Proof Abundance

Understanding, Attribution, and Reuse

Flyxion
Independent Researcher

November 2026

Preface

This monograph is a companion to *Portable Warrant* [8]. That work asks when a judgment, once certified in one setting, may be relied on in another. The present work begins at the moment of arrival. A certificate has crossed from its source to a recipient, it checks, and the recipient now holds a proposition that may be used within a stated scope. The question pursued here is narrower and, in practice, more demanding: under what conditions does a proof give its recipient the capacity to make a further mathematical move?

Proof abundance, the condition in which certified proofs are cheap to produce and to check, is what makes this question pressing. The phrase the use of a proof has two senses. In the first, to use a proof is to rely on what it establishes, to cite the theorem and proceed within the limits of what was shown. In the second, to use a proof is to be equipped by it, to adapt its construction to a neighbouring case, to see which hypothesis carries the weight, to recognize in advance where a similar argument will fail. A single certified judgment can be the product of many proofs, and the proofs differ in what they leave their readers able to do. The book studies how these two senses of use travel together, how they come apart, and what follows for attribution, for the economics of review, and for mathematical practice when proofs become cheap to generate and to check.

The independence of certified warrant and continuation gain is proved as a theorem about an explicit model in Appendix B, and the mathematical examples illustrate it without replacing the proof. The argument is built around one running example, small enough to be worked completely: the infinitude of the primes, established by Euclid's construction and by Euler's analytic argument, together with a family of continuation tasks drawn from primes in arithmetic progressions. Part I fixes the scope of a proof and what verification discards. Part II treats mathematical understanding as a resource-bounded capacity for continuation. Part III treats mathematical work as a gift whose value is returned through reuse, and examines attribution and early disclosure. Part IV turns to proof abundance and ends with a complete dossier for a single theorem. The formal apparatus is collected in three appendices: certificate syntax and reliance, the bounded continuation calculus, and the record forms for checking, transmission, and answerability. Wherever a formal claim is made it is proved or its proof is indicated; wherever an illustration is constructed rather than historical, it is labelled as constructed.

Contents

1 Introduction: Two Senses of Use **1**

I The Scope of a Proof **3**

2 The Certified Judgment **5**
2.1 Judgments, Certificates, and Scope 5
2.2 A First Failure of Transfer 6

3 What Verification Forgets **7**
3.1 Fibers of the Certification Map 7
3.2 Euclid and Euler 7
3.3 What the Fiber Contains 8

4 Transport and the Limits of Composition **9**
4.1 Interfaces for Warrant 9
4.2 A Composition Result 9
4.3 Individually Warranted, Jointly Unwarranted 10

II The Transfer of Mathematical Understanding **11**

5 Understanding as a Capacity for Continuation **13**
5.1 A Task-Relative Account 13
5.2 Why the Budget Is Essential 14
5.3 The Task Family for the Running Example 15

6 Aspect Relegation in Mathematical Exposition **17**
6.1 Checking a Choice and Making a Choice 17
6.2 Relegation 17
6.3 A Worked Example 18
6.4 The Failure Case 19
6.5 What the Reader Acquires 19

7 Sketches, Certificates, and Open Obligations **21**
7.1 Four Kinds of Artifact 21
7.2 The Independence Result 22
7.3 Mathematical Illustrations 22
7.4 Interpretation and Limits 23

III Mathematical Work as a Gift **25**

8 Reciprocity Through Reuse **27**
8.1 The Gift Reading 27
8.2 Two Graphs 27
8.3 A Constructed History 28
8.4 Reciprocity Without Repayment 29

9 Attribution and Answerability **31**
9.1 A Case With Four Parties 31
9.2 What a Signature Commits Its Bearer To 31
9.3 Recovering the Dependencies 33

10 The Conditions of Early Disclosure **35**
10.1 The Problem 35
10.2 A Modest Model 35
10.3 Interaction Among Holders 36
10.4 Assumptions and Evidence 36

IV Proof Abundance **39**

11 Generation, Selection, and Reuse **41**
11.1 Where the Costs Fall 41
11.2 A Finite Model 41
11.3 When Abundance Helps 42

12 A Worked Dossier **43**
12.1 Purpose and Status 43
12.2 The Claim and Its Scope 43
12.3 The Stipulated Artifact and the Decisive Construction 43
12.4 The Provenance Record 44
12.5 The Held-Out Case 44
12.6 A Boundary Case 44
12.7 Two Objections 45
12.8 What the Reader May Accept and What the Reader Can Do 45

12.9 After Verification 46

A Certificate Syntax and Scoped Reliance **47**

B A Bounded Continuation Calculus **51**

C Checking, Transmission, and Answerability **55**

Chapter 1

Introduction: Two Senses of Use

A proof is an artifact that moves between people. It is written by one mathematician, read by another, cited by a third, and formalized by a fourth, and at each handoff something is transferred and something is lost. The standard account of what is transferred is thin: the proof establishes its theorem, the theorem is thereafter available, and anyone who trusts the proof may build on the theorem. This account is correct as far as it goes, and it explains why mathematics can be cumulative at all. It does not explain why mathematicians read proofs they already believe, why a new proof of a known result can be a major event, or why a flawless certificate can leave a reader no better equipped than before.

Several strands of the literature press on exactly this gap. Rav argued that proofs are the vehicles of mathematical method and that their importance cannot be reduced to the verification of their conclusions [21]. Thurston described the progress of mathematics as the advance of human understanding and held that the transmission of ways of thinking, more than the accumulation of certified statements, is what a community of mathematicians produces [23]. De Millo, Lipton, and Perlis stressed that the confidence mathematicians place in a theorem is a social product, formed by reading, rederiving, and reusing, and that it differs in kind from the mechanical checking of a derivation [3]. Avigad, writing from within formal verification, proposed that understanding a proof is a matter of abilities, among them the ability to explain a step, to anticipate where an argument will go, and to adapt it to a new case [1]. Lakatos showed how proofs generate their own refinements by exposing the hypotheses they silently use [15].

The contribution of this book is to give these observations a common formal frame in which they can be stated precisely and, in some cases, proved. The frame has three components. The first is the certified judgment, the object that survives verification, together with the scope within which it may be relied on. The second is the continuation task, a mathematical move that a reader might want to make after receiving an artifact, such as adapting a construction or locating the premise on which it depends. The third is a resource bound, without which the frame cannot register that a proof teaches anything at all. A logically omniscient reader already has every consequence of what is known, and for tasks whose acceptance depends only on derivability, the only content of a proof for such a reader

is its conclusion. Real readers have bounded time, and the point of a good proof is that it shortens the route to the next result.

