarchy of economic research methods. In the natural sciences, the single case is typically the weakest form of evidence. In economics, by contrast, a rigorously executed case study can constitute a major scholarly contribution — not despite its focus on a single market, firm, or episode, but because of it. The case study's analytical power lies in its ability to examine a specific situation with a precision and granularity that broad cross-sectional analysis cannot achieve, to test theoretical mechanisms in a setting where the relevant variables are observable, and to generate welfare estimates or causal claims that are directly applicable to policy.

The ambiguity arises at the point of generalisation. A case study that makes no claims beyond its own setting is a description, not a contribution. A case study that claims universal generality has gone beyond what a single case can establish. The craft of the publishable case study lies in finding the right position between these two failure modes: stating what the case

establishes, what it suggests beyond its own setting, and where its findings depend on conditions that may not hold elsewhere.

This chapter uses Farronato and Fradkin's 'The Welfare Effects of Peer Entry: The Case of Airbnb and the Accommodation Industry' (Farronato & Fradkin, 2022), published in the *American Economic Review*, as its primary model text. The paper is unusually instructive for this chapter: it announces its own case-study status in its title, indicating that the authors knew exactly which methodological challenge they were taking on. Every structural decision in the paper reflects a deliberate response to that challenge.

4.1 The Model Text: Farronato and Fradkin (2022)

Farronato, C., and Fradkin, A. (2022). The Welfare Effects of Peer Entry: The Case of Airbnb and the Accommodation Industry. *American Economic Review*, 112(6), 1782–1817.

The paper studies the welfare effects of Airbnb's entry into the US accommodation market. Its core analytical contribution is a structural model of competition between two types of seller — flexible suppliers (peer hosts) and dedicated suppliers (hotels) — which differ in a theoretically significant way: peer hosts have alternative uses for their rooms and are therefore more responsive to demand fluctuations than hotels, whose capacity is fixed. The model generates a specific prediction: welfare gains from peer entry should be concentrated in times and places where hotel capacity is constrained, because it is precisely in those conditions that the flexibility of peer supply is most economically valuable.

The authors test this prediction using proprietary data from Airbnb and the hotel industry covering major US cities, and they find strong confirmation: welfare gains are concentrated in New York City and on high-demand dates such as New Year's Eve — exactly the conditions the model predicts. The welfare estimates are specific: Airbnb generated approximately \$41 of consumer surplus per room-night and \$26 of host surplus, whilst reducing variable hotel profits from accommodations by up to 3.7 per cent, for a total welfare gain of \$137 million in 2014 across the cities studied.

The paper is a good model for this chapter for several reasons. Its title announces the case study form, meaning the authors were explicitly working within the genre's constraints. Its structure — theoretical framework, data, reduced-form evidence, structural estimation, welfare analysis, heterogeneity — is a strong template for the publishable case study in applied economics. Its handling of the generalisation problem — what Airbnb establishes about peer entry more broadly — is precise and honest. And its policy implications section shows how to connect specific estimates to a live regulatory debate without overstating what the evidence supports.

4.2 What Makes a Case Study Publishable

A case study is publishable when the case is not merely interesting but analytically necessary — when its specific properties are what make the paper's contribution possible. This is a demanding standard. Many submitted case studies fail it not because the analysis is weak but because the authors have not asked — and therefore not answered — the question of why this specific case is the right vehicle for the paper's argument.

Farronato and Fradkin address this question directly. Their answer is two-fold: Airbnb is the right case because it provides unique data that make a welfare analysis otherwise infeasible, and because its market structure — the clean division between flexible and dedicated supply — maps directly onto the theoretical framework being tested. The case is not just interesting; it is analytically privileged. Identifying and stating this distinction is the first task of any publishable case study paper.

◆ Five types of case justification

There are five principal ways in which a specific case can be analytically justified for publication. They are not mutually exclusive — the strongest papers typically invoke more than one — but each makes a distinct claim about why the case is necessary rather than just convenient. The table below sets out each type with its logic and its illustration in Farronato and Fradkin.

Justification type	What it claims	Example from Farronato and Fradkin (2022)
Data availability	The case provides access to data not available elsewhere, enabling an analysis that would otherwise be infeasible	<i>Proprietary Airbnb transaction data and STR hotel data, covering major US cities, enable welfare estimation at a level of granularity unavailable from public sources</i>
Natural variation	The case exhibits variation in the variable of interest that allows causal or quasi-experimental analysis	<i>Airbnb's growth was uneven across cities and over time, providing cross-sectional and temporal variation that would be absent in a uniform national rollout</i>
Theoretical salience	The case is a particularly clear or extreme example of the phenomenon the theory predicts, making identification cleaner	<i>The peer-host / hotel distinction maps directly onto the flexible vs. dedicated supply framework; the case is analytically clean in a way that many sharing-economy settings are not</i>

Policy relevance	The case is directly implicated in an active policy debate, making the findings immediately applicable	<i>Airbnb's rapid growth prompted regulatory responses in multiple cities; the welfare estimates directly inform the benefit-cost calculus of such regulation</i>
Typicality	The case is representative of a broader class, enabling generalisation	<i>Airbnb is the dominant platform in most US cities and is structurally similar to other peer-to-peer marketplaces; findings are likely to generalise to comparable settings</i>

For writers, the practical implication is that the case justification must appear in the paper — not as an afterthought or as context-setting, but as part of the research design. A paper that moves directly from 'Airbnb is an important platform' to 'we study Airbnb' has asserted relevance without establishing analytical necessity. The justification must show that the case's specific properties — its data, its variation, its theoretical relevance, or its typicality — are what make the paper's contribution possible.

4.3 The Architecture of the Case Study Paper

A publishable case study paper in economics has a recognisable structure, and deviating from it has a cost. Editors and referees come to case study papers with specific expectations about what each section will do. A paper that does not meet those expectations creates friction that is often resolved by rejection. The table below sets out the standard sections of the economics case study paper, the function each is expected to perform, and the most common failure at each stage.

Section	Primary function	Most common failure at submission stage
Introduction	States the economic question, establishes why the case is the right vehicle for answering it, and positions the contribution relative to existing literature	<i>The case is justified by its familiarity or prominence rather than by its analytical properties; the research question is implicit rather than stated</i>

Case context	Provides the institutional, market, and historical background necessary to evaluate the analysis — and no more	Over-description: more background than the analysis requires, displacing the paper's own argument
Theoretical framework	States the economic model or conceptual structure that organises the analysis and generates testable predictions	The framework is described rather than justified; it is not connected to the specific properties of the case
Data and measurement	Documents the data sources, sample construction, and measurement decisions, with explicit justification for each choice	Measurement choices are stated without justification; their implications for the results are not acknowledged
Analysis and results	Presents the empirical or analytical findings in prose that leads and interprets, supported by tables and figures	Tables and figures are listed rather than interpreted; the results section narrates the data rather than directing the reader's attention
Heterogeneity and robustness	Tests the conditions under which results hold and vary; establishes the boundary conditions of the main finding	Robustness checks are perfunctory or mechanical; heterogeneity analysis is reported but not interpreted
Discussion and implications	Connects results to the theoretical framework, engages the existing literature, and states what the case study generalises to — and does not	Implications are overstated relative to the scope of the case; limitations are understated or absent

The most important structural principle — one that the anatomy table implies but does not state explicitly — is that the theoretical framework must come before the data and analysis, not after it. A paper that presents its data first and introduces its theoretical framework only later has, implicitly, allowed the data to generate the framework rather than the framework to guide the analysis. This is the case study equivalent of the research question problem discussed in

Chapter Two: the question must come before the method, and the framework must come before the evidence.

Farronato and Fradkin follow this sequence precisely. The theoretical model — differentiated competition between flexible and dedicated sellers — is introduced in the paper's second section, before any data are presented. The model generates predictions. The data are then used to test those predictions. This sequence is not just conventional; it is epistemologically necessary. A model introduced after the data have been described cannot be evaluated as a genuine prediction; it can only be evaluated as a post-hoc rationalisation.

We present a model of competition between flexible and dedicated sellers (peer hosts and hotels) who provide differentiated products. We estimate this model using data from major US cities and quantify the welfare effects of Airbnb on travelers, hosts, and hotels. The welfare gains are concentrated in specific locations (New York) and times (New Year's Eve) when hotel capacity is constrained. This occurs because peer hosts are responsive to market conditions, expand supply as hotels fill up, and keep hotel prices down as a result.

Farronato, C. and Fradkin, A. (2022). The welfare effects of peer entry: The case of Airbnb and the accommodation industry. *American Economic Review*, 112(6), p. 1782.

Abstract as architecture: the paper's logical sequence in four sentences

This abstract performs the research design of the paper in miniature. Sentence one introduces the theoretical framework (model of flexible vs. dedicated competition). Sentence two introduces the data and method (estimation using US city data). Sentence three states the main finding and its conditions (welfare gains concentrated under capacity constraint). Sentence four gives the causal mechanism (peer responsiveness to demand).

Notice that the finding is stated with its scope condition: 'concentrated in specific locations and times when hotel capacity is constrained.' The authors do not claim that peer entry always generates welfare gains; they claim that it does so under specific and theoretically predictable conditions. This is the generalisation discipline of the case study paper stated in a single sentence — and it appears in the abstract, which means the conditional framing of the finding is not a concession to limitations but the core of the argument.

4.4 Constructing the Theoretical Framework

The theoretical framework of a case study paper is not decorative. It is the source of the paper's testable predictions, the basis for interpreting the data, and the vehicle through which the case's findings connect to broader economic theory. A case study with no theoretical framework is a description. A case study with a weak framework is a description dressed in equations. The framework justifies its presence in the paper by generating specific predictions that the data can test.

◆ From market context to analytical model

The first move of the framework section is to identify the theoretically relevant features of the case — the characteristics of the market, the institutional setting, or the participants that connect the case to established economic theory. This is a selection problem: the case has many features, but only some are analytically relevant. The writer's task is to identify those features and show how they connect to the theoretical tradition being invoked.

Farronato and Fradkin's key theoretical move is the identification of the flexible/dedicated distinction as the market's most analytically important feature. Hotels have fixed capacity and are always available for booking; peer hosts have alternative uses for their space and are therefore more selectively available. This difference is not merely descriptive; it generates the paper's core prediction — that peer hosts will expand supply precisely when hotels are capacity-constrained — which is both theoretically derived and empirically testable.

A key difference between hotels and peer hosts is that while hotels have dedicated rooms that are always available for booking, peer hosts have alternative uses for their rooms, which make them more responsive to demand and price fluctuations.

Farronato, C. and Fradkin, A. (2022). The welfare effects of peer entry: The case of Airbnb and the accommodation industry. *American Economic Review*, 112(6), p. 1784.

Framework construction: the theoretical load of a single sentence

This sentence is doing more analytical work than its length suggests. 'Always available for booking' is not just a description of hotels; it is a statement about the nature of dedicated supply — supply whose cost structure means it is indifferent to whether a room is occupied or not in any given period. 'Alternative uses for their rooms' is not just a description of peer hosts; it is a statement about opportunity cost — the key variable that makes flexible supply price-sensitive at the margin.

The sentence is economically precise without being technically demanding. A reader without a formal training in industrial organisation can follow it; a reader with that training will immediately recognise the connection to models of competitive fringe entry and capacity-constrained incumbent pricing. The framework is accessible to a wide audience whilst remaining analytically rigorous — a balance that is harder to achieve than it appears and worth studying closely.

◆ Generating falsifiable predictions

A theoretical framework that does not generate falsifiable predictions is not functioning as a framework; it is functioning as background. The framework justifies its presence in the paper by making specific, directional predictions about what the data should show if the model is correct. The writer must state these predictions explicitly — not as hypotheses to be proved, but as consequences that follow logically from the model's assumptions and that the data will either confirm or challenge.

Farronato and Fradkin state their main prediction explicitly in the framework section: welfare gains from peer entry should be larger in locations and periods where hotel occupancy is high and hotel capacity is constrained. This is a conditional prediction — the effect is not uniform but varies systematically with an observable market characteristic. Conditional predictions are more valuable than unconditional ones precisely because they are harder to satisfy: they can fail in two ways (the effect may be absent altogether, or it may fail to vary in the predicted direction), and finding that the data satisfy both conditions simultaneously is stronger evidence than finding an unconditional average effect.

4.5 Data, Measurement, and the Credibility of Results

The data section of a case study paper must do three things: establish that the data are sufficient for the analysis being conducted; document the measurement decisions that transform raw data into the variables used in the model; and acknowledge the limitations of the data honestly without undermining the paper's credibility. A data section that is only descriptive —

that lists what data are available without explaining why they are appropriate — has not done the analytical work the section requires.

◆ Data sufficiency and provenance

Farronato and Fradkin's data section illustrates a particular challenge of case study research: the best data for studying a platform are often proprietary data held by the platform itself. The authors are explicit about this:

We use Airbnb data to study the welfare effects of facilitating peer entry in the accommodation market. While Airbnb is not the only company serving this market, it is the dominant platform in most US cities.

Farronato, C. and Fradkin, A. (2022). The welfare effects of peer entry: The case of Airbnb and the accommodation industry. *American Economic Review*, 112(6), p. 1785.

Data provenance: transparency about access and its implications

The paper is notable for the transparency of its conflict-of-interest disclosure: one of the two authors was previously employed by Airbnb and holds stock in the company. This disclosure appears in the paper's acknowledgements section rather than being concealed. The methodological lesson is not specific to this paper's circumstances; it is a general principle of data transparency: the provenance of proprietary data, and any relationship between the researchers and the data provider, must be stated. Readers evaluating the results need to know whether the data may have been filtered, sampled, or aggregated in ways that affect the findings.

The statement that Airbnb reviewed the paper to verify that confidential information was reported accurately is also instructive: it is an acknowledgement that the data provider had a review role, and that this role was limited to accuracy verification rather than editorial control over the findings. Making this distinction explicit is part of the paper's credibility architecture.

◆ Measurement decisions and their justification

Every decision about measurement in a case study is important. This includes how the main variable is defined, which time period is chosen, and how the sample is selected. Each of

these choices can influence the results. The writer should explain the main decisions clearly rather than presenting them as obvious.

If measurement choices are not explained, readers may think they were made to get a certain result. If the choices are explained and their possible effects are discussed, readers are more likely to trust the results.

In practice, it is important to understand which decisions have a strong effect on the results and which do not. Not every choice needs a long explanation. Good writing shows the ability to provide more explanation for important decisions and less for minor ones. If a writer explains all decisions the same way, it may suggest they do not understand which are more important.

◆ **Reduced form before structure**

In empirical case study papers, it is important to organise the analysis in a clear order. First, researchers usually present reduced-form analysis. This means they show simple results that connect the main variable with outcomes, without using a complex model. This helps to show the basic patterns in the data. After that, they present the structural analysis, which uses a model to explain these patterns.

Starting with reduced-form analysis has two advantages. First, it shows that the results are not solely due to the model's assumptions. Second, it makes the paper easier to understand for readers who may not be familiar with complex models.

Farronato and Fradkin follow this approach. They first show simple evidence about how Airbnb affects hotel prices and occupancy. This shows that the effect exists and goes in the expected direction. Then they use a model to measure the broader effects.

This order is not only about presentation. It also helps to show that the results are reliable. A simple analysis cannot explain everything, such as overall welfare, but it can show whether the model's main idea matches the data.

4.6 Presenting Results: Prose, Tables, and Figures

The results section of a case study paper is usually the most technical part, and it is often poorly written. A common problem is a simple description. The writer only explains what each table or figure shows, in order, without explaining what it means.

A results section like this does not help the reader understand the findings. It merely repeats the content of the tables rather than presenting a clear argument.

◆ **The prose leads**

In the results section, the text should guide the reader. The writer should first clearly explain the main finding. They should also describe how strong the result is, under what conditions it is true, and what it means for the main argument of the paper.

This explanation should precede or accompany the reference to a table or figure. Tables and figures provide evidence, but they do not replace explanation. The main ideas should always be presented in the text.

✗ Narrating the results	✓ Leading with the finding
<i>Table 3 presents the results of the welfare analysis. Column 1 shows consumer surplus per room-night. Column 2 shows host surplus per room-night. Column 3 shows the effect on hotel profits.</i>	<i>Airbnb generates substantial consumer surplus — approximately \$41 per room-night — with a smaller but positive gain to hosts (\$26) and a modest reduction in hotel profits of up to 3.7 per cent (Table 3). The net welfare effect is positive, but its distribution raises distinct policy considerations: the gains to consumers and hosts accrue broadly, whilst the cost is concentrated in the hotel sector.</i>
<i>Figure 2 plots the relationship between Airbnb supply share and hotel occupancy rate across cities. The x-axis shows the Airbnb supply share and the y-axis shows hotel occupancy rates.</i>	<i>The cross-city pattern in Figure 2 confirms the model's prediction: Airbnb's supply share is systematically larger in cities with higher baseline hotel occupancy rates. This is consistent with peer hosts entering precisely where the returns to flexible supply are highest — where hotels are most frequently at or near capacity.</i>

◆ Heterogeneity as argument

In a case study paper, heterogeneity analysis is not only a way to check results. It is an important part of the argument. When results differ across groups, time periods, or situations, this helps us understand why the result occurs and when it applies.

If a paper shows only average results and does not explore these differences, it misses an important part of the analysis.

In the study by Farronato and Fradkin, this type of analysis is especially important. It tests the model more directly. They find that the positive effects are strongest on New Year's Eve and in New York City. This is not a minor detail — it is a key result.

This finding supports the main explanation, because the effect is strongest exactly where the model predicts it should be.

Travelers benefit from Airbnb for two reasons. First, flexible sellers offer a differentiated product relative to hotels. Second, they also compete with hotels by expanding the number of rooms available. This second effect is particularly important in periods of high demand when hotels are capacity constrained and can thus charge higher prices.

Farronato, C. and Fradkin, A. (2022). The welfare effects of peer entry: The case of Airbnb and the accommodation industry. *American Economic Review*, 112(6), p. 1800.

Decomposing the mechanism: the two sources of welfare gain

This passage does something that a table cannot: it decomposes the overall welfare effect into its constituent mechanisms. The first mechanism (product differentiation) operates continuously — peer hosts always offer a somewhat different product from hotels, and consumers always value this variety. The second mechanism (capacity relief) operates conditionally — it matters most when hotels are full and can charge high prices, which is precisely when peer supply is most responsive.

By articulating both mechanisms and their different conditions of operation, the passage explains why the heterogeneity analysis takes the form it does: the paper looks for the effect of capacity constraint not as a robustness exercise but because the model says that the second mechanism — the more quantitatively important one — is visible precisely under those conditions. This is heterogeneity analysis motivated by theory rather than data mining, and it is considerably more persuasive as a result.

4.7 The Generalisation Problem

The generalisation problem is the defining methodological challenge of the case study paper. It has two components: the epistemic challenge (what does a single case establish about the broader world?) and the rhetorical challenge (how does the writer state the paper's implications without either understating or overstating what the evidence supports?). Both challenges must be addressed explicitly in the paper's discussion and conclusion sections.

◆ What a single case can establish

One case study cannot prove a general rule that is true in all situations. To do that, we need evidence from many different cases. However, a case study can demonstrate that an idea works in practice. It can demonstrate that a theory is possible in at least one situation. It can also provide specific results and calculations for that case.

The value of such a study depends on two things. First, how widely the main idea can be applied. Second, how relevant the specific case is to the problem being studied.

Farronato and Fradkin clearly explain both points. Their main idea — that flexible sellers can fill gaps left by more fixed sellers — could apply to many markets, such as ride-sharing, freelance work, or online services. However, their detailed results are based only on the US accommodation market in 2014.

The authors carefully separate these two levels. The general idea may apply to many situations, but the exact numbers may not.

◆ The register of appropriate generalisation

The language of generalisation in a case study paper must be calibrated to what the evidence supports. The table below identifies four types of generalisation claim, with appropriate formulations and the conditions under which each becomes problematic.

Claim type	Appropriate formulation	Warning: when this register becomes problematic
Conditional generalisation	<i>'Under conditions of capacity constraint, peer entry generates welfare gains concentrated amongst consumers.' The condition is stated; the claim is bounded.</i>	Becomes problematic when the conditions are not specified precisely enough to be testable or when they are presented as general rather than contingent
Mechanism-based generalisation	<i>'The mechanism driving these results — flexible suppliers expanding to fill capacity gaps that dedicated suppliers cannot — is likely to operate wherever flexible and dedicated supply co-exist.' The mechanism, not the finding, is generalised.</i>	Becomes problematic when the mechanism is described so abstractly that it loses the content needed to generate falsifiable predictions in new contexts

Negative generalisation	<i>'These results do not support the view that peer entry necessarily displaces incumbent sellers; displacement is concentrated in high-demand periods and appears modest in magnitude.' The claim is about the limits of an existing generalisation.</i>	Becomes problematic when negative results are used to dismiss a body of evidence rather than to refine it
Explicit scope limitation	<i>'These estimates are specific to the US accommodation market in 2014 and should not be applied mechanically to other markets or periods.' The writer states what the case does not establish.</i>	Becomes problematic only if the limitation is not accompanied by a positive statement of what the case does establish

◆ Implications for policy

The policy implications section of a case study paper is often the most difficult part. This is because it tries to move from one case to general advice about what should be done in other situations.