Two consequences organize what follows. First, certified warrant and continuation gain are independent in a precise sense: there are artifacts that supply the first without the second, the second without the first, both, and neither. Second, this independence has costs and obligations that fall in different places. The cost of generating and checking certified proofs is falling, and may fall much further, while the cost of identifying which methods a body of proofs contains, and of making them available to readers, is borne by people with finite attention. A literature that answers every question with a certificate and no account of how the certificate was found would be valid and nearly useless as a medium of instruction. The aim of the book is to say, with enough precision to be tested, why.

The subtitle's word reuse and the word continuation, which Part II uses, name the same thing from two sides. To continue from a proof is to make a further mathematical move with what it supplies, and to reuse a proof is to put what it supplies to work in a problem it was not written for. Two kinds of reuse should be kept apart. Reuse of a result is reliance: the theorem is cited and applied within its scope, which is the first sense of use. Reuse of a method is continuation: the construction, heuristic, or technique is taken up and applied to a new instance, which is the second sense. The book uses continuation for the capacity a single reader acquires, since that is what the continuation sets of Chapter 4 measure, and reuse for the circulation of that capacity among readers, since that is what Chapters 7 through 9 follow as methods pass between authors and attribution tracks them. Where the two words appear in the same passage, continuation names what a reader can do and reuse names what a community does with it.

Plan of the Book

Chapters 2 through 6 form the formal spine. Chapter 2 studies the map from proofs to certified judgments and the fibers it creates. Chapter 3 treats transport of warrant between environments and shows that composition of individually justified transfers can fail. Chapter 4 defines understanding as a resource-bounded capacity for continuation. Chapter 5 treats the relegation of search to routine and compares an author-supplied choice with a reconstructed way of making the choice. Chapter 6 proves the independence of certified warrant and continuation gain within an explicit model. Chapters 7 through 9 concern the social side of circulation: reciprocity through reuse, attribution and answerability, and the conditions of early disclosure. Chapters 10 and 11 treat proof abundance, first by a finite model and then by a worked dossier.

Part I

The Scope of a Proof

Chapter 2

The Certified Judgment

2.1 Judgments, Certificates, and Scope

Let $\mathcal{F}$ be a formal environment: a language, a set of axioms or definitions, and a set of inference rules, together with a checking procedure that decides whether a given object is a derivation. A judgment in $\mathcal{F}$ has the form $\Gamma \vdash_{\mathcal{F}} \varphi$, read as the assertion that φ follows from the hypotheses Γ in $\mathcal{F}$. A proof object is a derivation that the checking procedure accepts, and the certified judgment of a proof object is the judgment it derives. Write $\mathsf{Prf}_{\mathcal{F}}$ for the collection of proof objects of $\mathcal{F}$, $\mathsf{Jdg}_{\mathcal{F}}$ for the collection of judgments, and $\kappa\colon \mathsf{Prf}_{\mathcal{F}} \to \mathsf{Jdg}_{\mathcal{F}}$ for the map that sends a proof object to the judgment it certifies.

A certified judgment is not yet something a recipient may use. Reliance requires knowing the environment in which the judgment was established, the proposition itself, the use that is intended, and any obligations that remain open. The full Scope-Carrying Certificate of the companion volume binds a property identity, an observational projection, an admitted domain, a proof reference, environmental assumptions, temporal validity, composition scope, disposition, issuer, and provenance. Only a reduced form is needed for the mathematical setting, and it is a quintuple.

Definition 2.1 (Scope-carrying certificate)**.** A scope-carrying certificate is a tuple $c = (\mathcal{F}, \Gamma, \varphi, U, O)$ in which $\Gamma \vdash_{\mathcal{F}} \varphi$ is a certified judgment, U is a set of declared uses, each a pair (i, ψ) of an identifier i and a statement ψ of the source environment $\mathcal{F}$ for which the certificate is offered, and O is a set of open obligations, each a pair (i, χ) attaching to the use i a proposition χ whose truth that use depends on and which the derivation does not establish. Declared uses live in the source environment, and the recipient's environment may lack the vocabulary to state them; reliance therefore passes through a translation, as Appendix A makes precise. A recipient may rely on c for a use $u = (i, \psi)$ exactly when the reliance rule of Appendix A applies: the certificate is backed by a checked proof, the translation to the recipient's environment preserves the source judgment, the use is declared, the recipient derives the translated context and the translated obligations $(i, \chi) \in O$ indexed by i, and the translated conclusion together with those obligations entails the translated use. Obligations attached to other uses do not bear on u.

The definition is deliberately modest. It does not say how U and O are determined, since that depends on the environment and the community, but it records that they are part of what is certified. A certificate stripped of its scope has been converted from an object about which responsibility can be assigned into a bare string.

2.2 A First Failure of Transfer

The first example is a certificate that checks after import and nonetheless does not support the recipient's intended application. It is chosen to be unremarkable, because the point is that such failures do not require exotic circumstances.

Example 2.2 (Checks but does not serve (constructed))**.** Let c_0 be an imported certificate for the judgment $\vdash \forall n \in \mathbb{N}\ \exists p > n\ \mathrm{Prime}(p)$, established in an environment of first-order arithmetic. A recipient intends to use it to supply, for every n, a prime $p > n$ with $p \equiv 3 \pmod 4$. The certificate checks successfully in the recipient's environment, and no step of the derivation is in error. If the recipient's environment is ordinary first-order arithmetic, the intended statement is independently provable there, so the entailment premise of the reliance rule is met; the failure is one of declaration, since the use is not among the declared uses U of c_0. If the recipient's environment is a weaker theory whose only relevant axiom is the imported judgment, there is also a genuine entailment failure, with a countermodel, as Appendix A shows. The first case is worth keeping: it shows what the scope record contributes even when the recipient already has enough mathematics to prove the intended result.

Two features of this example recur. The failure is detected not by the checker but by a comparison between the imported judgment and the intended use, a comparison the checker was never asked to perform. And the repair is not to improve the certificate: any further proof of the same judgment, however elegant, leaves the gap exactly where it was. What the recipient needs is either a different judgment or a way of producing one, and which proof of the original judgment is on hand may determine whether the second is feasible. This observation is the seed of the next chapter.

Chapter 3

What Verification Forgets

3.1 Fibers of the Certification Map

The map $\kappa\colon \mathsf{Prf}_{\mathcal{F}} \to \mathsf{Jdg}_{\mathcal{F}}$ is, in every interesting environment, far from injective. Its fiber over a judgment J is the collection $\kappa^{-1}(J)$ of all proof objects that certify J. Proofs in the same fiber may differ in their intermediate lemmas, in the objects they construct, in the representation of the data, and in the choices that are made along the way.