Writers need to avoid two common problems. First, they should not make the implications too weak, or the paper will have little practical value. Second, they should not make them too strong by treating results from one case as if they apply everywhere.

Farronato and Fradkin provide a good example. They do not suggest one specific policy. Instead, they use their results to explain the possible costs and benefits of regulation.

They find that the positive effects are strongest during periods of high demand. This suggests that general restrictions on peer-to-peer accommodation may not be effective, because they would reduce benefits exactly when they are greatest.

This conclusion is useful for policy, but it stays within the limits of the study. It does not claim that the same results will apply in all countries or situations.

4.8 Limitations: Proportionate, Not Performative

Every case study paper has limitations, but this part is often poorly written. There are two common problems. First, some papers do not clearly discuss their limitations, so readers must find the problems themselves. Second, some papers list too many limitations without explaining which ones are important.

A good approach is to focus on importance. The writer should explain the most important limitations in more detail. If a limitation could strongly change the results, it should

be clearly discussed. The writer should explain why their approach remains reasonable or acknowledge that the results may be limited. Less important limitations can be mentioned briefly.

Farronato and Fradkin give a good example. They explain that their analysis looks only at short-term effects. Their model does not include long-term changes, such as the impact of Airbnb on housing supply and rent prices.

The authors clearly say this, explain why they did not include it, and suggest how it might affect their results. This is a good way to discuss limitations: it is clear, honest, and directly connected to the findings.

⚠ Four failure modes in case study writing — a pre-submission checklist

1. Case unjustified — the paper moves from 'this case is important/interesting' to 'we study this case' without establishing analytical necessity. Fix: state explicitly which analytical property of the case (data, variation, theoretical salience, typicality) makes it the right vehicle for the research question.
2. Framework follows data — the theoretical model is introduced after the data have been presented, implying post-hoc rationalisation. Fix: state the model, its assumptions, and its predictions before presenting any data.
3. Narrating results — the results section describes what tables and figures show rather than stating and interpreting findings. Fix: every results paragraph should open with the finding, state the relevant magnitude and condition, and connect it to the paper's argument.
4. Overgeneralised implications — specific estimates from a single case are presented as universal prescriptions. Fix: distinguish the level at which the mechanism generalises from the level at which the specific estimates apply, and state scope conditions explicitly.

4.9 Writing from a Ukrainian and Regional Context: Three Case Studies

The principles in sections 4.1–4.8 are based mainly on the standards of leading international economics journals. These are the expectations for papers submitted to top journals.

However, researchers in Ukraine and Eastern Europe often publish in different types of sources. These include regional journals, open-access conference series, short conference papers, and international journals. This section looks at the specific challenges and opportunities in this environment. It uses three papers by Chala and co-authors as examples.

These three papers are different in type and length. One is a literature review in a Ukrainian journal, one is a short empirical study in a conference series, and one is a very short conference paper. Together, they help answer important questions. For example: How do

regional standards compare to international ones? How can a short paper develop into a full article? And how can Ukrainian research be presented for an international audience?

The three model texts for this section

Paper A (Literature review): Chala, N., Pichyk, K., and Voropai, O. K. (2021). Practice of Neuromarketing Research on Artwork Consumption: Literature Review and Marketing Management Perspective. *Bulletin of Khmelnytskyi National University*, 1(290), 159–164.

Paper B (IOP conference paper): Lopuschnyak, H., Chala, N., and Poplavska, O. (2021). Socio-economic Determinants of the Ecosystem of Sustainable Development of Ukraine. *IOP Conference Series: Earth and Environmental Science*, 915(1), 012019.

Paper C (Conference proceedings): Chala, N. and Poplavska, O. (2025). Resilient Teams in a Volatile World: The Impact of Global Megatrends on Employee Motivation and HR Strategies. In: *Planning and Ensuring Sustainable Development of Socio-Economic Systems: Materials of the 1st International Scientific and Practical Conference*, 18 December 2025. Poznan: WSHIU Akademia Nauk Stosowanych, Part 1, pp. 42–43.

Regional vs. international publication standards

A common question for researchers who want to publish internationally is whether writing rules in regional journals differ significantly from those in top international journals, or if the only difference is prestige.

The answer is both. There are differences in prestige and access to data, but there are also differences in how papers present their arguments.

Paper A, a literature review on neuromarketing, is a good example. Its abstract clearly explains its aim: to review research on gallery visitors and to suggest how certain tools can be used in marketing. The paper's structure is logical. It reviews existing studies, groups marketing problems, proposes hypotheses, and discusses implications. It also presents five clear and testable hypotheses.

However, compared to top international journals, there are some differences. First, the review primarily describes existing studies rather than developing a strong central argument or identifying a key problem in the literature. Second, the hypotheses are presented at the end rather than clearly developed from the earlier discussion. Third, the practical implications are expressed as general possibilities rather than as clear, evidence-based conclusions.

These are not fundamental weaknesses and can be improved during revision. The main point is that the difference between regional and international publications is not usually in the topic, but in how clearly the paper explains the connection between previous research, the research gap, and its own contribution.

Regional register (Paper A, as published)	International register (revised formulation)
<i>'We suggest that neuroscience tools (fMRI and eye-tracking) can be beneficial for solving all three main types of marketing research problems.'</i>	<i>'The existing literature has demonstrated that eye-tracking reliably captures attentional patterns in gallery settings; what remains untested is whether these patterns predict subsequent purchase decisions. We argue that this gap is addressable using the causal research design proposed below.'</i>
<i>'Artworks are difficult to evaluate in terms of objective and measurable parameters. Therefore, people's preferences towards a specific piece of art are hard to explain.'</i>	<i>'Because artwork preferences are known to involve high emotional and subconscious components (Bauer & Schwan, 2018), self-report measures systematically underestimate the predictive variables — a measurement problem for which neuromarketing tools offer a specific and empirically grounded solution.'</i>

The contrast above is not a criticism of Paper A; it is a demonstration of the kind of revision a strong regional paper undergoes when it is developed for international submission. The research content is the same in both columns. What changes is the precision of the logical relationship between problem, method, and claim — and that precision is exactly what international referees are trained to evaluate.

◆ From conference paper to journal article: the development path

Paper C — the two-page conference proceedings piece on resilient teams and global megatrends — represents a genre that is widely misunderstood in the Eastern European academic context. A two-page conference proceedings contribution is not a short journal article; it is a research agenda statement. Its function is to stake a claim in a debate, attract potential collaborators, and test the reception of a set of ideas before the full analytical work is undertaken. Treating it as a complete publication — and citing it as such in a CV or research portfolio — overstates its academic standing relative to peer-reviewed research.

This does not mean conference proceedings contributions are without value; they are a normal and appropriate starting point for research that is not yet ready for journal submission. What matters is the development path — the sequence of steps by which the two pages of Paper

C become a publishable journal article. That path has four stages, and understanding them is practically useful for faculty at the beginning of a research programme.

Stage	What the stage produces	Applied to Paper C (Chala and Poplavska, 2025)
1 — Research question specification	A single, focused, falsifiable question that the eventual paper will answer. At this stage, the question replaces the thematic overview.	<i>From 'the impact of global megatrends on employee motivation and HR strategies' (a topic) to: 'Under conditions of rapid technological disruption, do HR resilience practices moderate the relationship between megatrend exposure and employee motivation outcomes? Evidence from Ukrainian firms, 2022–2024.'</i>
2 — Literature positioning	A synthesis of the 30–50 most relevant papers that establishes the gap the research question addresses. This is a full Chapter Three exercise, not a background section.	<i>The paper cites megatrend frameworks and HR strategy literature; a full positioning would map the specific debate about resilience as an HR outcome variable and identify where the Ukrainian context has not been studied.</i>
3 — Data and method design	A data collection protocol, identification strategy, and analysis plan connected explicitly to the research question. For Paper C, primary data collection would be required.	<i>The 2025 conference paper does not report original empirical data. The journal version requires: survey design or case study protocol, sample and field period, and a method appropriate to the causal or descriptive claim being made.</i>
4 — Full draft and outlet selection	A complete manuscript conforming to the target journal's conventions, with the research question, framework, data, results, and discussion all present and coherent.	<i>Appropriate target outlets given the topic and Ukrainian data: Journal of Business Research, International Journal of Human Resource Management, or Management and Organization Review (which specifically invites research from non-Western contexts).</i>

◆ Framing Ukrainian case material for an international audience

Paper B, a conference paper about sustainable development in Ukraine, shows a different challenge. It is about how to present research based on Ukrainian data so that it is interesting for an international audience. This is not only about method, but also about how the study is presented.

Many researchers begin by explaining why Ukraine is important or not well studied. This approach is not wrong, but it is not strong enough. It focuses on the country itself. A stronger approach is to focus on the research question and explain why Ukrainian data help to answer this question. International reviewers are more interested in the research problem than in the country.

Paper B studies the factors that influence sustainable development using Ukraine as an example. For a stronger international presentation, the order can be changed. First, the paper should introduce the main research question. Then, it should explain why Ukraine is a useful case to study this question, for example because of its economic situation or available data.

In this way, Ukraine is not only the subject of the study, but also a useful case that helps answer a broader research question.

✗ Geography-led framing	✓ Question-led framing with Ukrainian case justification
<i>'Ukraine is an important transitional economy whose sustainable development challenges have received insufficient attention in the international literature. This paper analyses the socio-economic determinants of sustainable development in Ukraine.'</i>	<i>'Identifying the socio-economic determinants of national sustainable development ecosystems requires variation in institutional context, development stage, and policy environment that cross-country panel studies cannot provide. Ukraine's post-2014 transition period offers a natural quasi-experiment in institutional restructuring that enables the causal identification of determinants that remain contested in the broader literature. This paper exploits that variation to test competing theoretical accounts of ..'</i>

This second approach does not reduce the importance of Ukraine. It makes it stronger. It shows that Ukraine is not only an interesting case, but also a very useful one for answering an academic question. Its special features make it a good example for research. This is similar to how some studies focus on a specific case to explain a broader idea. This approach helps turn a country study into research that can be published internationally.

Another important point is how to use Ukrainian-language sources in international papers. These sources are valid and should not be hidden or replaced. However, they should be presented clearly for international readers. Since some readers may not know these sources, the writer should explain where they come from, for example the journal and its academic status.

This helps readers understand the quality of the source. It is not about simplifying, but about making the research clear and accessible to an international audience.

⚠ Three practical principles for publishing from a Ukrainian research context

1. Lead with the question, not the country. Frame Ukraine as the analytical vehicle that makes the research question tractable, not as the object of study whose importance must be asserted. International referees are more responsive to 'this case enables us to identify X' than to 'Ukraine is important and understudied.'
2. Know the development path from conference paper to journal article. A two-page conference proceedings piece is a research agenda statement, not a short publication. Its value lies in what it becomes. The four-stage development path — research question specification, literature positioning, data and method design, full draft — is the translation of that agenda into a publishable contribution.
3. Close the precision gap. The most common difference between a strong regional paper and an internationally competitive one is not the research question but the precision with which the relationship between the existing literature, the identified gap, and the paper's own contribution is stated. Revising for this precision — in the introduction, in the literature review, and in the discussion — is the highest-return writing investment a faculty member can make.

4.10 Vocabulary and Sentence Models

The vocabulary below covers terms that are particularly important in the case study genre, where the relationship between a specific case and a general theoretical or empirical claim requires careful handling. The sentence models demonstrate the moves that distinguish publishable case study writing from case study reporting.

◆ Key vocabulary

Term	Meaning in case study writing	Example of the word in a sentence
------	-------------------------------	-----------------------------------

analytical leverage	The degree to which a case allows a researcher to test or refine a theoretical claim. A case with high analytical leverage is one whose specific institutional or structural features make the theoretical mechanism visible in the data.	<i>'Airbnb provides exceptional analytical leverage for testing the welfare effects of flexible supply entry: the platform's design creates a clean division between hosts who rent occasionally and those who operate professionally, which directly maps onto the model's key distinction between flexible and dedicated supply.'</i>
theoretical framework	The set of concepts, mechanisms, and relationships that structure the paper's analysis of the case. In economics, the theoretical framework typically includes a formal model or a set of propositions derived from existing theory, which the case study then tests or extends.	<i>'The theoretical framework draws on models of platform competition with multi-dimensional supply, following Rochet and Tirole (2003). The key mechanism — that flexible supply entry compresses price dispersion among incumbents without eliminating welfare-improving quality differentiation — is derived in Section 2 and tested in Sections 4–6.'</i>
external validity	The degree to which the findings from one case generalise to other contexts. External validity is a central concern in case study research, because a finding that holds only in the studied case adds little to the general literature.	<i>'External validity is supported by three features of the Airbnb case that are not idiosyncratic to the accommodation industry: the cost structure of the entering platform, the elasticity differential between flexible and dedicated supply, and the nature of quality as a non-price competitive dimension.'</i>
counterfactual	What would have happened in the absence of the treatment, intervention, or event being studied. In case study research, counterfactual reasoning is the basis of causal claims: the effect of X is the difference between what happened with X and what would have happened without it.	<i>'The counterfactual — what hotel revenues would have been without Airbnb entry — is estimated using the synthetic control method, which constructs a weighted average of control markets that closely tracks the pre-entry trend in treated markets.'</i>

mechanism	The causal process by which an independent variable produces an effect on a dependent variable. Identifying the mechanism is what separates a paper that shows a correlation from one that explains a causal relationship.	<i>'The paper identifies the mechanism through which Airbnb entry reduces hotel revenues: not price competition directly (hotel prices are largely unchanged) but occupancy effects, as the marginal consumer who would previously have chosen a budget hotel now prefers short-term rental accommodation.'</i>
heterogeneity	Variation in the size, sign, or form of an effect across different sub-groups, contexts, or time periods. Reporting heterogeneity is important because a single average effect can conceal patterns that are theoretically important or policy-relevant.	<i>'Heterogeneity by hotel category is the most important finding: Airbnb entry reduces occupancy and revenue for budget and mid-range hotels, but has no detectable effect on luxury hotels — consistent with the theoretical prediction that flexible supply competes most directly at the lower end of the quality distribution.'</i>
welfare	In economics, welfare refers to the well-being of economic agents — consumers, producers, or society as a whole — typically measured in terms of utility or surplus. Welfare analysis evaluates who gains and who loses from a change, and by how much.	<i>'Consumer welfare, measured as the area under the inverse demand curve above price, increases following Airbnb entry: the entry of lower-cost flexible supply expands accommodation options for price-sensitive consumers who were previously unable to afford the available hotel stock.'</i>
limitation	A specific constraint on what the paper's findings can establish — arising from the data, the identification strategy, the scope of the case, or the theoretical model. Limitations should be stated proportionately: acknowledging real constraints builds credibility; excessive qualification undermines it.	<i>'A limitation of this analysis is that it cannot directly measure quality: occupancy and revenue are observed, but the utility consumers derive from choosing between hotel and Airbnb accommodation is not. The welfare estimates therefore depend on the assumption that revealed preferences in booking data accurately reflect underlying utility.'</i>

◆ Sentence models

The models below cover the key moves specific to the case study paper, from justifying the case selection to handling limitations. Adapt by replacing bracketed elements with your own case and theoretical content.

Function	Model sentence	Notes for the writer
Justifying case selection	<i>This paper studies [case] not because of its prominence but because [specific structural or institutional feature] makes it the appropriate vehicle for testing [theoretical mechanism].</i>	Case justification must be analytical, not circumstantial. Connect the case's specific properties to the theoretical mechanism — this is the move Farronato and Fradkin make with Airbnb.
Introducing the theoretical framework	<i>The analysis is grounded in [framework], which predicts that [mechanism]. This prediction generates the paper's central hypothesis: that [testable claim] when [condition].</i>	Name the framework, state the mechanism it identifies, and convert it into a testable hypothesis. The hypothesis is what the case study will confirm, qualify, or challenge.
Justifying a data choice	<i>[Data source] is used to measure [variable] because [advantage over alternatives]. An alternative measure — [alternative] — is used in the robustness checks reported in [section/appendix].</i>	Every data choice requires justification. State the advantage of the chosen measure and acknowledge what the alternative would have offered.
Reporting a main result with mechanism	<i>[Outcome] [falls/rises/changes] following [treatment/event] (Table [N]). This is consistent with [mechanism]: [explanation of why the result is what the mechanism predicts].</i>	State the result, point to the evidence, then connect the result to the mechanism. This three-part structure is what distinguishes case study argument from case study description.

Reporting heterogeneity	<i>The effect is heterogeneous across [dimension]: [sub-group A] experiences [effect A], while [sub-group B] experiences [effect B] (Table [N], Columns [N]–[N]). This pattern is consistent with [theoretical prediction about heterogeneity].</i>	Heterogeneity results are more informative than average effects. Connect the pattern to the theoretical prediction — this is where the case study contributes most distinctively.
Addressing external validity	<i>The findings may extend to [other contexts] where [conditions that support generalisation] obtain. However, they should not be assumed to hold where [conditions that would change the result], because [theoretical reason].</i>	Frame generalisation as conditional, not universal. Name the conditions that support it and those that would limit it.
Stating a limitation proportionately	<i>A limitation of this analysis is that [specific constraint]. This limits the paper's ability to [specific thing it cannot establish]. It does not affect the main finding that [main finding], which depends only on [what the finding does depend on].</i>	The second sentence is what separates proportionate from excessive limitation statements. Show that the limitation is real but bounded — that it constrains a specific claim, not the whole paper.
Connecting the case to the broader literature	<i>This finding contributes to the broader literature on [theme] by showing that [specific claim], a pattern that [prior reviews/papers] had not identified because [reason — data, context, methodological approach].</i>	The contribution statement should specify what the case study adds that prior work missed. This is where regional papers often gain international traction — by showing a pattern in a context the general literature had not studied.

Farronato and Fradkin did not study Airbnb because it was prominent. They studied it because its market structure — the clean division between flexible and dedicated supply — made it the right vehicle for testing a theoretical mechanism that, if confirmed, had implications well beyond the accommodation industry. That is the standard for a publishable case study: not that the case is important in the world, but that its specific analytical properties are what make the paper's contribution possible. Getting that justification right — and sustaining it through the paper's framework, data, and findings — is what distinguishes a case study paper from a case study report.

→ **Chapter Five: Data Interpretation and Quantitative Writing**

CHAPTER FIVE

Data Interpretation and Quantitative Writing

Quantitative writing is the form of academic economics writing that most consistently exposes the gap between technical competence and communicative clarity. A researcher who has executed a carefully designed empirical study — assembled a large dataset, constructed

appropriate measures, addressed identification, estimated precise effects — can still produce a paper that fails to communicate its contribution if the writing does not match the quality of the analysis. Numbers do not speak for themselves. They require prose that states what they show, what they mean, how large the effects are relative to relevant comparators, and what the reader is justified in concluding from them. This chapter addresses that prose.

The model text for this chapter is Chetty, Hendren, Kline, and Saez's 'Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States' (Chetty et al., 2014), published in the *Quarterly Journal of Economics*. The paper uses administrative tax data covering nearly all US families with children born between 1971 and 1993 to construct measures of intergenerational income mobility for each of 741 commuting zones. It is chosen as a model text for three reasons: it is one of the most carefully written quantitative papers in modern economics; its measurement decisions are documented with a thoroughness that is instructive in itself; and its handling of the distinction between correlation and causation — a distinction that quantitative writing must manage with constant precision — is exemplary.

5.1 The Model Text: Chetty, Hendren, Kline, and Saez (2014)

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where is the land of opportunity? The geography of intergenerational mobility in the United States. *Quarterly Journal of Economics*, 129(4), 1553–1623.

The paper's central finding is that intergenerational income mobility varies enormously across the United States: children who grow up in some commuting zones reach substantially higher income ranks in adulthood than children from similarly situated families who grow up elsewhere. This finding is stated with precision from the outset. The abstract gives the range (absolute upward mobility — the mean income rank of children with parents at the 25th percentile — runs from 35.8 percentile points in Charlotte, North Carolina to 46.2 in Salt Lake City, Utah), gives a comparator (this gap is larger than the difference between the United States and Denmark in cross-country comparisons of mobility), and states the causal interpretation the paper is and is not claiming (the correlational findings identify associations but not causal mechanisms; the paper's companion study on movers addresses causality directly).

That discipline — state the finding, give its magnitude, give a comparator, state the causal status — is the template this chapter uses throughout. It is not unique to Chetty et al. (2014), but the paper executes it with unusual consistency and can be read, alongside this chapter, as a sustained demonstration of the principles discussed below. The full paper is openly available via the NBER working paper series (Working Paper 19843) and through institutional access to the *Quarterly Journal of Economics*.

5.2 Measurement Decisions and Their Justification

Every quantitative paper depends on a set of measurement decisions: how the key outcome variable is constructed, how the sample is bounded, how time periods are chosen, what

geographic unit is used, and how the main explanatory variable is operationalised. These decisions are not merely technical; they are intellectual commitments that shape what the paper can and cannot establish, and that experienced readers evaluate as carefully as they evaluate the statistical results.

Measurement documentation requires the writer to state each significant decision, justify it, and acknowledge what was traded off. A paper that states its measurement choices without justifying them invites the inference that the choices were arbitrary or, worse, that they were made to produce a desired result. A paper that justifies its choices and acknowledges the alternatives signals that the authors have thought carefully about the costs and benefits of each decision — which is itself evidence of the analysis's reliability.