Definition 3.1 (Fiber-invariant property)**.** A property Π of proof objects is fiber-invariant if $\kappa(p) = \kappa(q)$ implies $\Pi(p) \Leftrightarrow \Pi(q)$. Equivalently, Π factors through κ.

Properties such as "establishes a judgment that is true in the standard model" or "may be cited as a lemma in an environment that includes $\mathcal{F}$" are fiber-invariant. Properties such as "uses no appeal to the axiom of choice", "provides an explicit bound", "adapts to a modified hypothesis", or "exhibits the object whose existence is asserted" are not. The following observation is elementary and is recorded because it states the problem exactly.

Proposition 3.2. *Every property of an artifact that depends only on its certified judgment is constant on fibers of κ. Consequently, any account of the value of a proof that is stated in terms of certified judgments alone cannot distinguish proofs of the same judgment.*

Proof. If Π depends only on $\kappa(p)$ then $\Pi(p) = \Pi'(\kappa(p))$ for some Π' on judgments, and $\kappa(p) = \kappa(q)$ gives $\Pi(p) = \Pi(q)$. The second sentence is the contrapositive reading of the first applied to the property "is the more useful of the two proofs." □

The content of the proposition is not the triviality of its proof but the fact that the standard account of what a proof gives its recipient is of the kind it describes. Everything that varies within a fiber is invisible to it.

3.2 Euclid and Euler

Two proofs of the same statement illustrate what varies within a fiber. Let J_∞ be the judgment that for every n there is a prime larger than n. Euclid's proof [5] takes the primes

up to n, multiplies them, adds one, and observes that any prime factor of the result is not among them. Euler's argument [6] considers the product over primes of $(1 - p^{-1})^{-1}$, identifies it formally with the harmonic series, and concludes that the sum of $1/p$ over primes diverges, from which infinitely many primes follow. Both certify J_∞, and rigorous versions of each are routine in contemporary formal systems.

The proofs lead in different directions. Euclid's construction produces an explicit number $N = p_1 \cdots p_k + 1$, all of whose prime factors lie outside a given finite set, and so it invites variation in the choice of N. It gives, with a little more work, the bound $p_{k+1} \le p_1 \cdots p_k + 1$ and hence a doubly exponential upper bound on the kth prime. Euler's argument produces a statement about the density of the primes, namely that the series of reciprocals diverges, and so it rules out, for instance, that the nth prime is eventually bounded below by $n^{1+\varepsilon}$ for a fixed $\varepsilon > 0$. It also invites generalization along an entirely different axis, since replacing the harmonic series with a Dirichlet series twisted by a character leads toward the theorem of Dirichlet on primes in arithmetic progressions [4]. A reader of the first proof is positioned to ask what happens if N is chosen differently. A reader of the second is positioned to ask what happens if the sum is twisted.

Remark 3.3 (Relation to the fiber criterion). Proposition 3.2 is the fiber criterion of the companion volume applied to proof objects: a property Φ survives a transformation T exactly when $\Phi = \widehat{\Phi} \circ T$ on the admitted domain. Here T is the certification map κ and the properties of interest are those that bear on continuation.

Remark 3.4. Nothing in this comparison ranks the two proofs. The point is that the available continuations differ, and that the difference does not appear in $\kappa(p)$. The recurring example of the book is therefore the triple consisting of the single certified judgment J_∞, the two proofs E (Euclid) and H (Euler) in its fiber, and a family $\mathcal{T}$ of continuation tasks on which the proofs can be compared. The tasks are introduced in Chapter 4 and developed in Chapter 5.

3.3 What the Fiber Contains

A fiber contains more than the proofs a community would recognize as distinct. It contains, for instance, proofs that differ only in a choice of representation, and proofs that carry unnecessary hypotheses along. Within a fiber one can define an equivalence of proofs by what they permit a reader to do, and the structure of that equivalence is the object of study of Part II. It is worth noting at this stage that the relevant equivalence is coarser than identity of proof objects and finer than identity of certified judgments. A treatment of the fiber that quotients by the former loses the intensional information that makes a proof reusable, and a treatment that quotients by the latter loses everything.

Chapter 4

Transport and the Limits of Composition

4.1 Interfaces for Warrant

A warrant moves from a source to a destination through an interface. Typically the interface is a translation τ from the language of the source environment to that of the destination, together with a statement of the uses for which the translation is claimed to preserve what matters. Examples include the passage from a classical to a constructive environment, the embedding of a theory in a stronger one, the replacement of an informal definition by a formal one, and the re-expression of a result about real numbers in terms of a particular representation.

Definition 4.1 (Faithful transfer)**.** Let $\tau\colon \mathcal{F}_1 \to \mathcal{F}_2$ be a translation and let L_1 be a class of statements of $\mathcal{F}_1$ that the translation is declared to serve. The transfer is faithful on L_1 if for every $\varphi \in L_1$, $\vdash_{\mathcal{F}_1} \varphi$ implies $\vdash_{\mathcal{F}_2} \tau(\varphi)$ (soundness of transport), and every $\varphi \in L_1$ with $\vdash_{\mathcal{F}_2} \tau(\varphi)$ satisfies $\vdash_{\mathcal{F}_1} \varphi$ (reflection).

A classical example is the conservation result of Friedman for Peano arithmetic over Heyting arithmetic: for sentences of the form $\forall x \exists y\, \theta$ with θ decidable, provability in the classical theory coincides with provability in the intuitionistic one [9]. The transfer from classical to intuitionistic arithmetic is faithful on this class, and, as is well known, the conservation does not extend to the next level of the arithmetical hierarchy. The class L_1 is thus part of the warrant, and a transfer carries no guarantee beyond it.

4.2 A Composition Result

If a transfer τ_1 is faithful on L_1 and a second transfer $\tau_2\colon \mathcal{F}_2 \to \mathcal{F}_3$ is faithful on $L_2 \subseteq \mathrm{Stmt}(\mathcal{F}_2)$, the composite $\tau_2 \circ \tau_1$ should be faithful on some class, and the problem is to identify it.

Proposition 4.2 (Composition under compatibility)**.** *Let τ_1 be faithful on L_1 and τ_2 be faithful on L_2. Put $L_{12} = \{\varphi \in L_1 : \tau_1(\varphi) \in L_2\}$. Then $\tau_2 \circ \tau_1$ is faithful on L_{12}.*

Proof. Let $\varphi \in L_{12}$. If $\vdash_{\mathcal{F}_1} \varphi$ then $\vdash_{\mathcal{F}_2} \tau_1(\varphi)$ by faithfulness of τ_1 on L_1, and then $\vdash_{\mathcal{F}_3} \tau_2(\tau_1(\varphi))$ by faithfulness of τ_2 on L_2, because $\tau_1(\varphi) \in L_2$. For reflection, suppose $\vdash_{\mathcal{F}_3} \tau_2(\tau_1(\varphi))$. Since $\tau_1(\varphi) \in L_2$, reflection for τ_2 gives $\vdash_{\mathcal{F}_2} \tau_1(\varphi)$, and since $\varphi \in L_1$, reflection for τ_1 gives $\vdash_{\mathcal{F}_1} \varphi$. □

The proposition identifies the compatibility hypothesis: the image of the first transfer's class must lie inside the second transfer's class. The hypothesis is not automatic, and its failure is the mechanism of the counterexample.

4.3 Individually Warranted, Jointly Unwarranted

Example 4.3 (Non-transitive warrant (constructed))**.** Take three propositional environments $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3$ over the atoms P and Q. Let P be an axiom of $\mathcal{F}_1$ and $\mathcal{F}_2$ and not of $\mathcal{F}_3$, and let Q be an axiom of $\mathcal{F}_2$ and $\mathcal{F}_3$ and not of $\mathcal{F}_1$, with nothing else derivable that changes this. Use identity translations $\tau_1 \colon \mathcal{F}_1 \to \mathcal{F}_2$ and $\tau_2 \colon \mathcal{F}_2 \to \mathcal{F}_3$. The first transfer is faithful on $L_1 = \{P\}$, since P is derivable on both sides. The second is faithful on $L_2 = \{Q\}$, since Q is derivable on both sides. Each has a nonempty faithful scope. The composite is nevertheless unwarranted for P: P is derivable in $\mathcal{F}_1$ and not in $\mathcal{F}_3$, so soundness of transport fails. This is consistent with Proposition 4.2, since $L_{12} = \{\varphi \in L_1 : \tau_1(\varphi) \in L_2\} = \varnothing$, and it makes the role of the intermediate interface explicit: the second link was never offered for P.

The example does not rely on arithmetic. A real instance of the same pattern is the conservation result of Friedman noted above, which holds for one class of sentences and not the next. The scope guaranteed by the declared links of a chain is the intersection of the pulled-back classes. A composite may preserve additional statements, and a link may leave the guaranteed scope unchanged, but nothing beyond the intersection is guaranteed, and each link can shrink the set of statements on which reliance is justified. The composite should carry, in its scope record, the class L_{12} rather than either of the original classes. This is the bridge to the treatment of non-transitive warrants in the companion volume: the failure of transitivity is a failure of scope to propagate, and it is detectable because scope is part of the certificate.

Part II

The Transfer of Mathematical Understanding

Chapter 5

Understanding as a Capacity for Continuation

5.1 A Task-Relative Account

The preceding chapters show that the certified judgment omits what distinguishes proofs of the same statement. The question now is how to say what is omitted without importing a vague notion of insight. The account adopted here is task-relative. After receiving an artifact, a reader can attempt various mathematical moves: adapt a construction to a new instance, identify the premise on which a step depends, repair the argument after a hypothesis is changed, or produce a counterexample when a condition is dropped. A reader understands a proof, for the purposes of a given family of tasks, to the extent that the artifact enlarges the set of tasks the reader can complete. This follows the general direction of Avigad's treatment of understanding as a collection of abilities [1], and gives it a form in which comparisons can be made and propositions proved.

Three parameters are fixed. The first is a background B, the statements, definitions, and techniques already available to the reader. The second is a family $\mathcal{T}$ of tasks, each a pair consisting of an instance and a goal. The third is a resource budget r, a bound on the amount of derivation, search, and checking that the reader may spend on any one task. The budget may be measured in inference steps, in time, or in any other quantity that is monotone in effort.

Definition 5.1 (Continuation set and gain)**.** For an artifact A, write $\mathsf{Cont}_r(B; A) \subseteq \mathcal{T}$ for the set of tasks a reader with background B and access to A can complete within budget r, and $\mathsf{Cont}_r(B; \varnothing)$ for the corresponding set with no artifact. The gain of A is

$$\mathsf{Gain}_r(A) = \mathsf{Cont}_r(B; A) \setminus \mathsf{Cont}_r(B; \varnothing).$$

The empty-artifact baseline measures total gain, and under it receiving a theorem and applying it as a black box can produce gain. The comparison that matters for the difference between proofs is narrower. Let A_J be the artifact that supplies only the judgment J, with

the same scope and the same mode of access as A. The *method gain* of A is

$$\mathsf{Gain}_r^J(A) = \mathsf{Cont}_r(B; A) \setminus \mathsf{Cont}_r(B; A_J),$$

which asks what becomes feasible through the richer artifact. If reading or importing an artifact costs anything, the cost is charged on both sides of the comparison. Assumption (M1) below is the special case in which the judgment-only baseline completes none of the chosen tasks, so that $\mathsf{Cont}_r(B; A_J) = \varnothing$ and total gain and method gain coincide.

The definition is relative to B, $\mathcal{T}$, and r in all three parameters, and none of these relativizations can be dropped. A proof that is illuminating to a graduate student may teach an expert nothing, and a proof that teaches an expert everything may be opaque to a beginner for whom the relevant tasks lie beyond any feasible budget.

5.2 Why the Budget Is Essential

Proposition 5.2 (Omniscient collapse)**.** *Call a task* derivation-insensitive *if its acceptance condition depends only on which statements are derivable from the background and the certified judgment, and not on the particular derivation supplied. Suppose every task in $\mathcal{T}$ is derivation-insensitive and the budget is unbounded. Then $\mathsf{Gain}_\infty(A)$ depends on A only through $\kappa(A)$. In particular, Gain_∞ is constant on fibers of κ.*

Proof. With unbounded budget, the reader completes exactly the derivation-insensitive tasks whose goals follow from the derivable statements. If $\kappa(A) = \kappa(A')$ then the sets of derivable statements coincide, so $\mathsf{Cont}_\infty(B; A) = \mathsf{Cont}_\infty(B; A')$, and the empty-artifact term in the definition of gain is the same in both cases. □

The hypothesis is essential. Tasks such as identifying the decisive step in the supplied proof depend on the particular proof even with unlimited time, and for them unbounded budget does not collapse the fiber. The proposition is the formal counterpart of Proposition 3.2 for derivation-insensitive tasks. For such tasks every proof of the same judgment is equivalent to every other for a logically omniscient reader, because whatever Euclid's construction could teach about them is already available from the bare statement and the background. A model that cannot register the difference between E and H in Chapter 2 has idealized away the very resource that makes the difference visible. For real readers the budget is finite, and a proof is valuable because it lowers the cost of reaching tasks that would otherwise be out of reach.