◆ The four-column justification discipline

The most useful way to think about measurement documentation is to ask four questions for each significant decision: what choice was made, why was it made, what would the alternative have cost, and how sensitive are the results to reasonable variations in the choice? Not all four questions need to be answered in the main text — sensitivity analysis can often go in an appendix — but the writer must have worked through all four in order to write the justification confidently. The table below applies this discipline to five of the key measurement decisions in Chetty et al. (2014).

Decision	Choice made by Chetty et al. (2014)	Justification given	What the alternative would have cost
Income measure	<i>Mean household income at ages 26–30, measured using tax records</i>	Captures household resources more completely than individual earnings; age range chosen to capture settled adult income rather than early-career variation	Individual earnings would miss family formation effects; measuring too early (age 22) would pick up in-school non-earners; too late (40+) reduces sample size for recent cohorts
Mobility measure	<i>Rank–rank slope (child's income rank regressed on parent's income rank)</i>	Rank-based measure is comparable across areas with different income distributions; less sensitive to outliers than level-based measures	Level-based mobility would conflate differences in local price levels with real differences in economic opportunity; percentile ranks are meaningful on a common national scale

Geographic unit	<i>Commuting zone (CZ) — a functional economic area based on commuting patterns</i>	CZs capture local labour markets more accurately than counties or states; their boundaries reflect economic geography rather than administrative convenience	County boundaries often split metropolitan areas; state-level analysis obscures the enormous within-state variation that is the paper's central finding
Parent income measure	<i>Mean household income between 1996 and 2000</i>	Five-year mean reduces noise from single-year fluctuations; the period captures parents' income when most sample children are in their teens — the relevant developmental window	Single-year income would introduce substantial measurement error; using tax years before 1996 would lose coverage for younger birth cohorts
Sample restriction	<i>Children born 1980–1982 who lived with their parents in 1996</i>	Birth years provide sufficiently long outcome window; co-residence condition identifies parent-child pairs in tax data without relying on household surveys	Wider birth year range would require pooling cohorts with very different labour market entry conditions; dropping co-residence requirement would introduce family structure confounds

The table illustrates that measurement justification is not a formal exercise; it is a substantive argument about why the analysis is credible. Each of the five decisions in Chetty et al. (2014) is made for a specific reason that connects to the paper's research question, and the cost column makes clear what the alternative would have sacrificed. Writers who have worked through this logic — even informally, before writing the data section — produce justifications that are substantive rather than formulaic.

◆ Documenting sample construction

Sample construction — the sequence of decisions by which a raw dataset is reduced to the analysis sample — is an area of measurement documentation that is often handled too briefly. The standard practice in a top-journal paper is to present the sample construction as an explicit sequence of steps, with the number of observations at each stage, so that the reader can see how many observations were dropped and why. This is not a courtesy; it is a check. A

sample construction sequence that drops 60 per cent of the raw observations without explanation creates legitimate concern that the exclusions were not innocuous.

Chetty et al. (2014) document their sample construction with particular care because their data source — federal tax records — is unusual, and the co-residence condition that identifies parent-child pairs in administrative data requires explicit justification. The paper explains the birth year restrictions, the co-residence requirement, and the income thresholds used to exclude implausible observations, and it checks the representativeness of the analysis sample against external benchmarks. This level of documentation is the appropriate model for any paper whose sample construction involves non-trivial choices.

5.3 Communicating Magnitudes

The most common failure in quantitative writing — from dissertation to journal submission — is not incorrect results but uncommunicated results: findings reported in statistical terms without translation into terms that make them interpretable. A coefficient of 0.35 on a rank-rank regression is not self-evidently large or small; a slope of 0.35 means that children whose parents are 10 percentile points higher in the income distribution end up, on average, 3.5 percentile points higher themselves. The second formulation conveys a finding; the first conveys a parameter estimate. They are not the same thing.

Communicating magnitude requires three elements: a statement of the finding in units the reader can evaluate; a comparator that places the finding on a scale of large and small; and a statement of the conditions under which the finding holds. All three are present in the most effective quantitative prose. Chetty et al. (2014) use this structure consistently throughout their results section.

We find significant variation in upward mobility across CZs. Absolute upward mobility — the mean rank of children with parents at the 25th percentile — ranges from 35.8 in Charlotte to 46.2 in Salt Lake City among the 50 largest CZs. The differences in mobility across CZs are not simply driven by differences in current income levels or growth: high-mobility areas do not have higher mean incomes than low-mobility areas.

Chetty, R. et al. (2014). Where is the land of opportunity? Quarterly Journal of Economics, 129(4), p. 1562.

Three sentences, three functions: finding, comparator, scope condition

The first sentence announces the finding: mobility varies significantly. The second sentence gives the finding its content: not 'significant variation' in the abstract, but a specific range with named places attached. Charlotte and Salt Lake City are named not for vividness but because they are the endpoints of the distribution — the reader is being given the range, not a representative example.

The third sentence is a scope condition disguised as a negative finding: high mobility is not simply a proxy for current prosperity. This matters because the most obvious alternative explanation for variation in upward mobility is that rich places have more mobile children. By showing that the correlation between current income and mobility is weak, Chetty et al. (2014) establish that they are measuring something real and independent — and they do this in a single sentence, without relegating it to a footnote.

Notice that no standard error appears in this passage. Chetty et al. (2014) use sufficiently large samples that precision is not the primary concern, and cluttering a finding sentence with standard errors obscures the finding. Precision should be reported where it matters — in tables, or in prose when the statistical significance of a specific comparison is the point. It should not be inserted reflexively into every sentence that reports a number.

◆ Five registers for describing magnitude

Describing the magnitude of a finding is not a single act but a set of choices about how to contextualise a number. The table below identifies five registers for communicating magnitude, with examples from or calibrated to Chetty et al. (2014), and the conditions under which each is appropriate and potentially problematic.

Register	Example formulation	When to use / when problematic
----------	---------------------	--------------------------------

Absolute with comparator	<i>'Absolute upward mobility ranges from 35.8 in Charlotte to 46.2 in Salt Lake City — a gap of 10.4 percentile points, roughly equivalent to the difference between median US and Canadian mobility.'</i> (Chetty et al., 2014)	Use when the raw number has a natural comparator that the reader can evaluate. Problematic when the comparator is chosen to make the effect look larger or smaller than alternative comparators would suggest.
Standard deviation benchmark	<i>'A one standard deviation increase in CZ-level upward mobility is associated with a 0.2 standard deviation improvement in a child's expected rank.'</i> (Chetty et al., 2014)	Use when absolute magnitudes are hard to evaluate without context. Gives the reader a sense of the signal-to-noise ratio in the data. Problematic when standard deviations themselves are not reported and the reader cannot reconstruct the implied absolute magnitude.
Back-of-envelope translation	<i>'Children whose parents earn \$27,000 and who grow up in the top versus bottom quartile of CZs by mobility have mean adult incomes that differ by approximately \$10,000 per year.'</i>	Use to translate statistical findings into units that are meaningful for policy audiences. Requires explicit statement of the assumptions behind the translation. Problematic when the translation embeds assumptions that are not flagged.
Hedged comparative	<i>'The relationship between inequality and mobility is negative — areas with higher Gini coefficients tend to have lower upward mobility — but the relationship is modest in magnitude and not the primary driver of geographic variation.'</i> (Chetty et al., 2014, paraphrased)	Use when a finding is real but its effect size is small relative to other factors. Honesty about small effects is a scholarly virtue. Problematic when 'modest' is used to dismiss a finding that, at scale, is economically significant.

Decomposition statement	<i>'Approximately 50 per cent of the cross-CZ variation in upward mobility is explained by standard neighbourhood characteristics such as poverty rates; the remaining variation points to factors not yet captured by existing measures.'</i> (Chetty et al., 2014, paraphrased)	Use when the paper's analysis explicitly partitions variation across explanatory factors. Forces precision about what the paper does and does not explain. Problematic when decomposition shares do not add to a meaningful total or when residual variation is not discussed.
--------------------------------	--	--

The most important practical rule is that magnitude must always be described relative to something. A number stated in isolation — 'the rank-rank slope is 0.35' — is not a communicated finding; it is a reported parameter. The writer's job is to provide the context that makes the number meaningful: a range, a comparison, a translation into policy-relevant units, or an explicit statement of what the number implies for a specific real-world situation. Chetty et al. (2014) provide multiple contextualisations for their key findings — the Charlotte-to-Salt Lake City range, the comparison to Denmark, the implication for the American Dream — because different readers need different frames to evaluate the same number.

5.4 The Correlation–Causation Discipline

The most consequential writing discipline in quantitative economics is managing the boundary between correlation and causation. Getting it wrong in either direction is costly: understating a causal claim when the identification strategy supports it undersells the paper's contribution; overstating a causal claim when the analysis is correlational is one of the most reliable routes to a referee rejection and, in the worst cases, to retraction.

This discipline is not merely terminological — it is not simply a matter of replacing 'causes' with 'is associated with'. It is a substantive intellectual commitment about what the research design can establish, and it must be consistent throughout the paper. A paper that describes its design as correlational in the data section and draws causal conclusions in the discussion has contradicted itself. Careful referees will identify and penalise this.

◆ The causal language register

The table below sets out the principal causal language registers used in quantitative economics writing, with the conditions under which each is appropriate and the failure modes specific to each.

Causal language	Appropriate when...	Problematic when...
<i>'X causes Y' / 'the effect of X on Y'</i>	The identification strategy (RCT, IV, RDD, DiD) addresses selection and confounding and the paper has validated the identifying assumptions	The analysis is observational and no identification strategy has been stated; or the identifying assumptions have not been tested
<i>'X leads to Y' / 'X produces Y'</i>	The paper has quasi-experimental evidence of a directional effect, though not as clean as an RCT	Observational correlation only; the direction of causality is asserted but not established by the design
<i>'X is associated with Y' / 'X predicts Y'</i>	The analysis is correlational and the paper is explicit that causality is not claimed; the association is the finding	Used as a substitute for causal language when the writer intends a causal claim but wishes to avoid scrutiny — referees recognise this evasion
<i>'X is correlated with Y'</i>	Purely descriptive analysis where the co-movement of variables is itself the object of interest	Weaker than 'associated with'; appropriate only for descriptive statistics sections, not empirical findings
<i>'areas with high X tend to have high Y'</i>	Describing a pattern in cross-sectional data where no causal claim is made and heterogeneity across units is the subject	Appropriate and frequently used by Chetty et al. (2014) for their correlational findings; should not be used when the paper's actual identification strategy supports a stronger claim

Chetty et al. (2014) are unusually disciplined in their handling of this distinction. The paper's main analysis is correlational — it documents associations between CZ characteristics and mobility outcomes — and the authors are explicit about this throughout. They consistently use the fifth register in the table above ('areas with high X tend to have high Y') for their correlational findings, and they are explicit in the abstract that the paper does not establish causal mechanisms. The causal analysis is explicitly deferred to companion papers using the movers sample.

The paper's most instructive passage on this point appears in its treatment of the correlation between income inequality and mobility — a relationship the authors call the 'Great Gatsby Curve':

The finding of a negative correlation between inequality and mobility — areas with more inequality tend to have lower mobility — has attracted substantial attention. We wish to note that this correlation should not be interpreted as a causal relationship. The cross-sectional correlation may reflect common causes of both inequality and low mobility rather than a direct effect of inequality on mobility.

Chetty, R. et al. (2014). Where is the land of opportunity? *Quarterly Journal of Economics*, 129(4), p. 1596.

Explicit causal disclaimer: protecting a finding by bounding it correctly

The phrase 'we wish to note' signals that what follows is a deliberate methodological statement, not a passing observation. Chetty et al. (2014) have found a relationship that attracts attention precisely because it appears to support a causal interpretation that is politically salient — the claim that inequality itself causes low mobility. By explicitly stating that the correlation should not be read causally, they are not diminishing their finding; they are protecting it from overinterpretation that would make it vulnerable to criticism.

The final sentence — 'the cross-sectional correlation may reflect common causes of both inequality and low mobility' — is a specific causal alternative, not a vague disclaimer. It tells the reader what kind of confounding would need to be ruled out before the causal interpretation could be drawn, and it implicitly points to the additional empirical work that would be required. This is the model for how to handle a provocative correlational finding: report it accurately, protect it from overinterpretation, and indicate what further evidence would be needed.

◆ A contrasting case: Reinhart and Rogoff (2010)

To make these principles concrete, it is useful to compare Chetty et al. (2014) with a paper that made the opposite choices — and paid the consequences. Reinhart and Rogoff (2010), 'Growth in a Time of Debt', published in the *American Economic Review* Papers and

Proceedings, examined the relationship between public debt levels and economic growth across a large number of countries and historical periods. Its central finding — that countries with public debt above 90 per cent of GDP experience markedly lower growth — was widely cited in policy discussions and used to justify austerity policies in Europe and the United States. In 2013, Herndon, Ash, and Pollin (2014) published a replication that identified a series of errors in the original analysis and showed that the headline finding was not robust.

The Reinhart-Rogoff case is now a standard teaching example in economics methodology precisely because its failures are specific, teachable, and consequential. The table below identifies four of the most important, contrasting each with the approach demonstrated by Chetty et al. (2014).

Failure type	✗ Reinhart and Rogoff (2010)	✓ Chetty et al. (2014)
Causal language exceeding the design	The paper frames the 90% threshold as identifying a tipping point effect on growth, and policy audiences drew direct causal conclusions. The analysis is correlational. No identification strategy is provided that would separate the effect of debt on growth from the effect of slow growth on debt (reverse causality).	Explicit throughout that findings are correlational. The paper uses 'areas with high X tend to have low Y' language consistently. Causal claims are reserved for the companion movers analysis and explicitly labelled as such.
Selective data exclusion without justification	Several country-years were excluded from the dataset in ways that systematically affected the headline result. The exclusions were not disclosed in the paper and came to light only through Herndon et al.'s replication. The replication found that including the excluded observations substantially reduced the estimated growth differential.	Sample construction is documented step by step with observation counts at each stage. Exclusions are explained and justified. The representativeness of the analysis sample is checked against external benchmarks.

Inappropriate aggregation	Countries were weighted equally regardless of the number of years each contributed to the sample, so a single-year observation from one country counted as heavily as 19 years of data from another. This weighting is not justified in the paper and substantially drives the headline result.	Measurement decisions — including the choice of commuting zone as the geographic unit — are justified explicitly. The paper explains why this unit is more appropriate than the alternatives and what the alternatives would have cost analytically.
Magnitude communicated without context	The headline finding — that growth is roughly 1 percentage point lower in high-debt countries — is stated without a comparator that would allow the reader to evaluate its economic significance, and without a confidence interval that would allow its precision to be evaluated. Policy audiences read the finding as establishing a sharp threshold effect.	Findings are consistently contextualised: the Charlotte-to-Salt Lake City range provides the scale of the main finding; the comparison to Denmark provides the international benchmark; back-of-envelope calculations translate the rank-rank slope into dollar terms for specific family income scenarios.

What the Reinhart-Rogoff case teaches about quantitative writing

The failures identified in Reinhart and Rogoff (2010) are not, primarily, failures of statistical technique. They are failures of writing discipline — specifically, the failure to document choices clearly enough that a reader (or a replicator) could evaluate them. The selective data exclusions were consequential precisely because they were not disclosed; had the paper stated which country-years were excluded and why, the exclusions could have been evaluated on their merits.

The causal language failure is similarly a writing failure: the paper's correlational analysis was presented in a register that, combined with the policy urgency of the debt debate, invited causal interpretation that the data could not support. A paper that is explicit about the limits of its design — as Chetty et al. (2014) consistently are — provides its readers with the information they need to draw appropriate conclusions.

Note also what the Reinhart-Rogoff case illustrates about the stakes of quantitative writing. The paper was published in a prestigious venue and was technically sophisticated enough to pass peer review. Its errors were not obvious; they required a replication effort to identify. This is precisely why the writing disciplines described in this chapter matter: they are not just standards of academic good practice but protections against the kind of error that, at worst, influences policy with consequences for many people.

The full reference for the original paper and its replication is: Reinhart, C. M., and Rogoff, K. S. (2010). Growth in a time of debt. *American Economic Review: Papers and Proceedings*, 100(2), 573–578. The replication is: Herndon, T., Ash, M., and Pollin, R. (2014). Does high public debt consistently stifle economic growth? A critique of Reinhart and Rogoff. *Cambridge Journal of Economics*, 38(2), 257–279.

Tables and figures in a quantitative paper are not decorations or supplements; they are the primary vehicle for presenting evidence. A paper whose tables and figures are unclear, incomplete, or inconsistently formatted is a paper whose evidence cannot be properly evaluated — and a paper whose evidence cannot be evaluated is a paper that cannot be accepted.

The most important single principle for tables and figures is self-containment: every table and figure should be comprehensible without reference to the paper's text. A reader who encounters a table in isolation — as happens when tables are shared in presentations, cited in subsequent papers, or reproduced in policy documents — must be able to understand what the table shows, what data it uses, and what the column and row labels mean. Self-containment

requires complete labelling, explicit notes, and stated data sources. It is not difficult, but it requires consistency that many writers do not maintain.

Element	Standard practice	Common failures
Table title	States the content specifically: 'Rank–Rank Slopes by Commuting Zone: Estimates of Relative Intergenerational Mobility.' Not: 'Mobility Results.'	Generic titles ('Table 1: Results') that require reading the table to understand what it contains; titles that duplicate the section heading without adding specificity
Column and row headers	Self-explanatory labels with units stated (e.g. 'Mean Child Income Rank (percentile)'). Abbreviations defined in a note below the table.	Unlabelled columns; undefined abbreviations; units omitted; 'N' used without specifying what N counts
Table notes	Three types: (1) general note explaining sample and data sources; (2) specific notes flagging departures from the main specification; (3) significance stars defined (*, **, *** for $p < 0.10$, 0.05 , 0.01).	No notes at all; notes that repeat the table title; significance levels not defined; sample described in text but not in the table itself (making the table non-self-contained)
Figure axis labels	Both axes labelled with variable name and units. Scale chosen to display the full range of data without distortion.	Axes unlabelled; scale starting at a non-zero value to exaggerate trends; axis range truncated to exclude outliers without flagging this
Figure notes	Same three-type structure as table notes. Source of data stated. Any smoothing, binning, or transformation described.	Data source unstated; binning undisclosed; trend line shown without specifying the functional form used
Prose reference	Every table and figure explicitly referenced in the prose before or alongside its appearance: 'Figure II plots...' or '...(Table III, Column 2).' The prose states the finding; the table provides the evidence.	Tables and figures that are never referenced in the prose; prose that merely says 'see Table 3' without stating what Table 3 shows; tables and figures introduced only by footnote

◆ The prose-leads principle in practice

In a quantitative paper, the prose should always lead and the table or figure should provide the evidence. This means that the paragraph introducing a result must state the finding before or alongside the reference to the table — not after it, and not instead of it. The most common violation is the results paragraph that opens with a table reference ('Table 3 shows...') and then describes what each column contains, without ever stating what the table means. The reader arrives at the end of the paragraph knowing what numbers are in the table but not what to make of them.

✗ Table-led: describes without interpreting	✓ Prose-led: states finding, then directs to evidence
<i>Table II presents rank–rank slopes for children by commuting zone. Column 1 shows the slope for all children. Columns 2 and 3 show the slope separately for children in the bottom and top halves of the parent income distribution.</i>	<i>Relative mobility — measured by the rank–rank slope — varies substantially across commuting zones, with children in the Great Plains and Mountain West showing substantially lower intergenerational persistence than children in the South and parts of the Midwest (Table II, Column 1). The variation is not symmetric: the slope is steeper for families at the top of the income distribution, meaning that high-mobility areas disproportionately benefit children from lower-income families.</i>
<i>Figure III plots upward mobility against the Gini coefficient by commuting zone. The x-axis shows the Gini coefficient and the y-axis shows absolute upward mobility. A regression line is shown.</i>	<i>Areas with higher income inequality tend to have lower upward mobility — the scatter plot in Figure III confirms the negative cross-sectional correlation (slope: -0.57, $p < 0.001$). The relationship is, however, loose: the Gini coefficient alone explains less than a third of the cross-CZ variation in mobility, indicating that factors other than inequality are the primary drivers of geographic differences.</i>

5.6 Writing the Empirical Framework Section

The empirical framework section — sometimes called the econometric specification, empirical strategy, or methods section — performs a function that is frequently misunderstood: it does not describe what the authors did; it justifies why what they did is appropriate for the question they are asking. This distinction determines what belongs in the section and what does not.

What belongs: the statistical model and its connection to the economic question; the identifying variation and the assumptions it requires; the tests used to validate those assumptions; and the limits of what the design can establish. What does not belong: a description of regression software or computing procedures; a list of control variables without explanation of why each is included; a restatement of what the data section has already said.

◆ Stating the identifying variation

The most important sentence in the empirical framework section is the one that states the identifying variation: the source of variation in the data that is used to answer the research question and that is plausibly uncontaminated by selection or confounding. In an observational study, stating the identifying variation requires acknowledging that it may not be clean — that the analysis is correlational, that selection into treatment may be present, and that the estimates should be interpreted accordingly. In a quasi-experimental study, it requires explaining why the relevant variation is effectively random for the purposes of the analysis.