Remark 5.3 (Gain can rise and fall with the budget). A task that the artifact makes feasible at some budget enters the gain set at that budget. A task that the background alone eventually reaches leaves the gain set once r is large enough. A task that the background never reaches, however large the budget, stays in the gain set at every larger budget: if an artifact makes a task possible at budget 5 and the background alone never suffices, the task remains in $\mathsf{Gain}_r(A)$ for all $r \geq 5$. No single peak at an intermediate budget is therefore implied. As $r \to \infty$ the gain of each derivation-insensitive task stabilizes pointwise at its

fiber-level status, as in Proposition 5.2. On an infinite task family every task's status may eventually stabilize while the gain sets are still not eventually constant: a task may enter $\mathsf{Gain}_r(A)$ at some r_0 and stay thereafter, while a new task enters at a larger r_1, and so on without end. For tasks the background reaches by itself, the artifact's contribution is confined to a range of budgets, which is where a budget-limited reader benefits from it.

5.3 The Task Family for the Running Example

For the running example the task family $\mathcal{T}$ concerns primes in arithmetic progressions, and the background B is elementary number theory through quadratic residues, with the multiplicativity of residue classes. Formally, a task is a triple $t = (x, g, \alpha)$ of an instance x, a goal g, and an acceptance criterion α that decides whether a candidate output achieves g at x; completing t within budget r means producing an output accepted by α within that budget. The tasks are grouped by the kind of move they require.

Adaptation. Given a residue class a mod m, produce an argument that there are infinitely many primes $p \equiv a \pmod{m}$ by modifying a given construction. The case of interest is a *held-out* class, one for which the artifact supplies no construction, so that the task tests transfer and not reproduction.

Localization. Given a proof, identify the premise on which a designated step depends, for example the property of the constructed number that forces its prime factors into the class.

Repair. Given a proof and a change of hypothesis, such as replacing "prime" by "prime in a stated class," restore the argument or determine that it cannot be restored.

Counterexample. Given a proof and a dropped condition, produce an instance in which the conclusion fails or the argument breaks.

Density. Given a lower or upper growth hypothesis on the sequence of primes, decide whether the hypothesis is compatible with what the proof establishes.

The table is deliberately qualitative. Its purpose is to show that the comparison can be made task by task and that it comes out differently for different tasks, which is what the fiber analysis predicted. The remaining chapters of Part II examine why the entries differ and what would be required for an artifact to improve its row.

Task type	Euclid (E)	Euler (H)
Adaptation to classes with $a^2 \equiv 1 \pmod{m}$	within reach	not directly
Localization of the coprimality premise	within reach	not applicable
Repair after restriction to a class	within reach	via characters
Counterexample when condition dropped	within reach	within reach
Density hypotheses on the primes	not directly	within reach

Table 5.1: A constructed table, illustrative of the form of the comparison and not derived entry by entry, for a reader whose background includes quadratic residues but not L-functions. Two entries are illustrated in Chapter 6 under (M2), where the completion of the task is argued: the density row for H, as a conditional inference from the sketch's divergence claim, and the adaptation row for E, given the exposition of Chapter 5. Formal realizability is established only for the finite model of Appendix B.

Chapter 6

Aspect Relegation in Mathematical Exposition

6.1 Checking a Choice and Making a Choice

Many proofs turn on a single choice. In Euclid's argument the choice is the number N; in a proof of a convergence theorem it may be a sequence; in a proof of existence it may be an auxiliary object chosen from a family. Such a choice-point has a structure that the formalization of the certified judgment suppresses. At the choice-point the author supplies an object c in a space $C(x)$ that depends on the instance x, and the remainder of the proof verifies that c has the properties required. A reader can then do two different things. The reader can check that the supplied c does the job for the instance at hand. Or the reader can make a suitable choice in a new instance x', when no choice has been supplied.

Checking is cheap and local. Making is a search whose cost depends on how much the reader knows about what makes a choice suitable. A proof that only supplies $c(x^*)$ for the original instance x^* lets the reader check but not make. The point of the chapter is to compare it with an exposition that also states the criterion by which $c(x^*)$ was found.

Definition 6.1 (Two kinds of artifact at a choice-point)**.** A *certificate with choice* supplies a particular $c(x^*) \in C(x^*)$ and a derivation that it meets the requirement $\rho(x^*, c)$. An *exposition with search* supplies, in addition, a criterion σ stating what a choice must satisfy in order to meet ρ, in a form from which a reader can generate a suitable $c(x')$ for nearby instances x' within a stated budget.

6.2 Relegation

Relegation is defined here without reference to any wider theory. Let $\mathrm{cost}_t(n)$ be the budget a reader spends on a task t of a given type after n rehearsals on instances of that type. A capacity is *relegated* for the reader when $\mathrm{cost}_t(n)$ has fallen below a fraction ϵ of r, with ϵ fixed by the reader's context, so that exercising it leaves the budget free for the next move.

Search for a suitable object is the standard case: effortful at first, routine after practice. The notion is drawn from the author's work on dual-process accounts of cognition [7], but nothing below depends on that work. The relevant observation for exposition is that relegation is a property of readers rather than of proofs, but proofs can be written so as to promote or to prevent it. A certificate with choice presents the search as already completed and offers no handle on it. An exposition with search presents the criterion, and a reader who rehearses it on several instances will find that the choice has become automatic.

6.3 A Worked Example

Consider the problem of showing that there are infinitely many primes in the class $a \bmod m$, for classes with $a^2 \equiv 1 \pmod m$. Let P be the product of the primes in the class that lie below a given bound.

Example 6.2 (Three supplied instances and one held out)**.** For $a = 3$, $m = 4$ take $N = 4P - 1$. Then $N \equiv 3 \pmod 4$ and N is coprime to every prime in the list, since $N \equiv -1$ modulo each of them. If every prime factor of N were $\equiv 1 \pmod 4$ then $N \equiv 1 \pmod 4$, so some prime factor lies in the class 3 mod 4 and not in the list. This is the multiplication of residues.

For $a = 1$, $m = 3$ take $x = 3P$ and $N = x^2 + x + 1$. Since $N \equiv 1$ modulo every prime dividing $3P$, N is coprime to $3P$. A prime q dividing N divides $x^3 - 1 = (x - 1)N$. If $x \equiv 1 \pmod q$ then $N \equiv 3 \pmod q$, forcing $q = 3$, which is excluded. So x has multiplicative order 3 modulo q, and the order divides $q - 1$, giving $q \equiv 1 \pmod 3$. This is the forcing of a multiplicative order. For $a = 1$, $m = 4$ take $x = 2P$ and $N = x^2 + 1$, which is odd and coprime to P; a prime $q \mid N$ divides $x^4 - 1$, and if x had order dividing 2 then $N \equiv 2 \pmod q$ would force $q = 2$, so x has order 4 and $q \equiv 1 \pmod 4$.

The class $a = 1$, $m = 5$ is held out. The supplied exposition ends here, and the reader is asked to produce a suitable choice of N for this class; the worked solution is given in Chapter 12.

The supplied choices look unrelated when presented as certificates with choice, and the held-out choice is not among them. They instantiate a design criterion: choose N coprime to the listed primes, and such that an algebraic constraint on N forces its factorization to contain at least one prime of the class. A criterion of this kind states what success requires. It is not yet a procedure. The transferable material is how the exposition makes the condition attainable, through multiplication of residues in one family and forcing of a multiplicative order in another. A reader who has the order argument can produce the held-out choice for 1 mod 5, and a reader who has only the first certificate cannot.

Proposition 6.3. *Let $P_r(B; A)$ be the set of instances on which a reader with background B and budget r, holding the artifact A, can produce a valid choice, and let S_A be the set of instances for which A supplies a choice. The instances of new production beyond the supplied cases are exactly*

$$P_r(B; A) \setminus (P_r(B; \varnothing) \cup S_A).$$

A stated criterion is one way to make this set nonempty. It is not necessary, since worked examples can lead a reader to infer a new choice by analogy, and it is not sufficient, since a criterion can leave the search for a choice beyond the budget.

Proof. The set is the definition of production by A that the reader could not already achieve and that A did not simply hand over. For non-necessity, a reader shown the choices for $(a, m) = (3, 4)$ and $(1, 4)$ may infer the choice $N = 9P^2 + 3P + 1$ for $(1, 3)$ by analogy without a stated criterion. For non-sufficiency, a criterion that says only that N must be coprime to the listed primes and have a prime factor in the class, with no indication of how to find such an N, leaves a search whose cost may exceed r. □

The distinction the proposition preserves is between giving a reader a condition that a choice must satisfy and giving them a way to find such a choice.

6.4 The Failure Case

The criterion has a domain, and the examples locate part of its boundary directly. For the class 2 mod 5 a residue condition on N does not force a prime factor in the class: $87 = 3 \cdot 29$ has prime factors in the classes 3 and 4 modulo 5, and $3 \cdot 4 \equiv 2 \pmod 5$, so a number in the class 2 mod 5 can have no factor in that class. This much the examples establish, and it is what makes the class a plausible boundary case. The broader statement that no Euclidean-style argument exists for such a class is a different claim. Call a proof Euclidean-style if, from any finite list of primes in the class, it builds a value of a fixed integer polynomial whose prime factors must include a prime of the class outside the list. That claim is attributed to Murty [19], whose theorem concerns a specific notion of Euclidean proof. This book does not establish that its definition matches the hypotheses of that theorem, and nothing in the argument depends on it; the elementary obstruction carries the boundary example. Dirichlet's theorem establishes infinitude in that progression by a different route [4], one that descends from the ideas of Euler's argument. An exposition that states the criterion should state its boundary, since the boundary tells the reader when to stop looking for an N and to change method.

6.5 What the Reader Acquires

The conclusion of the chapter concerns capacity and not statement. What matters is that the reader acquires the ability to produce the choice at an instance the artifact left open, whether the rule is stated or reconstructed from examples, and that the reader knows where the ability ends.

Chapter 7

Sketches, Certificates, and Open Obligations

7.1 Four Kinds of Artifact

The two notions at issue can now be stated in the same frame. For an artifact A write $C(A) = 1$ if A is a complete certificate of the judgment J, that is, if it is a proof object accepted by the checker with no open obligations in the sense of Definition 2.1, and $C(A) = 0$ otherwise. Write $S(A) = 1$ if $\mathsf{Gain}_r(A) \neq \varnothing$ for the fixed background, task family, and budget, and $S(A) = 0$ otherwise. The first measures certified warrant, the second continuation gain.

The model is subject to one assumption, which expresses that the tasks of interest are not already solvable from the statement of the judgment.

(M1) No task in $\mathcal{T}$ is completed within budget r from $B \cup \{J\}$ alone, where $J = J_\infty$ is the judgment that there are arbitrarily large primes.

The positive cases require a reader to be specified, since continuation gain is a statement about what a reader can complete.

(M2) The reader has background B containing the comparison test for series, the multiplicativity of residue classes, and the facts that for a prime q and an integer x not divisible by q the multiplicative order of x modulo q divides $q-1$ and divides any k with $x^k \equiv 1 \pmod{q}$. The budget r permits derivations of about the length of Example 6.2. The assumption (M2) also supposes that the reader finds the derivations below within budget. A derivation that fits within the budget once found is evidence that an accepted output exists; the claim that the reader can find it within budget requires a search rule or a separate assumption, and (M2) is that assumption. Appendix B gives a second witness in which the search is made explicit. Two tasks in $\mathcal{T}$ are used. The task t_{dens} has instance the hypothesis H_ε that $p_n \geq n^{1+\varepsilon}$ for all large n, with goal to decide whether H_ε is compatible with the density claim made by the artifact; its acceptance criterion accepts a valid conditional inference that cites the claim as stated.

The density illustration is a smaller example of gain from a sketch: the divergence claim it conditions on is supplied by the sketch, and if the validator would accept the same conditional comparison when the reader formulates the antecedent independently, (M1) is a search assumption about that comparison and not a consequence of the missing claim. The task t_{hold} has instance the held-out class $1 \bmod 5$ and goal a Euclidean-style argument for it; its acceptance criterion accepts a valid derivation.

7.2 The Independence Result

Appendix B gives the cost model and a single-task finite instance of the four cases.

Theorem 7.1 (Independence of certified warrant and continuation gain)**.** *There exists a bounded continuation model with one fixed background, task family, and budget in which all four pairs $(C(A), S(A))$ occur. On the subdomain of certified artifacts, S does not factor through the map sending an artifact to its certified judgment, where the certified judgment of an artifact is that of its proof field, when it has one.*

Proof. This is Theorem B.1 of Appendix B. The model there is explicit, and the proof is a finite check of four artifacts against a stipulated reader. □

7.3 Mathematical Illustrations

The theorem is a statement about a model, and the appendix supplies its operational premise, namely what the reader can find within budget. In the illustrations the judgment-only baseline A_J of Chapter 4 is the artifact A_{10}, which exposes no method. The artifacts below give the four combinations mathematical content, for a reader specified by (M1) and (M2). They illustrate the model and exercise it on the running example; they do not replace the proof.