Chetty et al. (2014) are unusually careful about this. Their main analysis compares mobility across commuting zones using cross-sectional variation in CZ characteristics. They are explicit that this variation is not exogenous — some third factor may cause both low mobility and the CZ characteristics associated with it. The identifying variation for their causal claims comes from a separate analysis of families who move across CZs during childhood, where the age at which children move provides within-family variation that is more plausibly exogenous. The two analyses are kept separate, with different causal language applied to each.

We interpret the differences in upward mobility across CZs as reflecting differences in the causal effect of growing up in those areas on children's incomes in adulthood. This interpretation requires the assumption that families who move to higher-mobility areas do not differ from observationally similar non-movers in ways that would affect their children's outcomes. We evaluate this assumption using a set of placebo tests.

Chetty, R. et al. (2014). Where is the land of opportunity? *Quarterly Journal of Economics*, 129(4), p. 1579.

Stating the identifying assumption explicitly — and then testing it

The sentence 'this interpretation requires the assumption that...' is the correct template for stating an identifying assumption in empirical economics writing. It does three things: it names the interpretation being claimed (causal effect); it states the assumption under which that interpretation is valid; and it signals that the paper will provide evidence bearing on the assumption. The phrase 'we evaluate this assumption using a set of placebo tests' is not decorative — it tells the reader that the identification is not merely asserted but subjected to falsification tests.

The placebo test is the standard diagnostic for quasi-experimental identification: if the research design is valid, the estimated effect should be zero for groups or outcomes where no effect is theoretically predicted. Chetty et al. (2014) use children who move when they are old enough that the childhood exposure effects should be absent — if these older movers show the same estimated effects as younger movers, the design is likely picking up selection rather than genuine place effects. Stating that this test will be reported — before reporting it — is itself a signal of methodological confidence.

5.7 Robustness, Heterogeneity, and the Limits of the Analysis

Every quantitative paper's main results depend on a set of design choices whose justification, however careful, cannot be exhaustive. Robustness checks test whether the main findings survive reasonable variations in those choices. Heterogeneity analysis examines whether the findings are uniform across subgroups or whether they conceal substantial variation. Both are required in top-journal submissions, and both are frequently handled in ways that satisfy the formal requirement without providing genuine analytical content.

◆ Robustness checks that do and do not earn their place

A robustness check justifies its inclusion when it tests a specific, articulable concern about the main specification — a concern that a sceptical reader might reasonably raise. The appropriate way to introduce a robustness check is to state the concern before presenting the check: 'One concern with the baseline specification is that [X]; Table VI addresses this by [Y]. The results are qualitatively unchanged, though the coefficient falls from [A] to [B], suggesting that [concern] accounts for some but not all of the estimated effect.'

A robustness check that does not justify its inclusion is one that tests a variation that was never a plausible concern, or that reports 'results are robust' without stating what the results are. The phrase 'results are robust to alternative specifications' without a table or coefficient range

is not a robustness check; it is an assertion. Referees who cannot evaluate an assertion treat it as a request for further revision.

◆ **Heterogeneity as theory test**

As established in Chapter Four, heterogeneity analysis is most valuable when it is theory-motivated: when the economic model predicts that effects should be larger in some subgroups or conditions than in others, and when the heterogeneity analysis tests that prediction directly. Chetty et al. (2014) use this approach throughout. Their analysis of racial heterogeneity in mobility, for instance, is not simply a descriptive exercise — it is motivated by the theoretical question of whether neighbourhood effects operate differently by race, a question with direct implications for the mechanisms behind geographic variation in mobility.

The most instructive example of theory-motivated heterogeneity in Chetty et al. (2014) is their analysis of the five factors most strongly correlated with upward mobility: segregation, income inequality, primary school quality, social capital, and family structure. These five are not chosen arbitrarily from among all CZ characteristics; they are identified by their theoretical connection to the economic mechanisms most plausibly driving mobility differences. Presenting the analysis this way — theory first, then the five factors, then the correlational evidence — means that the heterogeneity analysis is structured as a contribution to understanding mechanisms, not as data-mining.

◆ **Stating the limits proportionately**

The limitations section of a quantitative paper should follow the same proportionality principle established for case studies in Chapter Four: depth of treatment proportional to consequentiality. A limitation that, if taken seriously, would substantially change the paper's main claims deserves explicit discussion of why the authors' approach is still defensible, or an acknowledgement that the claim should be qualified. A minor limitation can be noted in a sentence.

The most common failure in limitations discussions is the generic disclaimer: 'As with all cross-sectional studies, our findings should be interpreted with caution.' This statement is true of every cross-sectional study ever conducted, adds no information, and is not a limitations discussion. A genuine limitations discussion names the specific threat to validity in this paper, explains why it arises given the specific design, and states what additional evidence would be needed to address it.

5.8 Connecting Quantitative Findings to Policy

Quantitative economics papers frequently have policy implications, and stating those implications clearly is both an obligation to the policy community and a contribution to the paper's scholarly relevance. The challenge is the same as in case study writing: how to state implications that are specific enough to be useful without claiming more than the evidence supports.

Chetty et al. (2014) are widely cited in policy discussions of economic mobility in the United States, partly because of the quality of their evidence but also because their policy implications section is well crafted. They do not claim that their correlational findings identify policy levers — they are explicit that the associations they document have not been established as causal. What they do claim is that the geographic variation in mobility they document identifies places that are doing better and worse at generating upward mobility, and that this variation provides a starting point for more targeted policy intervention and further causal research. This is a defensible and useful policy implication that follows directly from correlational findings without overstating their causal status.

The appropriate model for connecting quantitative findings to policy is the conditional recommendation: 'if the correlation between [CZ characteristic] and mobility reflects a causal relationship — which our cross-sectional evidence cannot establish — then policies that [change the characteristic] may improve mobility outcomes.' The conditional is not timidity; it is precision. It tells the policy reader exactly what further evidence would be needed before the implication becomes a firm recommendation.

⚠ Five failure modes in quantitative writing — a pre-submission checklist

1. Undocumented measurement decisions — choices about outcome variables, sample construction, geographic units, or time periods are stated but not justified. Fix: for each significant decision, state what was chosen, why, and what the alternative would have cost.
2. Magnitudes without comparators — coefficients and estimates are reported without translation into interpretable terms or comparison to relevant benchmarks. Fix: every key finding should include a range, a comparison, or a back-of-envelope translation into units the reader can evaluate.
3. Causal language exceeding the design — 'X causes Y' language appears in papers whose design supports only 'X is associated with Y'. Fix: audit every causal claim against the identifying variation; replace causal language with association language where the design does not support causality.
4. Tables and figures without self-contained documentation — column headers are undefined, data sources are unstated, significance levels are not explained. Fix: every table and figure must be comprehensible without reference to the text; notes must include sample description, variable definitions, and significance conventions.
5. Generic robustness claims — 'results are robust to alternative specifications' without stated results or coefficient ranges. Fix: every robustness check must state the concern being addressed, the alternative specification tested, and the numerical result.

5.9 Vocabulary and Sentence Models

The vocabulary below covers the quantitative and methodological terms that appear most frequently in empirical economics writing. The sentence models demonstrate how to present measurements, report findings, and describe causal claims at the precise level that the design supports.

◆ **Key vocabulary**

Term	Meaning in quantitative writing	Example of the word in a sentence
operationalise	To define a concept in terms that can be measured with available data. Operationalisation is a key step between a theoretical concept and an empirical analysis, and it must be justified explicitly.	<i>'Intergenerational mobility is operationalised as the rank–rank slope — the coefficient from a regression of children's income rank on parents' income rank. This measure is chosen because it is robust to differences in the local income distribution and directly comparable across commuting zones.'</i>
coefficient	A number estimated from data that describes the relationship between two variables in a statistical model — for example, the change in Y associated with a one-unit change in X. A coefficient is not self-interpreting; the prose must state what it means in the units of the analysis.	<i>'The estimated coefficient on parent income rank is 0.35, meaning that children whose parents are 10 percentile points higher in the income distribution reach adult income ranks that are, on average, 3.5 percentile points higher.'</i>
standard error	A measure of the precision of an estimated coefficient — the degree to which the estimate would change across different samples drawn from the same population. Standard errors are the basis for confidence intervals and hypothesis tests.	<i>'The standard error on the mobility estimate is 0.012, implying a 95 per cent confidence interval of [0.33, 0.37] — a range narrow enough to rule out the possibility that the estimated variation across commuting zones reflects sampling noise rather than real differences.'</i>

fixed effects	Statistical controls for factors that are constant within a group (such as a firm, a region, or a time period) and that might otherwise confound the relationship being estimated. Fixed effects absorb all variation within the group, leaving only the variation between groups or over time.	<i>'City fixed effects absorb all time-invariant differences between cities — differences in geography, historical industry composition, and local institutions — so that the identification of the Airbnb effect relies only on variation within cities over time.'</i>
endogeneity	A problem that arises when an explanatory variable is correlated with the error term in a regression — typically because of omitted variables, reverse causality, or measurement error. Endogeneity means that ordinary regression estimates cannot be given a causal interpretation.	<i>'Neighbourhood characteristics are likely endogenous to mobility outcomes: families that value education may sort into high-mobility neighbourhoods, creating a positive correlation that does not reflect a causal effect of neighbourhood quality on children's outcomes.'</i>
instrument / instrumental variable	A variable that is correlated with the endogenous explanatory variable but affects the outcome only through that variable. Instrumental variables (IV) are used to identify causal effects in the presence of endogeneity.	<i>'The shift-share instrument isolates variation in local labour demand that comes from national industry trends rather than local conditions, providing a source of variation in employment that is plausibly exogenous to local wage-setting.'</i>
statistical significance	A statement that an estimated coefficient is unlikely to be zero purely by chance, given the data. Significance is conventionally judged at the 5 per cent level ($p < 0.05$), but significance is not the same as importance: a precisely estimated small effect is significant; a large effect estimated with low precision may not be.	<i>'The coefficient is statistically significant at the 1 per cent level ($p = 0.003$), but its magnitude — 0.02 standard deviations — is small relative to the cross-sectional variation in outcomes, suggesting that the variable is a real but modest contributor to the explained variation.'</i>

commuting zone (CZ)	A functional economic area based on commuting patterns, used as the geographic unit of analysis in several prominent labour economics and mobility studies. CZs better capture local labour market boundaries than counties or states.	<i>'The analysis uses commuting zones as the unit of observation because they capture functional economic areas — the markets in which workers actually compete for jobs — rather than administrative boundaries that may split or aggregate economically integrated areas.'</i>
----------------------------	--	--

◆ Sentence models

The models below cover the main moves in quantitative writing, from documenting measurement decisions to stating causal claims at the correct level of confidence.

Function	Model sentence	Notes for the writer
Justifying a measurement choice	<i>[Concept] is measured as [operational definition]. This measure is chosen because [advantage]; the main alternative — [alternative] — would [disadvantage or what it would miss].</i>	Every key measurement decision requires explicit justification. State the advantage of the chosen measure in terms of the research question, not in terms of data availability alone.
Translating a coefficient	<i>The estimated [coefficient] of [number] implies that [plain-language interpretation]: [a unit change in X is associated with / results in] a [magnitude] [unit] [change/increase/decrease] in [Y].</i>	Translate every key coefficient into interpretable terms. A coefficient of 0.35 says nothing by itself; a 3.5 percentile point difference for every 10-rank difference in parental income says something.

Communicating magnitude with a comparator	<i>[Finding] is [larger/smaller/comparable to] the [benchmark comparator] — [explanation of what the comparison implies].</i>	Magnitudes need benchmarks. Choose a comparator the reader can evaluate, and explain what the comparison means for the significance of the finding.
Stating an associational finding	<i>Areas with higher [X] tend to have higher [Y] (Figure [N]). This relationship is consistent with [mechanism], though the cross-sectional design cannot rule out that [omitted variable] drives both.</i>	Use associational language ('tend to', 'are associated with') when the design is correlational. Always note what the design cannot rule out.
Stating a causal finding	<i>We estimate that [treatment] [increases/reduces] [outcome] by [magnitude] ([unit]). This effect is identified from [source of variation], which is plausibly exogenous to [potential confounder] because [reason].</i>	Causal language ('increases', 'reduces') requires an identification strategy. State the source of identifying variation in the same sentence or paragraph that states the causal claim.
Describing a robustness check	<i>To address the concern that [specific threat to validity], we re-estimate the main specification [with/using/restricting to] [alternative]. The coefficient falls from [X] to [Y] — a [percentage] change that does not alter the substantive conclusion.</i>	Robustness checks must state the concern, the alternative, the result, and the implication. 'Results are robust' without numerical support is not a robustness check.
Reporting a null result	<i>[Outcome] does not respond detectably to [treatment] in [sub-group/specification] (Table [N], Column [N]; 95 per cent confidence interval [lower, upper]). This is consistent with [theoretical reason why the null is expected here].</i>	Null results are findings. Report confidence intervals so the reader can evaluate whether the null reflects a precise zero or an imprecise estimate.

Connecting findings to policy	<i>If [finding] reflects a causal relationship — which the [design/evidence] is consistent with but cannot fully establish — then [policy implication] may be effective in [context]. Confirming this would require [what further evidence is needed].</i>	Policy implications from correlational evidence must be conditional. The 'if' clause signals the causal status; the final sentence shows intellectual honesty about what remains to be established.
--------------------------------------	---	--

What makes Chetty et al. (2014) an instructive model is not the scale of the dataset or the sophistication of the econometrics; it is the discipline with which every number is translated into meaning, every causal claim is bounded by the design that supports it, and every finding is connected to a question that the reader already cares about. Quantitative competence is a precondition for writing a paper like this, but it is not sufficient. The writing itself — the measurement justifications, the magnitude descriptions, the correlation-causation boundary, the prose that leads the tables — is where the analysis becomes accessible, evaluable, and useful. That work is not a formality that follows the real work; it is part of the real work.

→ **Chapter Six: Citation, Attribution, and Academic Integrity**

CHAPTER SIX

Citation, Attribution, and Academic Integrity

Citation is the mechanism by which academic writing acknowledges its intellectual debts, positions itself within a scholarly conversation, and allows readers to verify the evidential basis of its claims. In economics, where empirical findings and theoretical arguments build on each other across decades of accumulated work, correct and complete citation is not a bureaucratic formality. It is part of the intellectual structure of the paper.

This chapter covers three things. First, the conventions of APA 7th edition citation as they apply specifically to economics writing — the in-text citation forms, the reference list format, and the handling of source types that economists frequently encounter but that general citation guides do not always cover well. Second, the question of when and what to cite — how to calibrate attribution across a paper so that the paper is neither under-cited (inviting charges of insufficient engagement with the literature) nor over-cited (a defensive reflex that substitutes footnotes for argument). Third, academic integrity: the spectrum of failures from sloppy attribution to fabrication, and the practical disciplines that prevent each.

6.1 APA 7th Edition: In-Text Citations

APA 7th edition is the citation standard used throughout this guide and adopted by a growing number of economics journals. Its in-text citation system uses the author-date format: the author's surname and the year of publication appear in parentheses immediately after the claim they support. This system has two forms — parenthetical, where both author and year appear in parentheses, and narrative, where the author's name appears as part of the sentence and only the year is parenthetical.

Parenthetical: 'The natural rate of unemployment is the rate that would be ground out by the Walrasian system of general equilibrium equations (Friedman, 1968, p. 8).'

Narrative: 'Friedman (1968) argued that monetary policy cannot permanently lower the unemployment rate below its natural level.'

Both forms are acceptable in economics. The choice is typically guided by whether the author's identity is important for the claim being made. If the claim matters because of who made it — because Friedman's authority is relevant, or because the paper is discussing Friedman's argument specifically — the narrative form is more appropriate. If the claim is a standard empirical finding for which the author's identity is incidental, the parenthetical form is more efficient.

◆ In-text citation quick reference

The table below covers the most common in-text citation situations encountered in economics writing. All forms are consistent with APA 7th edition.

Situation	In-text form	Notes
One author	<i>(Friedman, 1968) or Friedman (1968)</i>	Either parenthetical or narrative form is acceptable; choose one style and apply it consistently throughout the paper
Two authors	<i>(Farronato & Fradkin, 2022) or Farronato and Fradkin (2022)</i>	Use '&' inside parentheses; spell out 'and' in the running text
Three or more authors	<i>(Chetty et al., 2014) or Chetty et al. (2014)</i>	Use 'et al.' from the first citation for three or more authors
Direct quotation	<i>(Friedman, 1968, p. 9) or (Chetty et al., 2014, pp. 1560–1561)</i>	Page number required for all direct quotations; use 'p.' for single pages, 'pp.' for a range
Same author, same year	<i>(Chetty et al., 2014a) and (Chetty et al., 2014b)</i>	Add 'a', 'b', 'c' to year; the reference list must use the same suffixes
Work cited in another work	<i>(Okun, 1962, as cited in Friedman, 1968)</i>	Cite the source you actually read, not the original; use sparingly — locate the primary source where possible
Multiple citations in one parenthesis	<i>(Chetty et al., 2014; Goldfarb & Tucker, 2019)</i>	Order alphabetically by first author surname; separate with semicolons

Two common errors are worth noting explicitly. The first is treating 'et al.' as a substitute for full author information in the reference list: 'et al.' is used in the in-text citation, but the reference list entry must give all authors' surnames and initials. The second is inconsistency in punctuation: inside parentheses, use '&' between two authors; in the running text, use 'and'. Inconsistency is a mark of inattention that editorial assistants and referees notice.

6.2 APA 7th Edition: The Reference List

The reference list appears at the end of the paper and contains full bibliographic information for every source cited in the text. It must be: complete (every in-text citation has a

reference list entry and vice versa); consistent (the same formatting conventions applied throughout); and accurate (author names, titles, years, volume numbers, and DOIs are verified against the primary source).

◆ Journal articles — the core format

The standard APA 7th format for a journal article is:

Author, A. A., and Author, B. B. (Year). Title of article in sentence case. Journal Name in Title Case and Italics, volume(issue), first page–last page. <https://doi.org/xxxxx>

APA 7th edition — journal article format

Applied to the papers used in this guide:

Friedman, M. (1968). The role of monetary policy. *American Economic Review*, 58(1), 1–17.

Goldfarb, A., and Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1), 3–43. <https://doi.org/10.1257/jel.20171452>

Farronato, C., and Fradkin, A. (2022). The welfare effects of peer entry: The case of Airbnb and the accommodation industry. *American Economic Review*, 112(6), 1782–1817. <https://doi.org/10.1257/aer.20180260>

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where is the land of opportunity? The geography of intergenerational mobility in the United States. *Quarterly Journal of Economics*, 129(4), 1553–1623. <https://doi.org/10.1093/qje/qju022>

Several formatting points follow directly from these examples. Author surnames come first, followed by initials — not full first names. Multiple authors are separated by commas, with 'and' before the final author (not '&', which is reserved for in-text parenthetical citations). Article titles are in sentence case: only the first word, proper nouns, and the first word after a colon are capitalised. Journal names are in title case and italicised; volume numbers are italicised; issue numbers are not. The DOI is given as a URL (<https://doi.org/...>), not as a label ('DOI: ...'). When a DOI is available, it must be included.

◆ Non-standard sources in economics writing

Economists draw on a wider range of source types than many disciplines: NBER working papers, datasets, government reports, conference proceedings, and institutional publications all appear regularly in economics reference lists. The table below sets out the APA 7th format for each, with specific guidance on its appropriate use in an economics paper.

Source type	APA 7th reference list format	Guidance
NBER working paper	<i>Chetty, R., Hendren, N., Kline, P., and Saez, E. (2013). The economic impacts of tax expenditures (Working Paper No. 19843). National Bureau of Economic Research. https://doi.org/10.3386/w19843</i>	Include working paper number and NBER DOI where available. If subsequently published, cite the journal version and note the NBER WP number in a parenthetical.
Published dataset	<i>Chetty, R., Friedman, J., Hendren, N., Jones, M., and Porter, S. (2018). The Opportunity Atlas [Data set]. Opportunity Insights. https://opportunityinsights.org/paper/the-opportunity-atlas/</i>	Treat as a work with authors and a title. Include the label [Data set] in square brackets after the title. Provide the retrieval URL or repository DOI.

Government or institutional report	<i>World Bank. (2022). World development report 2022: Finance for an equitable recovery. World Bank Group. https://doi.org/10.1596/978-1-4648-1730-4</i>	Use the institutional author where no individual authors are named. Government reports and central bank working papers follow the same pattern as institutional publications.
Conference proceedings paper	<i>Chala, N., and Poplavska, O. (2025). Resilient teams and global megatrends. In Proceedings of the 7th International Conference on Economics and Social Sciences (pp. 45–46). WSHIU. https://ekmair.ukma.edu.ua/handle/123456789/38062</i>	Include conference name, page numbers, and publisher. Note: conference proceedings papers carry less evidential weight than peer-reviewed journal articles; use accordingly (see Chapter Three).
Preprint (not yet peer-reviewed)	<i>Author, A. A. (Year). Title of preprint. Repository Name. https://doi.org/xxxxx</i>	Always indicate preprint status; do not cite as if peer-reviewed. If the paper is subsequently published, update the citation to the journal version.