Case $(1, 0)$. Let A_{10} be a machine-checkable proof term of J_∞, delivered as an opaque object that the reader can check but whose internal construction is not exposed within budget r (for instance, a hash-resolved compiled term with no source). It is accepted by the checker with no open obligations, so $C = 1$. Within budget r the reader obtains only the fact that J_∞ holds, and by (M1) neither t_{dens} nor t_{hold} is completed, so $S = 0$.

Case $(0, 1)$. Let A_{01} be Euler's argument in the form of the heuristic manipulation of the product over primes $\prod_p (1 - p^{-1})^{-1}$ and the harmonic series, presented without justification of the interchange of sum and product and without a treatment of convergence. The open obligations are nonempty, so $C = 0$. The sketch asserts, flagged as resting on the unjustified interchange, that $\sum_n 1/p_n$ diverges. The reader completes t_{dens} as follows. If $p_n \geq n^{1+\varepsilon}$ for all n beyond some n_0, then $\sum_{n>n_0} 1/p_n \leq \sum_{n>n_0} n^{-(1+\varepsilon)} < \infty$ by the comparison test, so $\sum_n 1/p_n$ converges. The output "if the sketch's divergence claim holds, then H_ε fails" is a valid conditional inference accepted by the criterion, and it remains valid however the open obligation is later settled. It is an accepted output within budget that is not available from $B \cup \{J_\infty\}$ by (M1), so $\mathsf{Gain}_r(A_{01}) \neq \varnothing$ and $S = 1$.

Case $(1,1)$. Let A_{11} be the proof object of A_{10} accompanied by an exposition that gives the construction $N = 4P - 1$ for 3 mod 4 together with the order argument for 1 mod 3 and 1 mod 4 of Example 6.2, and states the boundary at classes with $a^2 \not\equiv 1 \pmod{m}$. Then $C = 1$ because the proof object is complete. For S, the held-out class 1 mod 5 is not among the supplied constructions. The reader completes t_{hold} by transferring the order argument, including the choice of the polynomial and the treatment of the exceptional prime, both left to the reader; the solution is worked in Chapter 12. The derivation is accepted within budget under (M2), and by (M1) it is not available from $B \cup \{J_\infty\}$, so $S = 1$. This is a transfer to an instance whose solution the artifact leaves open, which is the claim the independence theorem is meant to illustrate, and it returns in Chapter 11.

Case $(0,0)$. Let A_{00} be the assertion that the infinitude of the primes is well known. No proof object is supplied, so $C = 0$, and the assertion gives the reader nothing beyond J_∞, so by (M1) no task is completed and $S = 0$.

In these illustrations A_{10} and A_{11} certify the same judgment and differ in S, as the theorem asserts for the model. Relative to the judgment-only baseline A_{10}, the method gain $\mathsf{Gain}_r^J(A_{11})$ contains the held-out task, and $\mathsf{Gain}_r^J(A_{10})$ is empty.

7.4 Interpretation and Limits

Case $(0,1)$ is historically instructive and not only logically available. Euler's argument circulated for roughly a century in a form that did not meet modern standards, and it was nonetheless the source of the idea from which Dirichlet derived his theorem on primes in arithmetic progressions [4, 6]. The continuation gain preceded the certification, and the open obligations were discharged afterward by others. A policy that admitted into the literature only artifacts with $C = 1$ would have excluded the sketch from which the most productive continuation descended. Case $(1,0)$ is the opposite hazard, now increasingly common: the formal certificate is complete, the reader is warranted in relying on the theorem, and nothing has been transmitted about how the theorem was reached.

The theorem is established within an explicit model and should be read as such. It shows that the independence is possible under stated assumptions, and it identifies the construction of the model's artifacts. It does not estimate how often the cases occur in practice, and the assignment of S depends on the choice of B, $\mathcal{T}$, and r in ways that a different community might set differently. What it does settle is that no policy defined on certified judgments alone can guarantee continuation gain, and that the distinction between the cases is one that a community which cares about the continuation of mathematics must track by other means. The case $(1,1)$ is therefore the standard to which exposition should aspire: a checked proof accompanied by an account of the decisive construction and the conditions under which it generalizes. Chapters 7 through 11 examine the social and economic conditions under which such artifacts are, and are not, produced.

Part III

Mathematical Work as a Gift

Chapter 8

Reciprocity Through Reuse

8.1 The Gift Reading

Part II treated what an artifact gives a reader. Part III asks what the reader owes in return, and the question is not as odd for mathematics as it might be for other trades. Hyde's study of the gift describes an economy in which a gift is not repaid to the donor but passed on, so that its value is realized by movement and diminished by hoarding [13]. Mauss's earlier analysis of gift exchange stressed that what is given carries something of the giver, and that the obligation it creates is to continue the circulation, not to cancel it [17]. Harris has read the practice of mathematics in this light: an idea received from one person is repaid, in the first instance, by using it in work addressed to someone else [12].

The reading fits the observations of Part II. A method can be passed on only to a reader who has received it, and Theorem 7.1 shows that receiving a method is not guaranteed by receiving a certificate. A community in which certificates circulate freely and methods do not is a community in which the gift is withheld while the ledger is kept scrupulously balanced. The gift reading also supplies a vocabulary for what the dependency structure of the literature records and what it leaves out.

8.2 Two Graphs

Record forms for both relations, and for answerability, are given in Appendix C.

Fix a collection of mathematical works W. Two relations between works are of interest.

Definition 8.1 (Checking and transmission relations)**.** The checking relation D_{chk} has an edge $X \to Y$ when the proof object of Y discharges a premise by reference to a scope-carrying certificate issued for X, with the use made of it inside the declared uses of that certificate, so that the final judgment of Y cannot be checked without it. The edges concern the dependencies of the submitted derivation: a proof reference establishes an authorized import in that derivation, not that the import is indispensable to the judgment, since the judgment may have another derivation and the submitted one may contain a removable reference.

For an informal work, the edges are those of the lemmas on which the argument logically depends, whether or not they are cited. The transmission relation D_{trn} has an edge $X \to Y$ when the author of Y acquired from X a construction, heuristic, or technique that equipped the author to make a move in Y that would otherwise have exceeded the author's budget, in the sense of Definition 5.1. Reliance on a result of X as a black box is recorded in D_{chk} and not in D_{trn}, although access to the result may itself enlarge what the author can do; the relation records how a method traveled, not which statements were available.