Policy or news source used for context	<i>European Central Bank. (2023, September). Economic bulletin, issue 6/2023.</i> https://www.ecb.europa.eu/pub/economic-bulletin/html/eb202306.en.html	Policy and journalistic sources may be cited for institutional context or aggregate data but should not be used as evidence for theoretical or causal claims (see section 6.2).
---	---	--

◆ Reference list mechanics: a checklist

Before submitting: verify your reference list

- ☐ Every in-text citation has a matching reference list entry — including self-citations.
- ☐ Every reference list entry is cited at least once in the text.
- ☐ Author names match exactly between in-text citations and the reference list.
- ☐ Article titles are in sentence case; journal names are in title case and italicised.
- ☐ Volume numbers are italicised; issue numbers are not.
- ☐ DOIs are formatted as <https://doi.org/...> URLs (not as labels).
- ☐ For papers with four or more authors: the reference list gives all surnames and initials; 'et al.' is used only in the in-text citation.
- ☐ Working papers that have since been published are cited in their published journal form.
- ☐ The reference list is ordered alphabetically by first author surname, then by year for the same author.
- ☐ For the same author in the same year: suffixes a, b, c are added to the year, consistently in both in-text citations and the reference list.

6.3 When and What to Cite

Citation is not a mechanical process of adding a reference whenever another work is consulted. It is a substantive decision about intellectual provenance: which works are the sources of the ideas, findings, methods, or data being used, and which works are relevant enough to the argument to belong in the paper's scholarly conversation. Getting this calibration right — neither too sparse nor too dense — is a mark of disciplinary maturity.

◆ What always requires a citation

Some things always require a citation, without exception: a direct quotation (with page number); a specific empirical finding from another paper; a theoretical result attributed to a named economist or paper; a dataset that was collected or constructed by others; a method or identification strategy first proposed in another paper; and a prior claim that the paper is challenging, extending, or building on. These are non-negotiable. Failing to cite any of these is, at minimum, sloppy attribution and, at worst, plagiarism.

A particularly common error in submitted manuscripts is the uncited stylised fact. A statement such as 'intergenerational income mobility has declined substantially in the United States over the past 50 years' is an empirical claim with a specific evidential basis. It is not general knowledge, and it requires a citation to the work that established it — in this case, to Chetty et al. (2014) and the companion papers that document the trend. Presenting specific empirical findings as if they were general knowledge is a form of under-citation that erodes the paper's evidential accountability.

◆ What does not require a citation

Not everything requires a citation. Standard results from microeconomic or macroeconomic theory — that demand curves slope downward, that competitive equilibria are Pareto-efficient, that monetary expansion is inflationary in the long run — are part of the shared technical knowledge of the discipline and do not need to be sourced. The test is whether a well-trained economist would recognise the claim as a standard result. If they would, no citation is needed. If the claim is specific, contested, or recently established, a citation is required.

The description of the paper's own analysis never requires a citation to prior work, unless a specific method or dataset from elsewhere is being used. 'We estimate a linear probability model' does not need a citation; 'We use the bartik shift-share instrument proposed by Goldsmith-Pinkham, Sorkin, and Swift (2020)' does.

◆ Calibrating citation density

Reference density — the number of citations per unit of text — varies across economics subfields and journal conventions. Applied microeconomics papers tend to have denser citation practices than theory papers; empirical literature reviews (like Goldfarb & Tucker, 2019) are more densely cited than single-paper empirical studies. Calibrating correctly is a matter of judgement guided by two questions: does this citation tell the reader something they need to know about the claim being made; and would a specialist referee notice its absence?

Over-citation — citing three or four papers for a claim that one paper would support — is not a virtue. It signals a writer who is uncertain whether any single citation is adequate and is hedging by accumulation. It also buries the most important references in a list that gives them no more prominence than peripheral ones. The discipline is to cite the most important and most direct source, and to add others only when they contribute something the first citation does not: a different method, a different context, a more recent update, or a directly contradictory result.

Under-citation is the more common failure and the more consequential one. A paper that makes strong empirical claims without supporting citations, or that positions itself as novel without engaging the relevant prior work, will be rejected on that basis alone. The literature review functions partly as a display of citation competence: referees check whether the paper knows the field, and gaps in the citation record are read as evidence that it does not.

6.4 Quoting, Paraphrasing, and Summarising

Academic writing in economics involves three ways of using material from prior sources: direct quotation, paraphrasing, and summarising. Each has a distinct function and distinct conventions, and choosing among them is a substantive decision rather than a stylistic preference.

◆ Direct quotation

A direct quotation reproduces the exact words of a source and must appear in quotation marks with a page citation. Direct quotation is appropriate in economics writing when the exact wording matters — when the claim being quoted is contested and precision about what was actually said is necessary; when the source's phrasing is unusually precise or memorable; or when the quotation is itself the object of analysis. It is not appropriate as a substitute for paraphrasing, and it should not be used simply because the source said something well.

The most common misuse of direct quotation in economics manuscripts is the block quotation deployed where a paraphrase would serve better. A lengthy excerpt from a prior paper, dropped into the text without analysis, shifts intellectual labour from the writer to the reader: the reader must do the interpretive work of connecting the quotation to the argument. Quotations that are not explicitly analysed in the surrounding prose are not contributing to the argument; they are filling space.

◆ Paraphrasing

Paraphrasing restates a source's idea in the writer's own words. It still requires a citation — the idea originated with the source, even if the wording is new — but it does not require quotation marks or a page number (though adding a page number is good practice when the original is long and the paraphrase corresponds to a specific passage). Paraphrasing is the dominant mode of source use in economics writing, and it should be: it forces the writer to understand the source well enough to express it independently, and it integrates the source's contribution into the paper's own voice rather than interrupting it.

A paraphrase that closely mirrors the original's sentence structure or specific phrasing without quotation marks is not a paraphrase; it is disguised quotation, and it constitutes a form of plagiarism even when the source is cited. The test is whether a reader who did not have access to the original could distinguish the writer's voice from the source's. If they could not, the paraphrase is too close.

◆ Summarising

A summary compresses a larger body of material — a paper, a section, or a set of related findings — into a brief statement. Summaries are appropriate when the details of the source's argument or evidence are less important than its general conclusion, and when a single paragraph of the original would otherwise require two or three pages of paraphrase. Like paraphrasing, summarising requires a citation. Like paraphrasing, it should not reproduce the original's structure or phrasing.

In literature reviews, summarising is the mode of engagement with sources that are relevant but not central. The sources that deserve paraphrase — or even selective direct quotation — are those whose specific arguments, methods, or findings are directly engaged by the paper. The majority of cited sources in a literature review will be summarised in a sentence or two and integrated into a synthesis move (see Chapter Three). This is appropriate: the synthesis, not the summary, is where the literature review's intellectual contribution is located.

6.5 Academic Integrity: The Spectrum of Failures

Academic integrity is not a topic that most researchers expect to encounter personally. The failures described below — fabrication, falsification, plagiarism — are widely understood as belonging to pathological cases. The evidence, however, suggests that several of the failures on the spectrum are considerably more common than the field's self-image acknowledges. Systematic replication studies in economics and adjacent disciplines have found that a significant proportion of published results are not robustly replicable, and that failures of documentation, disclosure, and attribution are widespread rather than exceptional.

Understanding the spectrum of integrity failures — and the practical disciplines that prevent each — is therefore not a matter of learning where the obvious lines are. It is a matter of understanding how ordinary research practices, under competitive pressure, can shade into the less serious failures without a deliberate decision to cross a line. The table below sets out the spectrum from the most to the least serious, with guidance on detection and prevention for each.

Failure type	What it involves	How it is detected and its consequences
--------------	------------------	---

Fabrication	Inventing data, results, or citations that do not exist. This includes citing papers that were never written, or reporting results from an analysis that was never conducted.	Detected through replication attempts, data audits, and citation verification. Consequences include retraction, institutional investigation, and permanent damage to professional standing. The most serious form of academic misconduct.
Falsification	Manipulating real data, results, or methods to misrepresent the findings. This includes selective exclusion of inconvenient observations without disclosure (as in Reinhart & Rogoff, 2010), and post-hoc adjustment of analysis to improve results.	Detected through replication and data sharing requirements. Top journals increasingly require data and code deposits. Selective exclusion is the most common form and is often not deliberate — which is why documentation discipline (Chapter Five) is the primary prevention.
Plagiarism	Using another's words, ideas, or data without attribution. In academic writing, this includes uncited paraphrasing — closely following another author's argument structure or phrasing without acknowledgement, even if individual words are changed.	Detected by editors using text-matching software and by specialist referees who recognise their own work. Consequences range from paper rejection to investigation depending on severity. The safest practice is to cite generously and paraphrase substantially.
Self-plagiarism	Republishing material from one's own previous work without disclosure. This includes submitting substantially overlapping papers to multiple journals simultaneously, or reproducing passages from published papers without indicating that they appear elsewhere.	Detected through text-matching software and by editors familiar with the author's prior work. Journals increasingly require disclosure of related prior work at submission. The appropriate practice is to cite your own prior work and to explain, where relevant, how the new paper extends it.

Ghost and gift authorship	Ghost authorship omits a contributor who made a significant intellectual contribution. Gift authorship adds an author who did not. Both misrepresent the intellectual provenance of the work.	Increasingly detected through author contribution statements required by journals. The standard is that authorship should reflect a genuine intellectual contribution to the conception, design, analysis, or writing of the paper — not seniority, administrative role, or lab membership alone.
Sloppy attribution	Citing inaccurately: misquoting, misattributing a finding to the wrong paper, citing a secondary source as if it were the primary, or listing a paper in the reference list that is not cited in the text. Less serious than the failures above, but corrosive to the paper's credibility.	Detected by referees and readers who check citations. The prevention is simple: verify every citation against the primary source before submission, and run a final check that every in-text citation has a reference list entry and vice versa.

◆ Self-plagiarism: a specific discipline for faculty

Self-plagiarism deserves particular attention in the context of faculty development, because it arises in patterns of work that are specific to productive academic researchers. A researcher who has published a series of papers on a related topic will naturally have written introductions, literature reviews, and data descriptions that cover similar ground. The temptation to reuse this material — from an unpublished working paper, a conference proceedings contribution, or a prior published paper — is understandable and, in some limited forms, acceptable.

The boundary is disclosure. Reusing a paragraph of background material from a prior paper without indication that it appears elsewhere is a form of self-plagiarism; doing so with a footnote that reads 'this section draws on material from Smith (2022), the author's prior work on related themes' is not. The principle is that the reader — and the editor — should be able to form an accurate picture of what is new in the submitted paper and what has appeared before. Journals that discover substantial overlap with a prior published paper will retract or decline the new submission; journals that are informed at submission have the information to make an editorial judgement.

A specific application of this principle is the conference paper that is developed into a journal article. It is entirely standard, and expected, that a conference contribution will grow into a full paper. The journal submission should cite the conference paper in the text ('this paper

develops the analysis presented in Chala and Poplavska, 2025') and should not reproduce substantial passages from the conference paper without acknowledgement. The two-page conference contribution that becomes a twelve-page journal article poses no integrity concerns as long as the relationship between them is disclosed.

◆ **Citing AI-generated material**

The use of large language models as writing aids raises citation questions that existing guidelines have not yet resolved consistently. The emerging consensus, reflected in the author guidelines of a growing number of journals, is that AI-generated text is not citable as a source: it has no author, no institutional accountability, no peer review, and no stable version that can be retrieved. Text generated by an AI system and included in a paper without disclosure raises the same concerns as ghostwriting.

Where AI tools are used in the preparation of a paper — for language editing, for generating initial drafts that are subsequently rewritten, or for summarising prior literature — this use should be disclosed in the acknowledgements, with a brief description of how the tool was used. The intellectual responsibility for the paper's claims, arguments, and evidence remains entirely with the named authors; AI use does not change this, but concealing it is a form of misrepresentation about the paper's provenance.

6.6 Open Science and Reproducibility

The norms of academic publication in economics have shifted substantially over the past fifteen years towards greater openness about data, code, and methods. Understanding these norms — and meeting them — is now part of what editors and referees expect from a submission to a leading economics journal. It is also, less obviously, a form of scholarly communication: depositing data and code is part of the argument that the paper's findings are reliable.

◆ **Data and code availability**

The American Economic Association and most leading economics journals now require authors to deposit the data and code necessary to replicate the paper's results as a condition of publication. The American Economic Review, the Quarterly Journal of Economics, the Review of Economic Studies, and the Journal of Political Economy all have explicit data availability policies, and the AEA Data Editor conducts pre-publication replication checks for all AEA journals. Meeting these requirements is not optional for journal submission, and preparing for them from the beginning of a research project — by keeping clean, documented code and maintaining a clear chain from raw data to final tables — is considerably less work than reconstructing this documentation at the revision stage.

For faculty working with proprietary or administrative data — as Chetty et al. (2014) did with federal tax records — full data deposit is often not possible. In these cases, journals require instead a detailed description of the data and its sources, a statement of the access requirements,

and, where feasible, the deposit of the analysis code applied to the data. Some journals accept synthetic or simulated data that preserve the statistical properties of the original for replication purposes. The key obligation is transparency: the reader should be able to understand exactly what data were used and, in principle, how to obtain access to data of the same type.

◆ **Pre-registration and registered reports**

Pre-registration — the practice of publicly committing to a research design, hypotheses, and analysis plan before data collection or analysis begins — is standard in experimental economics and clinical research and is becoming more common in applied microeconomics. Pre-registration does not prevent exploratory analysis; it distinguishes confirmatory tests (specified in advance) from exploratory findings (discovered in the data), so that the reader can assess the evidential weight of each category of result.

Some journals now offer registered reports: a format in which the research design is peer-reviewed and provisionally accepted before the results are known. If the study is conducted as pre-registered, acceptance of the final paper does not depend on the direction or significance of the findings. This format directly addresses the publication bias problem — the tendency for significant results to be published and null results to be filed away — and is increasingly valued by editors who are aware of the problem.

Faculty at the beginning of a research programme do not need to pre-register every project. But understanding the rationale for pre-registration — that it provides a credible commitment against the unconscious optimisation of analysis towards significant results — is important for understanding why referee reports sometimes ask about the relationship between the paper's stated hypotheses and its reported results, and why transparency about the history of specification choices is valued.

⚠ **Pre-submission citation and integrity checklist**

- ☐ Every direct quotation is in quotation marks and has a page number citation.
- ☐ Every specific empirical finding from another paper is cited to its source.
- ☐ Every dataset, method, or identification strategy from another paper is cited.
- ☐ Every in-text citation has a reference list entry; every reference list entry is cited in the text.
- ☐ Author names, years, titles, volume numbers, and DOIs are verified against the primary source.
- ☐ Working papers that have since been published are cited in their published journal form.
- ☐ Overlap with the author's prior published work is disclosed in the submission letter.
- ☐ Data and code availability requirements of the target journal have been checked and met.
- ☐ Any use of AI writing tools is disclosed in the acknowledgements.
- ☐ The reference list is complete, consistently formatted, and alphabetically ordered.

6.7 Vocabulary and Sentence Models

The vocabulary below covers the terms specific to citation practice and academic integrity as they arise in economics publishing. The sentence models demonstrate how to attribute, quote, paraphrase, and disclose at the level of precision that journals expect.

◆ Key vocabulary

Term	Meaning in citation and integrity writing	Example of the word in a sentence
attribution	The act of identifying the source of an idea, finding, or data used in a paper. Attribution is required whenever a claim derives from another person's work, whether quoted directly or paraphrased. Failure to attribute is plagiarism.	<i>'Proper attribution does not mean citing every paper that touches the topic; it means identifying the specific work from which each claim, finding, or method derives — and doing so with enough precision that the reader can locate and verify the source.'</i>
paraphrase	A restatement of another author's idea in different words. Paraphrasing still requires a citation. A paraphrase that closely follows the original's sentence structure or phrasing, even without quotation marks, is not acceptable.	<i>'A good paraphrase does more than change individual words: it reconstructs the idea in the writer's own syntactic structure and vocabulary, so that the reader is encountering the writer's understanding of the source, not a verbal disguise of the original text.'</i>

in-text citation	A brief reference to a source that appears within the body of the paper, typically as an author-date parenthetical (Friedman, 1968) or a narrative form (Friedman, 1968, argued that...). The in-text citation connects the claim in the text to the full reference in the reference list.	<i>'Every specific empirical finding or theoretical claim drawn from another work requires an in-text citation placed immediately after the claim — not at the end of a paragraph that contains multiple claims from different sources.'</i>
reference list	The complete bibliography at the end of the paper, containing full bibliographic information for every source cited in the text. In APA 7th edition, the reference list is ordered alphabetically by first author surname.	<i>'The reference list must be complete and accurate: every in-text citation must correspond to an entry in the reference list, and every entry must be verifiable against the primary source. Inaccurate references undermine the paper's credibility.'</i>
DOI	Digital Object Identifier — a permanent URL used to identify and access scholarly publications online. APA 7th edition requires DOIs where available, formatted as https://doi.org/... rather than as a label ('DOI: ...').	<i>'The paper is available at https://doi.org/10.1093/qje/qju022 — the DOI provides a permanent link that does not depend on journal website structures remaining stable.'</i>
sentence case	A capitalisation convention in which only the first word of a title, proper nouns, and the first word after a colon are capitalised. APA 7th edition requires sentence case for article and book chapter titles in the reference list.	<i>'Article titles in the APA reference list use sentence case: the title of Friedman's paper is written as <i>The role of monetary policy</i>, not <i>The Role of Monetary Policy</i> — though the journal name is written in title case: <i>American Economic Review</i>.'</i>

pre-registration	A public commitment to a research design, hypotheses, and analysis plan made before data collection or analysis begins. Pre-registration creates a record against which the published paper can be compared, reducing the scope for undisclosed post-hoc adjustment of the analysis.	<i>'Pre-registration of the analysis plan before the dataset was assembled provides a credible commitment that the paper's hypotheses were not chosen after inspecting the data — a guarantee that referee reports increasingly look for in experimental and quasi-experimental work.'</i>
self-plagiarism	The republication of one's own previously published material without disclosure. Self-plagiarism includes submitting overlapping manuscripts to multiple journals simultaneously and reproducing substantial passages from prior published work without acknowledgement.	<i>'This paper develops the argument first presented in the author's conference contribution (Chala and Poplavska, 2025). The present version extends that analysis by including panel data on three additional regions and adding the welfare decomposition reported in Section 4.'</i>

◆ Sentence models

The models below cover attribution, quotation, paraphrasing, and disclosure. They are the practical tools for meeting the citation and integrity expectations described in this chapter.

Function	Model sentence	Notes for the writer
Narrative in-text citation	<i>[Author(s), year] [showed / argued / estimated / demonstrated] that [specific finding or claim].</i>	Use narrative form when the author's identity is relevant to the claim — when the contribution is associated with a named economist whose authority matters to the argument.

Parenthetical in-text citation	<i>[Claim] ([Author(s), year, p./pp. page number if quoting]).</i>	Use parenthetical form when the source is supporting a claim but the author's identity is not the focus. Always add a page number for direct quotations.
Citing multiple sources	<i>[Claim] ([Author A, year]; [Author B, year]; [Author C, year]).</i>	Order multiple citations alphabetically by first author surname. Use when multiple sources support the same claim — but cite the most direct source first where one is clearly stronger.
Paraphrasing a source	<i>According to [Author(s), year], [idea restated in the writer's own words and sentence structure].</i>	The 'according to' formula signals that the idea is not the writer's own. The paraphrase that follows must be genuinely rewritten — not a word-for-word copy with pronouns changed.
Quoting directly	<i>As [Author] ([year], p. [N]) [writes/argues/states], '[exact words in quotation marks]'. This [matters/is significant] because [explanation].</i>	Direct quotations must be followed by analytical engagement — not just presented and left. The sentence after the quote should explain why the exact wording matters.
Citing prior own work	<i>This paper extends the analysis in [Author (year)], which [brief description of prior work]. The present version [adds/revises/extends] by [what is new].</i>	Cite your own prior work as you would any other — with a full citation. Add a brief description of how the present paper differs from or extends the prior one.
Disclosing AI tool use	<i>The authors used [tool name] to [specific function: assist with language editing / generate an initial draft structure / summarise background literature]. All intellectual content, analysis, and conclusions are the authors' own.</i>	Disclose AI use in the acknowledgements or the methods section. Be specific about the function. The disclosure does not diminish the paper — concealment does.

Disclosing data availability	<i>The data used in this analysis are [available from / restricted by / subject to access conditions at] [source or repository]. The analysis code is deposited at [URL/repository]. Researchers seeking access to the underlying data should [contact / apply through] [institution].</i>	Data availability statements should be specific: say what is available, where it is, and what access conditions apply. Vague statements ('data available on request') are increasingly insufficient for top journals.
-------------------------------------	--	---

Citation is not a bureaucratic obligation that surrounds the real work of research. It is the mechanism by which scholarship is held accountable — to the work it builds on, to the evidence it uses, and to the readers who need to verify its claims. A paper that cites correctly, quotes honestly, and discloses openly is a paper that trusts its readers with the information they need to evaluate it. That trust is the foundation of the discipline's ability to accumulate knowledge.

→ **Chapter Seven: The Writing Process**

CHAPTER SEVEN

The Writing Process

The preceding six chapters have addressed what to write: argument structure, the genres of economics publication, literature synthesis, case study construction, quantitative writing, and citation. This chapter addresses how to write — the practical process by which a research project becomes a submittable paper.

The writing process in academic economics is not a linear sequence that begins once the analysis is complete and ends at submission. It is iterative, recursive, and interwoven with the research itself. Writing forces the researcher to be precise about what they know, what they have established, and what they are claiming — and this precision frequently reveals gaps in the analysis that were invisible before the writing began. The researcher who writes only at the end of the project, once all the results are in, is forgoing one of the most powerful diagnostic tools available. Writing is not the transcription of a completed analysis; it is part of the process by which the analysis is completed.

This chapter covers the planning, drafting, and revision of an economics paper; the process of seeking and using feedback at each stage; and the specific challenges of academic writing in English as an additional language. Its recommendations apply primarily to research papers and literature reviews, but the principles extend to any extended piece of academic writing.

7.1 Planning and Sequencing the Draft

Before writing the first sentence of a draft, it is worth establishing three things: the research question the paper will answer, the claim it will make in response, and the evidence that supports that claim. These three elements — question, claim, evidence — are the skeleton of the paper. Everything else, including the literature review, the methodology section, the tables, and the discussion, is there to establish that the claim is credible and that the question is worth asking. If the skeleton is unclear before drafting begins, the draft will be unclear; the writer will be discovering the argument as they write rather than communicating it.

◆ The one-paragraph plan

A useful planning discipline before drafting is to write a single paragraph that answers three questions: What is the paper's research question? What is its answer? Why does the answer matter? This paragraph need not appear in the paper — indeed, it usually should not, because the rhetorical conventions of economics writing distribute this information across the introduction in a specific way (see Chapter Two, section 2.3). Its purpose is to force the writer to commit to a clear claim before committing 8,000 words to supporting one.

A research paper without a clear one-paragraph answer to these three questions is not ready to be drafted. The question may be interesting. The data may be collected. The regressions

may be run. But if the writer cannot state, in plain prose, what the paper shows and why it matters, the draft will not be coherent. Drafting under these conditions produces a common failure mode: a paper that describes what was done rather than what was found, and that accumulates evidence without ever arriving at a claim.

◆ Sequencing: which sections to write first

The introduction is the last section to write, not the first. The introduction makes promises about what the paper will establish; those promises can only be stated accurately once the paper's argument is complete. Writing the introduction first — a common instinct, because it comes first in the paper — typically produces introductions that must be substantially rewritten at the end, because the paper that was planned is not the paper that emerged from the analysis.

A more productive sequence for most writers is: (1) the data and methodology section, because it is largely descriptive and forces early precision about the measurement decisions the paper depends on; (2) the results section, where the core findings are first assembled into prose; (3) the discussion section, where the findings are interpreted and connected to the literature; (4) the literature review, once the paper's contribution is clear enough to position it accurately; (5) the conclusion, which answers the research question directly; (6) the introduction, which is now written to match the paper that exists rather than the paper that was planned.

Abstract last. The abstract is a compressed version of the complete paper and can only be written accurately once the paper is complete. It is not a plan; it is a summary.

7.2 The First Draft: Speed Over Perfection

The most counterproductive approach to drafting is to write slowly, editing each sentence as it is produced. This approach conflates two distinct cognitive activities — generating ideas and evaluating them — that are better kept separate. A writer who edits as they draft will produce more polished individual sentences but a weaker overall argument, because the editorial mind interrupts the generative one before the argument has had time to develop.

The purpose of the first draft is to have a complete text — a text from which something can be made. A first draft that is rough, that has placeholder sentences ('DEVELOP THIS'), that has sections in the wrong order, that has arguments stated too bluntly and evidence described too briefly, is a more useful starting point than a half-finished draft whose existing paragraphs are perfectly written. The revision process transforms the complete rough draft into the final paper. There is no equivalent process for transforming a half-finished polished draft into a complete one.

◆ The blank page problem

Many writers find the transition from planning to writing difficult — the moment of sitting down to produce the first sentence of a section. Several strategies help. One is to write towards the draft rather than for it: before opening the document, write a paragraph of notes in plain language about what the section needs to establish and how it will do so. This paragraph,

written without the pressure of producing publishable prose, frequently becomes the section's first draft with light editing.

Another is to start in the middle — to draft the section or part of the section that the writer is most confident about, rather than the section that comes first. Economics papers do not need to be drafted in the order they will be read. The results section, for many writers, is the easiest to start with because it is closest to the analysis. Starting there breaks the blank page paralysis and generates material that the other sections can then be written around.

◆ **Writing and revising are different activities**

The first draft should be produced as quickly as the writer's understanding allows, with revision postponed to a separate stage. This does not mean writing carelessly; it means not allowing the care that revision requires to interrupt the momentum that drafting requires. A first draft produced in three days and then revised carefully over two weeks will usually be better than one produced slowly over four weeks without a clear separation between drafting and revision.

The separation between drafting and revision also helps with a specific problem that affects writers in English as an additional language: the temptation to stop and look up vocabulary during the draft. In the drafting stage, it is better to write a placeholder ('WORD FOR THIS IN ENGLISH?') or to use the first-language equivalent, marked for later, than to interrupt the drafting to search for the correct English term. The revision stage, conducted without time pressure, is the right place for this vocabulary work.

7.3 Revision: Argument, Structure, and Sentence Level

Revision is not proofreading. Proofreading is the correction of mechanical errors — spelling, punctuation, formatting — and takes place at the final stage of the revision process. Revision is the substantive reworking of the draft at multiple levels: the argument, the structure, the evidence, the sentence construction, and finally the mechanics. Each level is best addressed separately and in order, from the largest to the smallest. There is no point in polishing individual sentences in a section that will be deleted when the argument is revised.

Level	What you are checking	Practical method
-------	-----------------------	------------------

Argument revision	Does the paper answer its research question? Is every section necessary for that answer? Is every claim supported by evidence in the paper? Are there claims without evidence, or evidence without claims?	Write a one-sentence summary of each section's contribution to the argument. If you cannot write one, the section may need to be cut or restructured. Read the introduction and conclusion only: do they match? Does the conclusion answer the question the introduction posed?
Structure revision	Is the paper in the right order? Does each section connect logically to the next? Are the transitions between sections explicit? Does each paragraph have a clear topic sentence that can be read as a summary of its content?	Read only the first sentence of each paragraph in order. This sequence should read as a coherent summary of the paper's argument. Where it does not, the paragraph structure or order needs revision.
Evidence revision	Is every claim supported by the right evidence? Are the tables and figures correctly described in the prose? Are measurement decisions documented? Are causal claims justified by the identification strategy?	Read the results and discussion sections against the tables and figures. Verify that every number cited in the text matches the table. Check that every causal claim has an identified source of variation behind it.
Sentence revision	Are sentences as short as they can be without losing meaning? Are there nominalised verb phrases that could be verbs? Are there passive constructions that obscure agency? Is the hedging calibrated to the evidence?	Read aloud. Sentences that are hard to say without pausing are usually too long or too complex. Mark every sentence over 35 words for rewriting. Find and cut throat-clearing phrases: 'It should be noted that', 'In order to', 'Due to the fact that'.
Language and mechanics	Are there spelling, grammar, or punctuation errors? Are British English conventions applied consistently (analyse, behaviour, programme, per cent, -ise/-isation endings)? Are all citations formatted correctly?	Use spell-check as a starting point but not a substitute for reading. Run the reference list check from Chapter Six. For non-native writers: identify your three most common error patterns from prior feedback and search explicitly for each one.

The five-level revision process described above is not a single pass through the document. Each level may require multiple passes, and revision at one level often generates new work at another: restructuring the argument may require new evidence to be added, which will require new sentence-level writing. This is normal and expected. Academic papers are typically revised substantially between first draft and submission, and substantially again between initial submission and publication. The writer who expects the first draft to be close to the final paper will be disappointed and may make the mistake of submitting a draft rather than a revised paper.

◆ **The one-sentence test**

A useful diagnostic for argument revision is the one-sentence test: write one sentence summarising what each section contributes to the paper's central argument. If you cannot write the sentence, the section's contribution is unclear, either because the section does not have one or because the section is doing something different from what the paper's argument requires. If the sentence you write for one section repeats the sentence you wrote for another, the sections may be doing the same work and one of them may be redundant.

The same test applied to paragraphs within a section identifies the paragraphs that are contributing to the argument and the ones that are describing without contributing — the ones whose removal would not affect the reader's understanding of the section's claim. These paragraphs are candidates for revision or deletion.

◆ **Cutting and the 10 per cent rule**

Most first drafts are too long. The material that belongs in the paper is usually present in the first draft, embedded in material that does not. A disciplined approach to length is to set a target reduction — typically 10 per cent of the word count — and to treat the cutting process as an argument-strengthening exercise rather than a loss. Every sentence that is cut because it does not contribute to the argument is a sentence that was diluting it. The paper that remains after 10 per cent has been cut is usually more persuasive than the original.

The material most susceptible to cutting is: long introductory paragraphs that do not advance to the section's substance; passages that restate what has already been established earlier in the paper; lengthy descriptions of standard procedures that can be handled in a sentence with a citation; and the throat-clearing phrases that precede each section's actual content ('In this section, we will examine...', 'Having established the above, we now turn to...').

7.4 Getting and Using Feedback

Feedback is the mechanism by which the paper's argument is tested against readers who do not share the writer's assumptions, blind spots, and familiarity with the material. No amount of self-revision fully substitutes for it: the writer cannot read their own paper as a reader, because they know what they meant to say and will tend to read that rather than what is on the

page. Feedback from other readers exposes the gap between what the writer intended and what the text actually communicates.

The quality of feedback depends substantially on when it is sought, from whom, and with what specific questions. The table below covers the main types of feedback available to economics researchers at different stages of the writing process.

Feedback type	What it addresses	How to seek and use it
Specialist co-author or colleague	Argument quality, methodology, interpretation of results, engagement with the literature. The most intellectually valuable feedback because the reader knows the field.	Share the full draft, not a summary. Ask specific questions: 'Is the identification strategy convincing?', 'Does the contribution claim hold?', 'Which parts of the literature review are missing?' General 'what do you think?' requests produce general responses.
Research seminar presentation	Which parts of the argument are unclear, which findings surprise the audience, where the contribution is not yet communicated, where the methodology raises questions. Seminar feedback reveals problems that written feedback tends to miss.	Present before the paper is fully polished — not as a way of submitting unfinished work, but because seminar audiences respond to argument gaps and unclear claims in ways that proofreaders do not. Take notes on where the audience lost the thread.
Writing-focused reader (non-specialist)	Readability, clarity of exposition, structure, sentence-level difficulties. A reader without deep knowledge of the subfield will tell you where the argument is unclear in a way that specialist readers, who fill in gaps automatically, will not.	Ask the reader to mark every sentence they had to read twice, and every passage where they lost track of the argument. Do not ask them to evaluate the economics — that is not their role. Their task is to identify where the writing fails to carry the reader.

Language or writing adviser	Grammar, syntax, academic register, British English conventions, sentence structure. Particularly useful for writers working in English as an additional language.	Prepare a list of your known recurring problems before the session (from prior feedback on earlier drafts or papers). A good language adviser will work more efficiently on specific patterns than on a general reading.
Journal referee reports	Everything: argument, evidence, contribution, methodology, exposition, citation, and compliance with the journal's conventions. The most formal and consequential feedback the paper will receive.	Read referee reports at least twice before responding: once to understand the criticism and once to plan the response. See Chapter Two (section 2.11) for the discipline of responding to referee reports. Referee comments that seem wrong or unfair should be addressed, not ignored.

◆ Responding to feedback without losing the argument

A draft that has been read by three people will typically return with three different and partly conflicting sets of suggestions. This is normal. Readers bring different priorities, different knowledge of the field, and different ideas about what the paper should be. The writer's task is not to implement all suggestions but to evaluate each one against the paper's central argument and to make a considered decision about which changes strengthen the paper and which would compromise it.

The most common mistake in responding to feedback is to treat all suggestions as equally valid and to implement them uncritically, producing a revision that satisfies no reader well because it has tried to satisfy all readers equally. A sharper approach is to ask, for each piece of feedback: does this criticism identify a genuine weakness in the argument, or does it reflect a reader preference that would pull the paper away from its central claim? The first kind of feedback deserves a substantive response. The second deserves consideration and, often, a polite decision not to act on it.

When feedback is contradictory — one reader asking for more technical detail, another asking for less — the writer must decide which audience the paper is primarily written for and resolve the conflict in favour of that audience. This decision should be explicit, not avoided.

7.5 Writing in English as an Additional Language

Most of the researchers and postgraduate students this guide is written for are working in English as an additional language, with Ukrainian language as their first. This section addresses the specific challenges that arise from this situation — not to suggest that EAL writers are at a disadvantage, but to identify the patterns that are most likely to affect the clarity of the writing and to offer specific revision strategies for each.

The fundamental point is this: writing in an additional language is a skill that develops through practice, feedback, and systematic attention to recurring patterns. The errors that EAL writers make are not random; they follow predictable patterns related to the structural differences between their first language and English. A writer who has identified their three or four most common patterns — through feedback from supervisors, language advisers, or referee reports — and who checks for those patterns explicitly at the revision stage will improve their English academic writing faster than one who tries to correct everything at once or who relies entirely on spell-checkers and grammar tools.

◆ Six recurring patterns for Ukrainian first-language writers

Pattern to watch for	Typical form and why it occurs	Revised form
Absent or misplaced article	'The study analyses behaviour of consumers in market.' Articles (the, a, an) are not used in the same way across all languages; omission or misplacement is one of the most common EAL patterns in academic writing.	'The study analyses the behaviour of consumers in the market.' Add 'the' before specific, previously mentioned, or unique nouns. Add 'a/an' for first mention of a countable singular noun in a general sense.
Nominalisation without a verb of substance	'The realisation of the increase of the efficiency of production.' A chain of abstract nouns without a main verb of substance leaves the sentence grammatically complete but semantically weak.	'Production became more efficient.' Or: 'Efficiency in production increased.' Identify the main action and make it the verb. Cut the nominalised chain.
Calque from first language	'The authors make the conclusion that...' (literal translation from Ukrainian: 'роблять висновок'). Translated idioms produce phrasing that is grammatically possible but not idiomatic in English academic writing.	'The authors conclude that...' or 'The authors argue that...'. Use the verb directly. Common calques to watch for: 'make the analysis' (analyse), 'carry out the research' (conduct the research / investigate), 'give the definition' (define).

Overlong relative clause	'The paper, which was published in 2014 and which used administrative tax data from the United States Internal Revenue Service covering children born between 1971 and 1993, finds that...' Multiple stacked relative clauses delay the main verb and make the sentence hard to follow.	Break into two sentences: 'The paper uses administrative tax data from the US Internal Revenue Service covering children born between 1971 and 1993. It finds that...'
False friend in register	'The actual situation in the Ukrainian economy...' In Ukrainian, 'actual' (актуальний) means 'relevant' or 'pressing'. In English it means 'existing in reality, as opposed to hypothetical'.	'The current situation in the Ukrainian economy...' or 'This pressing issue in the Ukrainian economy...'. Common false friends for Ukrainian speakers: actual (current/pressing), sympathetic (understanding), control (check/monitor), correspond (match/align).
Stiff or formal connectors used too heavily	'Furthermore, it should be noted that... Moreover, it is worth mentioning that... Additionally, it can be observed that...' These connectors are grammatically correct but their accumulation produces writing that feels mechanical.	Use connectors selectively and vary them. Often the clearest option is no connector at all — the logical relationship between sentences should be visible from the content. Reserve 'furthermore' and 'moreover' for cases where the addition is genuinely significant.

◆ Building academic vocabulary in English

Vocabulary in academic English is not a separate skill from academic writing; it is embedded in the sentence-level and genre conventions that the rest of this guide describes. The most effective way to build academic vocabulary is through reading: reading the exemplar papers in this guide (Friedman, 1968; Goldfarb and Tucker, 2019; Farronato and Fradkin, 2022; Chetty et al., 2014) with attention not just to the arguments they make but to the specific language they use to make them. The vocabulary and sentence models sections at the end of each chapter in this guide are designed to make this implicit learning explicit.

Three practical strategies reinforce reading. The first is to keep a personal vocabulary record: when you encounter a phrase in a paper that performs a function you need (introducing a finding, hedging a claim, acknowledging a limitation), record it with the context in which it appeared and use it in your own writing. The second is to write regularly in English even outside formal academic projects — summaries of papers, responses to seminars, notes on your own research — so that the language becomes familiar in lower-stakes contexts before it is tested in a submission. The third is to read your own writing aloud: the ear catches register errors and awkward constructions that the eye misses, particularly for writers whose English reading is stronger than their spoken English.

◆ Using translation and AI tools responsibly

Machine translation and AI writing tools have become widely available and are used by EAL writers at various stages of the drafting process. Used well, they can help: they can provide vocabulary alternatives, suggest more idiomatic phrasing, and identify constructions that are grammatically possible but register-inappropriate. Used badly, they can produce writing that is polished at the sentence level but incoherent at the argument level — because translation tools optimise for local fluency, not for the overall coherence of an argument or for the specific conventions of a genre.

The most reliable approach is to draft in English, even imperfectly, and to use translation or AI tools in the revision stage for language improvement rather than as a substitute for drafting. A paper drafted in Ukrainian and then machine-translated into English will typically have better individual sentence fluency than one drafted imperfectly in English — but it will usually have worse argument structure, because the translation process does not carry the discourse-level conventions of English academic writing. The writer who drafts in English, even roughly, develops the argument in the language in which it will be read.

All AI-assisted language editing should be disclosed as described in Chapter Six (section 6.5). Disclosure is not a disadvantage; it is part of the paper's scholarly transparency.

7.6 Vocabulary and Sentence Models

The vocabulary below covers terms specific to the academic writing process. The sentence models provide ready-to-use frames for the moves most commonly needed when describing, framing, or reflecting on the writing and revision process — particularly in cover letters, response-to-reviewers documents, and methodological notes.

◆ Key vocabulary

Term	Meaning in writing process contexts	Example of the word in use

revise	To rework a draft substantially — restructuring arguments, strengthening evidence, improving clarity. Distinct from 'proofread' (correcting mechanical errors) and 'edit' (improving sentences without changing the argument). 'Revision' is usually the larger activity.	<i>'The paper has been substantially revised in response to the referees' comments. Section 3 has been reorganised to address the concern about the identification strategy, and the literature review in Section 2 has been extended to include the recent contributions on market frictions.'</i>
iterative	Involving repeated cycles of a process, each building on the output of the previous one. Used to describe both the writing process (draft → revise → redraft) and the research process (hypothesis → test → revised hypothesis).	<i>'The argument was developed iteratively: an early presentation at the departmental seminar identified a gap in the theoretical framework, which was addressed before the data analysis was conducted, and the resulting findings required a further revision of the framing.'</i>
register	The level of formality and the set of conventions appropriate to a particular context or genre. Academic writing has a formal register; policy briefs have a different register; newspaper writing has a different register again. Register errors occur when vocabulary, tone, or sentence structure from one context appears in another.	<i>'The phrase 'it is worth noting' is acceptable in an academic register but signals to the reader that the point is secondary. If the point is central to the argument, state it directly: not 'it is worth noting that mobility varies' but 'mobility varies substantially across commuting zones'.'</i>
throat-clearing	Phrases and sentences at the beginning of sections or paragraphs that announce what will be done rather than doing it. 'In this section, we examine...', 'Having discussed X, we now turn to Y...' These phrases add words without adding content.	<i>'The draft opened every section with a throat-clearing sentence announcing the section's topic. In revision, all of these were cut: a section heading already tells the reader what the section covers; a paragraph's first sentence should make its argument, not describe itself'</i>

nominalisation	The conversion of a verb or adjective into a noun form: 'analyse' becomes 'the analysis of', 'improve' becomes 'the improvement of'. Nominalisation is a natural feature of academic writing but in excess it buries the action in a noun chain and removes the grammatical subject who performed it.	<i>'The sentence 'there was a realisation of the improvement of the productivity of workers' uses three nominalisations where one verb would serve: 'workers became more productive'.'</i>
calque	A word-for-word translation of an expression from one language into another, producing a phrase that is grammatically possible but not idiomatic in the target language. From French 'calque' (tracing). Common in writing by speakers of Ukrainian, and other Slavic languages writing in English.	<i>'The phrase 'make the analysis' is a calque from Ukrainian (зробити аналіз): in English the idiomatic form is 'analyse' or 'conduct the analysis'. Calques are best identified by asking a native speaker whether a phrase sounds natural.'</i>
false friend	A word in a second language that resembles a word in the first language but has a different meaning. False friends cause errors because the writer applies the first-language meaning to the second-language word.	<i>'The word 'actual' is a false friend for Ukrainian speakers: in Ukrainian, 'актуальний' means 'relevant' or 'pressing', while in English 'actual' means 'existing in reality, as opposed to hypothetical or theoretical'.'</i>
hedge	A linguistic device that limits the strength of a claim, signalling uncertainty or conditionality. Hedges include modal verbs ('may', 'might', 'could'), adverbs ('typically', 'generally', 'in most cases'), and qualifying clauses ('subject to the caveats described above'). Appropriate hedging is a marker of academic precision, not weakness.	<i>'The paper's causal claims are hedged appropriately: where the identification strategy supports a causal interpretation, the paper uses causal language; where the analysis is correlational, it uses 'is associated with' and 'tends to'. The hedges track the evidence, not the writer's confidence.'</i>

◆ Sentence models

The models below cover the moves most commonly needed when describing the writing and revision process — in cover letters, response-to-reviewers documents, acknowledgement sections, and methodological notes. They are equally useful for academic writing about writing.

Function	Model sentence	Notes for the writer
Framing the paper's contribution (cover letter)	<i>This paper [investigates / examines / estimates] [specific question], using [data/method]. We find that [central finding], which [implies / suggests / shows] [why it matters].</i>	The cover letter introduction should answer in three sentences what the abstract answers in 150 words. Keep it specific: editors make desk-rejection decisions quickly.
Stating a revision (response to reviewers)	<i>In response to Referee [N]'s comment on [specific issue], we have [revised / extended / clarified / added] [specific change]. The revised text appears on page [N], paragraph [N].</i>	Every substantive revision should be traceable to a specific location in the manuscript. Vague responses ('we have improved the discussion') without a page reference give referees no way to evaluate whether the concern has been addressed.
Acknowledging a limitation (in the paper)	<i>A limitation of this analysis is [specific limit]. This arises because [reason]. It implies that [what the reader should be cautious about], though it does not affect [what the paper's main finding does establish].</i>	State the limit, explain it, and bound it. A limitation statement that does not say what the limit does not affect is an invitation to reject the paper's entire contribution.

Signposting structure	<i>Section [N] [develops the theoretical framework / presents the data and estimation strategy / reports the main results / discusses the findings in relation to the literature]. Section [N+1] [concludes / extends the analysis to / examines the robustness of the results].</i>	Roadmap sentences work best when each section description states the section's function in the argument, not just its topic. 'Section 3 establishes that identification holds' is more informative than 'Section 3 presents the econometric strategy'.
Qualifying a claim	<i>[Finding], though [this pattern / this relationship / this effect] is [stronger / more pronounced / clearest] among [specific group or context] and should be interpreted with caution when applied to [what it does not cover].</i>	Qualifications placed after the finding, not before it. A qualification that precedes the finding buries the finding; a qualification that follows it bounds it appropriately.
Crediting a co-author's contribution	<i>The [theoretical framework / identification strategy / dataset construction] in this paper was developed by [name], who also [contributed specific element]. The [empirical analysis / literature review / policy interpretation] was led by [name].</i>	Author contribution statements are increasingly required by journals. They should be specific enough to be meaningful — 'all authors contributed equally' is increasingly regarded as an insufficient statement.
Noting a translation or language assistance	<i>This paper was written in English by authors whose first language is Ukrainian. Language editing assistance was provided by [name or tool]. All intellectual content and conclusions are the authors' own.</i>	A transparent note in the acknowledgements about language assistance is a professional courtesy to referees and editors, not an admission of weakness. It sets appropriate expectations and is standard practice for many EAL research groups.

Connecting a section to the paper's argument	<i>Together, these findings establish that [claim derived from results], which addresses the paper's central question: [research question restated].</i>	Used in the discussion and conclusion to make the connection between the section's local findings and the paper's global argument explicit. Many papers leave this connection implicit; making it explicit is one of the most effective single revisions a writer can make.
---	--	---

Writing is not what happens after research is finished. It is how research is made precise, made communicable, and made accountable to the readers who will evaluate it. The process described in this chapter — plan, draft, revise at multiple levels, seek feedback, revise again — is not a formality that surrounds the intellectual work. It is part of the intellectual work. The gap between a finding that a researcher understands and a paper that communicates that finding clearly to a sceptical reader is closed by this process, and by no other.

→ **Conclusion**

CHAPTER EIGHT

Writing the Abstract

The abstract is the most widely read part of any academic paper. Before a reader decides whether to download an article, request it through a library, or spend an hour with it, they read the abstract. It is the paper's public face: the compressed argument that stands in for the full text in search results, citation databases, conference programmes, and grant applications. A paper with a weak abstract will be underread regardless of the quality of the analysis it contains. A paper with a strong abstract will attract the right readers and create the right expectations before they reach page one.

Despite this, abstracts are consistently among the most poorly written sections of economics papers. They are frequently drafted last, treated as a formality, and left to summarise rather than argue. This chapter explains what an abstract must do, how to do it in the word count most journals allow, how to adapt the abstract to different journal conventions and different genres of economics writing, and how to revise a weak abstract into a strong one.

8.1 What the Abstract Must Do

The abstract performs four functions simultaneously, and a well-written abstract performs all four within the word count. Understanding these functions — and recognising when one of them is absent — is the foundation of abstract revision.

The first function is to state the research question or the paper's central purpose. The reader must know within the first one or two sentences what the paper is about in the most precise sense: not the general topic area, not the data used, but the question being investigated. An abstract that begins 'This paper studies inflation expectations in emerging markets' has stated a topic. An abstract that begins 'This paper examines whether anchored inflation expectations reduce the output cost of disinflation in emerging markets' has stated a question. The difference is not stylistic; it is the difference between orienting the reader and informing them.

The second function is to describe the method or approach. The reader needs to know how the question was addressed: what data were used, what identification strategy was employed, what analytical framework was applied, and what the key methodological decision was that makes the paper's findings credible. This does not require methodological detail — the abstract is not a methodology section — but it requires enough specificity for the reader to evaluate whether the approach is appropriate and whether the paper is likely to be relevant to their own work.

The third function is to state the main findings. Economics abstracts should report quantitative results where the paper is empirical. A finding described only as 'significant' or 'positive' is not a finding as far as the abstract is concerned; it is a promissory note. 'Inflation expectations reduce the output cost of disinflation by approximately 30 per cent in countries

with independent central banks' is a finding. The specificity is what gives the abstract its value as a standalone document.

The fourth function is to state the implication or contribution. What does this paper add to what was previously known? Why does the finding matter? This is the move that converts a summary into an argument — that positions the paper within the field and tells the reader not just what was found but why it is worth knowing. For theoretical papers, the contribution replaces the empirical finding: the novel mechanism, the generalisation, the reconciliation of previously conflicting results.

The four-move abstract: a diagnostic checklist

Move 1 — Question: Does the abstract state the research question or central purpose precisely (not just the topic)?

Move 2 — Method: Does the abstract specify the data, identification strategy, or analytical approach?

Move 3 — Finding: Does the abstract report the main result with quantitative specificity (for empirical papers)?

Move 4 — Contribution: Does the abstract explain what the paper adds to existing knowledge and why it matters?

8.2 The Model Abstract: Friedman (1968) Annotated

Friedman's 1968 presidential address does not have a formal abstract — AER convention at the time did not require one — but the opening paragraph of the paper functions as an abstract in everything but name. It is worth examining because it executes all four moves with unusual economy and precision, and because it demonstrates what abstract-level argument looks like in practice.

What monetary policy can and cannot do. In recent years, economic theory has returned to the analysis of monetary policy. There is widespread agreement about certain descriptive facts. There is also, I believe, agreement about some of the basic principles of monetary policy. Yet to judge by stated objectives, there is sharp disagreement about what monetary policy can do. Finally, there is equally sharp disagreement about how monetary policy should be conducted. Let me then summarize what I see as the state of this debate, and use it to frame the paper's central claim: that there is a natural rate of unemployment that monetary policy cannot alter in the long run, and that attempts to do so systematically will generate accelerating inflation.

Friedman (1968, p. 1) — paraphrased for illustration

A contemporary abstract of Friedman (1968) — written to current AER conventions — would look something like the following. The annotated version below shows how the four moves are distributed across the abstract's sentences.

Sentence from the abstract	Move performed
<i>This paper analyses the long-run neutrality of monetary policy and the limits of the Phillips curve as a guide to policy.</i>	Move 1 — Question / central claim
<i>Using the distinction between the nominal and real interest rate, and between the actual and natural rate of unemployment, I develop a theoretical framework that distinguishes the short-run and long-run effects of monetary expansion.</i>	Move 2 — Method / analytical approach
<i>The analysis demonstrates that the long-run Phillips curve is vertical: monetary policy can lower unemployment temporarily, but only by generating accelerating inflation.</i>	Move 3 — Main finding
<i>These results imply that central banks should target monetary aggregates rather than unemployment, and that attempts to exploit any short-run trade-off will progressively worsen the inflation-unemployment trade-off over time.</i>	Move 4 — Implication / contribution

Several features of this abstract are worth noting. The question is stated in the first sentence as a claim about what the paper analyses, not as an open question — a convention common in theoretical economics where the answer is known before the reader begins. The method sentence specifies the conceptual tools used (the distinction between nominal and real rates, actual and natural rates) without describing their mathematical form. The finding sentence contains a specific, testable claim (the vertical long-run Phillips curve) rather than a general description of results. The implication sentence moves directly from the finding to its policy relevance.

8.3 Structure and Word Count

Most economics journals specify abstract word counts in their author guidelines. The most common limits are 100 words, 150 words, and 200 words, with some journals allowing up to 250. A small number of journals — including some in applied economics and economic history — allow structured abstracts with labelled sections. Understanding the implications of each constraint is necessary for effective abstract writing.

◆ **The 100-word abstract**

A 100-word abstract forces maximum compression. Each of the four moves must be accomplished in one sentence or less. This constraint rules out any background or contextualisation: the abstract must begin with the question, not with the problem that motivates the question. Method is reduced to a single identifying feature (the data source, the key instrument, the analytical framework). The finding is a single quantified result. The contribution is a single sentence. Writers working under a 100-word limit should draft to 150 words and then cut, rather than attempting to write 100 words directly.

◆ **The 150-word abstract**

A 150-word abstract allows one background sentence before the question, one or two sentences on method, one or two sentences on findings (covering the main result and one key robustness or heterogeneity result), and one sentence on contribution. This is the most common format at leading economics journals and the format for which the guidance in this chapter is primarily written.

◆ **The 200-word abstract**

A 200-word abstract gives more space to explain the study. The writer can describe the purpose, explain the method in more detail, and present several results. This type of abstract is often used in applied journals and working papers, where there are fewer limits on length.

However, there is a risk. Writers may use this extra space to add general description instead of clear ideas. They may include too much background or unnecessary comments.

Even in a longer abstract, it is important to follow a clear structure. The extra space should be used to give more precise information, not to add unnecessary text.

◆ **Structured abstracts**

Some journals — particularly in health economics and economic policy — require structured abstracts with labelled sections such as Background, Objective, Methods, Results, and Conclusions. These are common in interdisciplinary outlets that also publish medical and public health research. If the target journal requires a structured abstract, follow the journal's specified labels exactly. The four-move framework maps directly onto most structured abstract formats: Background and Objective correspond to Move 1, Methods to Move 2, Results to Move 3, and Conclusions to Move 4.

8.4 Abstracts Across Genres

The four-move structure applies to all genres of economics writing, but the relative weight given to each move, and the content of each move, varies by genre. A literature review abstract looks different from an empirical paper abstract, and both differ from a theoretical model abstract.

◆ The empirical research paper abstract

This is the paradigm case described in sections 8.1 to 8.3. The method move is the most demanding: the abstract must specify the identification strategy clearly enough for a specialist reader to evaluate the paper's credibility. 'Using an instrumental variables approach' is not sufficient; 'using rainfall as an instrument for agricultural income' is. The finding move must be quantitative. The contribution move should position the paper against the specific literature it advances, not the general field.

◆ The theoretical paper abstract

In a theoretical paper, the method move describes the model's key structure rather than a dataset or identification strategy. The finding move states the paper's main propositions or results — what the model shows, what equilibrium exists, what comparative statics hold. The contribution move explains what the theory adds: a new mechanism, a generalisation of prior results, a reconciliation of apparently contradictory predictions. Quantitative specificity is not possible in the findings move, but logical precision is: 'we show that X holds if and only if condition Y is satisfied' is more informative than 'we show that the model has an equilibrium'.

◆ The literature review abstract

A standalone literature review abstract (for a paper of the type examined in Chapter Three) differs from an empirical abstract in that the contribution move carries most of the argumentative weight. The method move describes the scope and selection criteria of the review. The finding move synthesises the state of the literature — what has been established, where there is disagreement, what the major open questions are. The contribution is the synthesis itself: what this review adds by virtue of its scope, its organising framework, or its identification of the field's open questions.

◆ The case study abstract

A case study abstract (for papers of the type examined in Chapter Four) must handle the generalisation problem within the abstract itself. The method move identifies the case and explains why it is informative. The finding move reports the case-specific result. The contribution move states the generalisation — what can be inferred beyond the specific case, and why. An abstract that presents case study findings without addressing their external validity is incomplete, regardless of word count.

◆ **The working paper abstract**

Working paper abstracts are typically less constrained by word count and are read in a context where the reader has already decided to engage with the paper. The temptation is to use the additional space for comprehensiveness rather than argument. Resist this: the working paper abstract should be written to the same four-move standard as a journal abstract, using any additional space to add quantitative precision to the findings move rather than to expand the motivation or contextualisation.

8.5 Common Failures in Economics Abstracts

The following failures appear with sufficient regularity in submitted manuscripts to warrant explicit treatment. Each represents a different way of writing an abstract that functions as a description of the paper rather than as a compressed argument.

◆ **The topic announcement**

The topic announcement describes the paper's subject matter without stating a question or claim. 'This paper studies the relationship between financial development and economic growth in sub-Saharan Africa' is a topic announcement. It tells the reader what the paper is about but not what it finds or what question it addresses. Topic announcements are common as opening sentences even in otherwise adequate abstracts; the problem is when the abstract never moves beyond them. Revise by converting the topic into a question ('whether financial development causes economic growth') or a claim ('that financial development causes economic growth only where contract enforcement is adequate').

◆ **The methodology advertisement**

The methodology advertisement devotes excessive space to describing the analytical approach at the expense of the findings. 'This paper employs a generalised method of moments estimator with heteroskedasticity-robust standard errors, controlling for country fixed effects and time trends, using a panel dataset of 47 countries over the period 1990 to 2020' has used the entire method allocation on technical detail. The reader learns how the paper was written but not what it found. Abstract-level method description should specify the key identifying feature — the strategy that makes the result credible — not the full methodological apparatus.

◆ **The finding hedge**

The finding hedge replaces a specific result with a directional description. 'The results suggest that financial development has a positive and statistically significant effect on growth' is a finding hedge. It conveys direction and significance but not magnitude. 'A one standard deviation increase in private credit is associated with a 0.4 percentage point increase in annual GDP growth, equivalent to approximately one third of the within-country standard deviation in growth rates' is a finding. The hedge is often a symptom of abstract-writing anxiety — a reluctance to commit to a specific claim that could be challenged. It should be resisted.

◆ The absent contribution

An abstract that performs Moves 1 to 3 but omits Move 4 is incomplete. Without the contribution move, the reader cannot assess why the paper matters or how it relates to the existing literature. This is the most common structural failure in economics abstracts, and it is particularly common in papers by researchers who have not yet internalised the distinction between a finding and a contribution. Finding: 'inflation expectations reduce the output cost of disinflation'. Contribution: 'these results resolve the empirical conflict between studies using aggregate and firm-level expectations data, and suggest that standard Phillips curve models underestimate the role of credibility in disinflationary episodes'.

◆ The five-move abstract

Some abstracts attempt to include a background or context move before stating the question. 'In recent years, central bank credibility has attracted renewed attention among both policymakers and researchers. This paper examines...' The background move is almost always unnecessary at the abstract level and almost always comes at the cost of space that should be given to findings and contribution. Cut it. If the context is essential, it can occupy half a sentence: 'Motivated by the post-2008 concern with central bank credibility, this paper examines...'

Five abstract failures and their diagnosis

Topic announcement — the abstract states a subject but not a question, claim, or finding.

Methodology advertisement — the method move is technically detailed but the findings move is absent or vague.

Finding hedge — the result is described as 'positive and significant' without quantitative specificity.

Absent contribution — Moves 1–3 are present but the abstract never states what the paper adds to existing knowledge.

Five-move abstract — context or background is included at the cost of space needed for findings and contribution.

8.6 Revision: From Draft to Final Abstract

Abstracts should be written twice. The first version is written before or during the drafting of the paper and serves as a planning document: it forces the writer to articulate the question, the method, and the expected contribution before the full analysis is completed. This version is

almost always wrong in some of its specifics, but the exercise of writing it is analytically valuable. The second version is written after the paper is complete, when the findings are known and the contribution can be stated precisely.

The revision process should follow three stages. In the first stage, check structure: identify which of the four moves is present and which is absent. A move that is absent needs to be added; a move that is present but vague needs to be specified. In the second stage, check precision: every finding should be quantitative, every method description should name the key identifying feature, every contribution should be positioned against specific prior work rather than the general literature. In the third stage, check economy: every sentence that does not advance one of the four moves should be cut.

◆ A revision example

The following before-and-after pair illustrates what systematic abstract revision produces. The first version is representative of a typical first draft; the second applies the four-move framework.

First draft	Revised abstract
<i>This paper studies the effect of trade liberalisation on labour markets in developing countries. We use data from a large number of countries and employ econometric techniques to examine wages and employment. The results show that trade liberalisation has mixed effects, with some workers benefiting and others facing negative outcomes. We discuss the policy implications of these findings.</i>	<i>This paper examines whether trade liberalisation raises or lowers wages for low-skilled workers in developing countries. Using tariff reductions driven by multilateral WTO accession negotiations as an instrument for trade exposure, we estimate wage effects for 64 countries over 1995–2018. Low-skilled wages fall by 4.2 per cent on average in the five years following liberalisation, but recover fully within a decade in countries with flexible labour market institutions. These results reconcile conflicting prior estimates by showing that labour market flexibility, not liberalisation per se, determines long-run distributional outcomes.</i>
<i>Missing: specific question, named method, quantified finding, stated contribution. (53 words)</i>	<i>All four moves present. Instrument named. Result quantified with magnitude and time horizon. Contribution positions paper against prior literature. (90 words)</i>

8.7 Practical Considerations: Submission, Search, and Reuse

◆ Keywords

Most journals ask for four to eight keywords submitted alongside the abstract. Keywords should be chosen to maximise discoverability in the journal's indexing system and in general academic databases (EconLit, Google Scholar, Scopus). Effective keywords for an economics paper typically include: the main topic or phenomenon studied; the methodology or identification strategy (e.g., instrumental variables, regression discontinuity, difference-in-differences); the geographic or institutional context; and the main outcome variable. Avoid keywords that are so general as to be uninformative ('economics', 'growth') or so specific as to be unsearchable. Check whether the target journal provides a keyword taxonomy — many field journals use JEL classification codes, which should accompany rather than replace substantive keywords.

◆ JEL classification codes

Journal of Economic Literature (JEL) codes classify papers by topic and are required by most economics journals and working paper repositories. Each code begins with a letter identifying the primary classification and continues with two digits specifying the sub-field. A typical empirical paper will carry two to four codes. Choose codes at the two-digit level (e.g., J31 — Wage Level and Structure) rather than only at the letter level (J — Labor and Demographic Economics), which is too broad to be useful for indexing. The full JEL code taxonomy is available on the AEA website. Selecting appropriate codes is not merely administrative: the codes determine in which journal sections and database filters the paper appears, and thus shape who encounters it.

◆ Writing the abstract for a working paper versus a journal submission

Working paper abstracts are read in a lower-commitment context than journal abstracts — they are often the entire basis on which a reader decides whether to open the file. The abstract for a working paper should therefore be as argumentatively complete as the word count allows: all four moves should be present, the finding should be quantified, and the contribution should be stated. Some researchers write longer working paper abstracts (150 to 200 words) than they would for a journal submission, on the grounds that the additional specificity increases download rates. This is a reasonable practice, provided the abstract is shortened to meet journal requirements when the paper is submitted.

◆ Reusing the abstract across submissions

It is common for a paper to be submitted to multiple journals before acceptance. The abstract should be reviewed and if necessary revised at each submission. Different journals have different conventions about the appropriate level of technical language in an abstract, the expected length, and the balance between methodological and finding content. An abstract written for a specialist field journal may be too technical for a general interest journal. One written for a 100-word limit should not simply be resubmitted to a journal that allows 200 words — the additional space should be used to add specificity, not to leave the abstract as it was.

8.8 Vocabulary and Sentence Models

The terms below are used consistently in high-quality economics abstracts and in the critical vocabulary for discussing abstract quality. The sentence models that follow provide ready-to-use frames for each of the four abstract moves, with register notes for adapting them to different genres and contexts.

◆ Key vocabulary

Term	Usage note
contribution	What the paper adds to the existing literature — distinct from a 'finding', which is a result. 'Our contribution is to show...' vs 'We find that...' The former positions the paper; the latter reports a result.
identification strategy	The approach used to establish causal rather than merely associational claims. Abstract-level reference: 'exploiting variation in X as an instrument for Y' or 'using a regression discontinuity in Z'.
implication	The consequence of the finding — policy, theoretical, or methodological. 'These results imply that...' or 'The findings suggest that policy X may...'
mechanism	The causal pathway through which an effect operates. Particularly important in theoretical and structural papers. 'The mechanism is...' or 'We show that X operates through Y'.
motivate	To provide the reason for the research question. 'Motivated by...' signals the background claim without adding a full context sentence. Use sparingly.
novel	Signals a contribution that has not previously been made. Use only when the claim is accurate: 'the first paper to...' or 'we provide the first systematic evidence on...'

robust	Results that hold across alternative specifications. In abstracts: 'the effect is robust to...' followed by the most important alternative tested.
trade-off	A situation where improving one outcome worsens another. Common in policy-relevant abstracts. British spelling: trade-off (hyphenated).
welfare	Economic wellbeing measured in terms of consumer surplus, utility, or equivalent variation. Welfare claims require careful hedging in abstracts: 'increases consumer welfare by...' not 'improves welfare'.
heterogeneity	Variation in the effect across subgroups. 'The effect is concentrated among...' or 'We document significant heterogeneity by...' Heterogeneity results are valuable in abstracts because they clarify who is affected.

◆ Sentence models

Move	Sentence model	Register note
Question (empirical)	<i>This paper examines [whether / the extent to which] __ [causes / affects / determines] __ [in context].</i>	Standard opening for empirical abstracts. 'Whether' signals an open question; 'the extent to which' signals the answer is known to be non-zero.
Question (theoretical)	<i>This paper [develops / analyses / characterises] a model of __ in which __.</i>	For theoretical papers. Name the setting and the key feature of the model.
Question (review)	<i>This paper [surveys / synthesises] the literature on __ with a focus on __.</i>	For standalone literature reviews. Specify scope and organising focus.

Method (instrumental variables)	<i>We [exploit / use] __ as [an instrument / exogenous variation] for __, using [data source] over [period].</i>	Name the instrument explicitly. 'Exploit' signals identification design; 'use' is weaker but acceptable.
Method (difference-in-differences)	<i>We compare __ before and after __ using a difference-in-differences design, with __ as the control group.</i>	Specify treated and control groups. Time period and data source should follow.
Method (regression discontinuity)	<i>We [exploit / use] a [sharp / fuzzy] regression discontinuity in __ to estimate the effect of __ on __.</i>	Specify the threshold and the forcing variable in the full sentence if space allows.
Method (theoretical)	<i>The model features __, and we [characterise / derive] equilibrium conditions for __.</i>	Describes model structure without full mathematical detail.
Finding (quantified)	<i>A [one-unit / one standard deviation] increase in __ [raises / lowers] __ by __ [per cent / percentage points / units], equivalent to __.</i>	Provide a translation of the coefficient into interpretable terms where possible.
Finding (conditional)	<i>The effect is [concentrated among / larger for / absent in] __, where __ [increases / decreases] by __.</i>	For heterogeneity results. Specifies the subgroup and the magnitude.
Contribution	<i>These results [contribute to / extend / resolve] __ by [showing / demonstrating / reconciling] __.</i>	Position the paper explicitly against prior work. Name the literature or the unresolved question.
Contribution (first-evidence)	<i>We provide the first [systematic / causal / large-scale] evidence on __, using __ to address the identification challenges that have limited prior work.</i>	Use only when accurate. 'First' claims are scrutinised by referees.

Implication	<i>The findings [imply / suggest] that __ [policy / mechanism / assumption] __, and [may / should / have implications for] __.</i>	'Imply' is stronger than 'suggest'. Use 'imply' for logically necessary consequences; 'suggest' for plausible but uncertain ones.
--------------------	--	---

CONCLUSION

Writing and Thinking as a Single Practice

This guide began with a distinction: between argument and description. It ends there too, because that distinction is not merely one principle among several — it is the organising logic of everything that follows from it. A paper that describes without arguing, that presents findings without claiming anything about them, that engages the literature without positioning itself within it: this is not a paper that has done good economics poorly. It is a paper that has not yet become economics writing. The discipline of academic writing in economics is the discipline of converting analysis into argument — of moving, at every level from the sentence to the paper as a whole, from 'here is what I found' to 'here is what it means, why it is credible, and why it matters.'

The eight chapters of this guide have addressed that discipline from different directions. It is worth drawing the threads together — not to repeat what each chapter said, but to show how the principles connect across the genres and tasks the guide has covered.

◆ **Argument is the common thread**

The four-move sequence introduced in Chapter One — claim, evidence, analysis, implication — is not a formula specific to the research paper. It operates at every scale and in every genre. In the literature review (Chapter Three), the synthesis move that carries the argument from 'here is what the literature has established' to 'here is what it has not yet established' is the analysis move applied to other people's work. In the case study (Chapter Four), the generalisation problem — what a single case can and cannot establish — is the implication move made harder by the evidential constraints of a N=1 design. In quantitative writing (Chapter Five), the correlation–causation discipline is the analysis move applied to statistical findings: what does this number mean, and what does it justify claiming? In every chapter, the question 'what are you arguing?' is the question that determines whether the writing is working.

The corollary is that most of the failures described in this guide are argument failures dressed up as technical failures. Undocumented measurement decisions (Chapter Five) are not failures of data management; they are failures to complete the analysis move — to connect the data choices to the claims the paper makes. A literature review that catalogues rather than synthesises (Chapter Three) is not a failure of thoroughness; it is a failure to advance an argument through the literature. A results section that describes tables without interpreting them (Chapter Two) is not a formatting failure; it is a failure to make the claim that the analysis supports. Technical and methodological competence is a precondition for economics writing, but it is not the writing itself.

◆ **Precision serves argument, not performance**

A recurring theme across the chapters is precision — in measurement documentation, in causal language, in hedging, in citation, in the statement of limitations. Precision in academic economics writing is sometimes misunderstood as a form of self-protection: qualifying every claim extensively, documenting every decision in exhaustive detail, hedging every finding until nothing falsifiable remains. This misunderstanding produces writing that is technically accurate but intellectually inert — writing that cannot be wrong because it has refused to claim anything.

The kind of precision this guide describes is different. It is the precision that serves the argument by matching the strength of the claim to the strength of the evidence; that documents measurement decisions because the reader needs that information to evaluate the paper's credibility; that hedges causal language because the identification strategy, not the writer's confidence, is what determines the appropriate claim. Friedman's 1968 paper — the guide's first model text — is precise in this way. It is also one of the most assertive papers in the macroeconomics literature. The two properties are not in tension. Precision of this kind is what makes bold claims credible.

◆ **Genre shapes argument but does not determine it**

Chapters Two through Five addressed specific genres — the research paper, the literature review, the case study, the quantitative report — and in each case described the conventions that govern structure, evidence, and prose. These conventions are real and their mastery matters: a literature review structured as a research paper will fail on its own terms; a case study that does not engage the generalisation problem will be rejected on those grounds regardless of its analytical quality.

But genre conventions are not a substitute for argument. They are the framework within which argument is made. The IMRaD structure of a research paper is not a template that produces a good paper when filled in; it is a set of expectations that a well-made argument satisfies for good reasons. The reason the introduction comes first is that the reader needs to know what is being claimed before they can evaluate the evidence for it. The reason the results section is separated from the discussion is that the reader needs to see the findings before they encounter the interpretation. Understanding why the conventions are what they are — rather than just what they are — is what makes it possible to depart from them intelligently when the argument requires it.

◆ **Clear writing and rigorous thinking are the same achievement**

The introduction to this guide proposed that clear writing and rigorous thinking are not separate achievements — that they are the same achievement, approached from different directions. Eight chapters of detailed guidance later, that proposition is easier to see concretely. The writer who cannot state the paper's research question in a single clear sentence does not have a writing problem; they have a research question problem. The writer who cannot write the paper's contribution without citing five other papers in the process does not have a prose problem; they have a contribution problem. The writer whose results section describes what is

in the tables but never says what the tables show does not have a communication problem; they have not yet finished the analysis.

This is why writing cannot be separated from research and treated as a finishing stage — a polish applied to completed work. The act of writing the measurement justification reveals whether the justification is actually there. The act of writing the contribution statement reveals whether the contribution is actually distinct. The act of writing the discussion reveals whether the findings actually support the interpretation. Writing does not merely communicate the results of rigorous thinking; it is one of the primary instruments by which rigorous thinking is carried out. The discipline of this guide is the discipline of taking that seriously.

Clear writing and rigorous thinking are not separate achievements. They are the same achievement, approached from different directions.

APPENDIX A

APA 7th Edition Quick-Reference Card

◆ A1 In-text citations

Use the author-date format. Two forms are acceptable: parenthetical — (Author, Year) — and narrative — Author (Year). Choose one style per paper and apply it consistently.

Situation	Parenthetical form	Narrative form
1 author	<i>(Friedman, 1968)</i>	<i>Friedman (1968)</i>
2 authors	<i>(Farronato & Fradkin, 2022)</i>	<i>Farronato and Fradkin (2022)</i>
3+ authors	<i>(Chetty et al., 2014)</i>	<i>Chetty et al. (2014)</i>
Direct quote	<i>(Friedman, 1968, p. 9)</i>	<i>Friedman (1968, p. 9)</i>
Same author/year	<i>(Chetty et al., 2014a)</i>	<i>Chetty et al. (2014a)</i>
Multiple sources	<i>(Chetty et al., 2014; Goldfarb & Tucker, 2019)</i>	<i>— order alphabetically; separate with semicolons</i>

Key rules for in-text citations

Use '&' between authors inside parentheses; spell out 'and' in the running text.

For direct quotations: always include a page number — (Author, Year, p. N) or (Author, Year, pp. N–N).

For same author, same year: add a, b, c to the year in both the in-text citation and the reference list entry.

For a work you have not read directly: '(Original Author, Year, as cited in Author you read, Year)' — then list only the work you read in the reference list.

◆ A2 Reference list formats

The reference list is ordered alphabetically by first author surname. Use 'and' (not '&') between the last two authors in reference list entries. Include a DOI as a URL (<https://doi.org/...>) whenever one is available.

Source type	APA 7th format	Example
Journal article	<i>Author, A. A., and Author, B. B. (Year). Title in sentence case. Journal Name in Title Case, volume(issue), pp–pp. https://doi.org/xxxxx</i>	Friedman, M. (1968). The role of monetary policy. <i>American Economic Review</i> , 58(1), 1–17.
Book	<i>Author, A. A. (Year). Title of book in sentence case. Publisher.</i>	Neugeboren, R., and Jacobson, M. (2005). <i>Writing economics</i> . Harvard University Press.
Working paper	<i>Author, A. A. (Year). Title (Working Paper No. N). Institution. https://doi.org/xxxxx</i>	Chetty, R., et al. (2013). The economic impacts of tax expenditures (Working Paper No. 19843). NBER. https://doi.org/10.3386/w19843
Dataset	<i>Author, A. A. (Year). Title of dataset [Data set]. Repository/Publisher. https://doi.org/xxxxx</i>	Include [Data set] label in square brackets after title. Provide DOI or stable URL.
Government / institutional report	<i>Institution. (Year). Title of report. Publisher. https://doi.org/xxxxx</i>	World Bank. (2022). <i>World development report 2022</i> . World Bank Group. https://doi.org/10.1596/978-1-4648-1730-4

Formatting rule	Detail
Article / chapter titles	Sentence case — capitalise only the first word, proper nouns, and the first word after a colon.
Journal / book titles	Title case and italics — capitalise all major words. Italicise the full title.

Volume numbers	Italicised. Issue numbers are not italicised and appear in parentheses.
DOI format	Always as a URL: https://doi.org/10.xxxx/xxxxx — not as a label 'DOI: ...'
Author names	Surname first, then initials. 'Smith, J. A.' not 'John A. Smith'.
Number of authors	List all authors in the reference list regardless of how many. Use 'et al.' only in in-text citations for three or more.
Pages	Use en dash: pp. 1553–1623, not pp. 1553-1623.

APPENDIX B

Pre-Submission Checklist

This checklist consolidates the pre-submission checks from all seven chapters. Work through it in order — from argument to mechanics — before submitting to a journal or circulating as a working paper. Each section references the chapter where the relevant guidance appears.

Argument and contribution (Chapters One and Two)

- ☐ The research question is stated explicitly and answered directly in the conclusion.
- ☐ The paper's contribution is distinct from the prior literature — not just related to it.
- ☐ Every section advances the central argument; sections that only describe without contributing have been cut or revised.
- ☐ The four-move structure (claim → evidence → analysis → implication) is complete in each major section.
- ☐ The introduction's four functions are fulfilled: context, gap, contribution, roadmap.
- ☐ The conclusion answers the research question in the first paragraph and does not introduce new evidence.

Literature review (Chapter Three)

- ☐ The review synthesises the literature — it does not catalogue it.
- ☐ The gap identified in the review is specific, not generic ('this topic has not been studied before').
- ☐ The paper's contribution is positioned against the gap: the reader can see exactly what is new.
- ☐ Sources are attributed accurately and in the appropriate register (strong claim vs. weak claim vs. contested finding).
- ☐ The review does not cite papers not yet read in full; secondary citations are flagged as such.

Data, measurement, and results (Chapters Four and Five)

- Each significant measurement decision is stated, justified, and its cost acknowledged.
- Sample construction is described as an explicit sequence with observation counts at each stage.
- Every table and figure is self-contained: headers defined, data source stated, significance conventions explained.
- Causal language ('causes', 'leads to', 'produces') is used only where the identification strategy supports it.
- Every key finding includes a magnitude description — not just a coefficient, but a translated, interpretable quantity.
- Robustness checks state the concern addressed, the specification tested, and the numerical result.
- Limitations are stated proportionately: specific, bounded, and not used to undermine the paper's main finding.

Citation and integrity (Chapter Six)

- Every direct quotation is in quotation marks and has a page number.
- Every specific empirical finding from another paper is cited to its source.
- Every in-text citation has a matching reference list entry, and vice versa.
- Author names, years, titles, volume numbers, and DOIs are verified against the primary source.
- Working papers that have since been published are cited in their published journal form.
- Overlap with the author's own prior published work is disclosed in the submission letter.
- Data and code availability requirements of the target journal have been checked and met.
- Any use of AI writing tools is disclosed in the acknowledgements.

Writing process and language (Chapter Seven)

- The paper has been revised at argument level before being revised at sentence level.
- The first sentence of each paragraph states the paragraph's argumentative contribution.
- Throat-clearing phrases have been cut from the openings of sections and paragraphs.
- Sentences over 35 words have been reviewed and rewritten where possible.
- Nominalisations have been checked: does each one serve a purpose, or can it be replaced by a verb?
- For EAL writers: the three most common recurring error patterns from prior feedback have been searched for explicitly.
- British English conventions are applied consistently throughout: -ise/-isation, analyse, behaviour, programme, per cent.
- The paper has been read aloud at least once during the final revision stage.

Before submission: journal-specific requirements

- The paper's length, reference style, and abstract word count conform to the target journal's author guidelines.
- Figures and tables are provided in the format (resolution, file type, placement) specified by the journal.
- The covering letter states the paper's contribution, confirms it is not under review elsewhere, and discloses any relevant overlap with prior work.
- Any required ethics statement, data availability statement, or conflict of interest declaration is included.
- Blind review requirements are met: no author names, institutional affiliations, or self-identifying references in the manuscript.

REFERENCES

All Works Cited in This Guide

Citations follow the American Psychological Association (APA) 7th edition style throughout. Entries are grouped by the chapter in which they are first cited, then listed alphabetically by first author within each group. Where a work is cited in more than one chapter, it appears under the chapter of first citation only.

APA 7th edition: key formatting rules applied in this guide

Author names: Last name, followed by initials only (e.g., Friedman, M.). Up to 20 authors are listed; for 21 or more, list the first 19, then an ellipsis (...), then the final author.

Publication year: In parentheses immediately after the author(s), followed by a full stop.

Article titles: Sentence case only — capitalise the first word, the first word after a colon, and proper nouns. No italics.

Journal names: Title case and italicised. Volume number italicised; issue number in parentheses, not italicised.

Page ranges: Use an en-dash (–), not a hyphen. For in-text citations, include page numbers for direct quotations.

DOI: Formatted as a hyperlink beginning <https://doi.org/> where available.

Conference proceedings (published volume): Treated as a book chapter: Author(s). (Year). Title of contribution. In Editor(s) (Eds.), Title of volume (pp. xx–xx). Publisher.

Hanging indent: All reference entries use a 1.27 cm (0.5 in) hanging indent — first line flush left, subsequent lines indented.

Chapter One — The Architecture of Economic Argument

Friedman, M. (1968). The role of monetary policy. *American Economic Review*, 58(1), 1–17.

Neugeboren, R. & Jacobson, M. (2005). *Writing economics*. Harvard University Department of Economics.

<https://www.economics.harvard.edu/files/economics/files/WritingEconomics.pdf>

Chapter Two — The Research Paper for Publication

Note: Chapter Two draws on Friedman (1968) and Neugeboren & Jacobson (2005), both cited under Chapter One above. No additional sources are introduced in this chapter.

Chapter Three — The Literature Review as a Standalone Publication

- Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic Theory*, 3(2), 156–168. [https://doi.org/10.1016/0022-0531\(71\)90013-5](https://doi.org/10.1016/0022-0531(71)90013-5)
- Goldfarb, A. & Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1), 3–43. <https://doi.org/10.1257/jel.20171452>
- Varian, H. R. (1980). A model of sales. *American Economic Review*, 70(4), 651–659.

Chapter Four — The Case Study as an Academic Publication

Primary model text

- Farronato, C., & Fradkin, A. (2022). The welfare effects of peer entry: The case of Airbnb and the accommodation industry. *American Economic Review*, 112(6), 1782–1817. <https://doi.org/10.1257/aer.20180260>

Section 4.9 — Ukrainian and regional context papers

- Bauer, D. & Schwan, S. (2018). Expertise influences meaning-making with renaissance portraits: Evidence from gaze and thinking-aloud. *Psychology of Aesthetics, Creativity, and the Arts*, 12(2), 193–204. <https://doi.org/10.1037/aca0000117>
- Chala, N. & Poplavska, O. (2025). Resilient teams in a volatile world: The impact of global megatrends on employee motivation and HR strategies. In O. Sylkin (Comp.), *Planning and ensuring sustainable development of socio-economic systems: Materials of the 1st International Scientific and Practical Conference, 18 December 2025* (Part 1, pp. 42–43). WSHIU Akademia Nauk Stosowanych. <https://ekmair.ukma.edu.ua/handle/123456789/38062>
- Chala, N., Pichyk, K., & Voropai, O. K. (2021). Practice of neuromarketing research on artwork consumption: Literature review and marketing management perspective. *Bulletin of Khmelnytskyi National University / Visnyk Khmelnytskoho Natsionalnoho Universytetu*, 1(290), 159–164. <https://doi.org/10.31891/2307-5740-2021-290-1-28>
- Lopuschnyak, H., Chala, N., & Poplavska, O. (2021). Socio-economic determinants of the ecosystem of sustainable development of Ukraine. *IOP Conference Series: Earth and Environmental Science*, 915(1), 012019. <https://doi.org/10.1088/1755-1315/915/1/012019>

Chapter Five — Data Interpretation and Quantitative Writing

- Chetty, R., Hendren, N., Kline, P., & Saez, E. (2014). Where is the land of opportunity? The geography of intergenerational mobility in the United States. *Quarterly Journal of Economics*, 129(4), 1553–1623. <https://doi.org/10.1093/qje/qju022>
- Herndon, T., Ash, M., & Pollin, R. (2014). Does high public debt consistently stifle economic growth? A critique of Reinhart and Rogoff. *Cambridge Journal of Economics*, 38(2), 257–279. <https://doi.org/10.1093/cje/bet075>

Reinhart, C. M. & Rogoff, K. S. (2010). Growth in a time of debt. *American Economic Review: Papers and Proceedings*, 100(2), 573–578. <https://doi.org/10.1257/aer.100.2.573>

Chapter Six — Citation, Attribution, and Academic Integrity

Note: Chapter Six uses examples drawn from sources already cited in earlier chapters — Friedman (1968), Goldfarb & Tucker (2019), Farronato & Fradkin (2022), Chetty et al. (2014), Chala & Poplavska (2025), and the Reinhart & Rogoff (2010) / Herndon, Ash, & Pollin (2014)

Chapter Seven — The Writing Process

Note: Chapter Seven introduces no new sources. The exemplar papers cited — Friedman (1968), Goldfarb & Tucker (2019), Farronato & Fradkin (2022), and Chetty et al. (2014)

In-text citation quick reference (APA 7th edition)

Situation	Format	Example from this guide
-----------	--------	-------------------------

One author, paraphrase	(Author, Year)	(Friedman, 1968)
------------------------	----------------	------------------

Two authors, paraphrase	(Author & Author, Year)	(Goldfarb & Tucker, 2019)
-------------------------	-------------------------	---------------------------

Three or more authors, first and subsequent citations	(First Author et al., Year)	(Lopuschnyak et al., 2021)
---	-----------------------------	----------------------------

Direct quotation — must include page	(Author, Year, p. X) or (Author, Year, pp. X–Y)	(Farronato & Fradkin, 2022, p. 1782)
--------------------------------------	---	--------------------------------------

Author named in sentence, paraphrase	Author (Year) ...	Goldfarb and Tucker (2019) argue that...
--------------------------------------	-------------------	--

Author named in sentence, direct quotation	Author (Year, p. X) "..."	Friedman (1968, p. 8) writes that "..."
--	---------------------------	---

Work with no DOI (print only)	Omit DOI; include volume/issue/pages as normal	Neugeboren and Jacobson (2005)
-------------------------------	--	--------------------------------