The first relation is determined by the artifact when the work is a formal proof object, and can then be read off from its references; for an informal work it requires a reconstruction of logical dependence; it is a relation of logical dependence and not of citation practice. The second is a counterfactual relation about the author and cannot.

Proposition 8.2. *Neither $D_{\mathrm{chk}} \subseteq D_{\mathrm{trn}}$ nor $D_{\mathrm{trn}} \subseteq D_{\mathrm{chk}}$ holds in general, and D_{trn} is not determined by the collection of proof objects.*

Proof. For the first non-inclusion, let Y import a lemma from X and use it as a black box; the edge lies in D_{chk}, and the author of Y acquired from X the statement of the lemma but no construction, heuristic, or technique, so the edge is absent from D_{trn}. For the second, let the author of Y learn a construction from the lecture notes of X, which prove nothing that Y cites, and then rederive the needed statements independently; the edge lies in D_{trn} and not in D_{chk}. For the last assertion, take two histories with identical collections of proof objects, in one of which the author of Y read X and in the other of which the author arrived at the construction alone. The proof objects do not distinguish the histories, and D_{trn} does. □

8.3 A Constructed History

The following small history is constructed to make the distinction concrete, and it should not be read as a report of events. A lecturer states in notes the criterion of Chapter 5 for choosing the auxiliary number in Euclid-style arguments, together with its boundary. A first paper uses the criterion to prove the cases of Example 6.2. A second paper, addressed to a different problem, concerns the prime divisors of the values of a quadratic polynomial, and it applies the criterion in the following way: the author recognizes that the constraint forcing prime factors into a class is a constraint on the discriminant, and uses that recognition to organize the whole argument. The second paper cites the first paper for a lemma and does not cite the notes. In D_{chk} the history has a single edge, from the first paper to the second. For the transmission relation the history stipulates the counterfactual budget facts: without the notes, neither author could have completed the key task within the author's actual budget, and the second author could have completed it without the first paper. With these stipulations D_{trn} has an edge from the notes to the first paper and an edge from the notes to the second paper, and no edge from the first paper to the second. Without the stipulations the same acquisitions would be recorded only as reported acquisitions, in the sense of Appendix C. The difference is the content of the unacknowledged part of the gift.

A history-sensitive account earns its place here. The checking graph is sufficient for the question of what the final judgment depends on. It is not sufficient for the question of how the method traveled, and that is the question whose answer determines where to look for the next person able to continue the work.

8.4 Reciprocity Without Repayment

The gift reading suggests a norm with two parts. The first is that a reader who has been equipped by a work should pass on the equipment, in a form usable by others, in work addressed to third parties. The second is that the transmission edge should be acknowledged where known, since the acknowledgment is the information that lets later readers trace the method to its source. Neither part requires repayment to the donor. Both are ordinarily honoured imperfectly, and the quantity of mathematics that has circulated in lecture notes, seminar talks, and correspondence without leaving a trace in D_{chk} is a measure of how much the literature underreports. The collaborative project that proved the density version of the Hales–Jewett theorem [20] is an instance of an institution designed to make D_{trn} visible, since contributions were made in a public forum in which the order and origin of ideas were recorded.

Chapter 9

Attribution and Answerability

9.1 A Case With Four Parties

Consider a published theorem in which a definition comes from one source, the central construction from another, one step of the proof was proposed by a language model, and a human author has checked the whole and signed the result. The case is chosen because it contains every difficulty that attribution presents, and its structure is not specific to any particular technology. A reader encountering the published claim has three needs. The reader needs to recover the intellectual dependencies, to know who contributed what, and to find someone able to respond to an error or an objection. The first two needs concern the record of the work. The third concerns the standing of the person who signed it.

9.2 What a Signature Commits Its Bearer To

Authorship conventions in the sciences distinguish contribution from accountability. The recommendations of the International Committee of Medical Journal Editors, for example, tie authorship to substantial contribution and to agreement to be accountable for all aspects of the work, and state that tools which cannot assume accountability do not qualify as authors [14]; the Committee on Publication Ethics takes a similar position on artificial intelligence tools [2]. The point of interest here is the structure of the commitment. A signature on a mathematical claim binds its bearer to answer for the claim within its declared scope, and the scope is the one carried by the certificate in the sense of Definition 2.1.

Definition 9.1 (Answerability)**.** A party p is answerable for a claim φ with scope (U, O) if for every well-formed objection o to φ within that scope, p can, within a reasonable budget, discharge the objection, repair the claim, or show the objection sound and withdraw the claim; withdrawal without such a showing does not count, as made precise in Appendix C. The party is modelled as a reader in the sense of Chapter 4, with its own background B_p and budget r_p, and holding the artifact A; to discharge, repair, or withdraw within a reasonable budget means to complete the corresponding task within r_p. Objections are of two kinds. A

certificate-level objection alleges that the claim does not follow, that the formal statement does not match the informal one, or that an obligation is open. The first and third are decidable from the checked artifact and its trusted base; the second is not, since the checker decides derivability of the formal statement and cannot decide whether that statement expresses the published claim. A *method-level* objection alleges that the construction does not apply in a stated nearby case, that the premise credited with the result is not the one that carries it, or that a stated generalization fails.

Proposition 9.2. *Let a signer have access to a checked artifact and its trusted base, with background B_p and budget r_p. For an objection o the signer is answerable exactly when at least one response task is completed within budget, that is, when*

$$\{t(o, a) : a \in \mathrm{Response}(o)\} \cap \mathsf{Cont}_{r_p}(B_p; A) \neq \varnothing.$$

(a) A certificate-level objection decidable from the artifact and its trusted base, such as the objection that a proof term does not establish its stated formal judgment, is answered provided the cost of resolving dependencies, running the check, and returning an accepted response fits within r_p; decidability alone supplies no such bound. (b) An objection that the formal statement does not express the published claim is also certificate-level, but the checker alone cannot settle it; the signer must compare the formal statement with the informal claim, which the artifact does not supply. (c) A method-level objection may admit a repair even when it cannot be discharged, so one completed response suffices. Membership may hold from background alone, with a response task in $\mathsf{Cont}_{r_p}(B_p; \varnothing)$, in which case the response is feasible and contributes no gain. For a specified method-level objection, the signer cannot answer if every accepted response to it lies outside the signer's budget given the judgment and the trusted base; this is a fact about that objection and not a consequence of (M1), which concerns the chosen prime continuation tasks and places no restriction on responses to objections about the formal judgment.

Proof. The displayed condition is the quantifier of Appendix C read objection by objection: some response a must be completed within r_p. For (a), rerunning the checker decides the objection, and the budget condition is the stated one. For (b), the checker decides derivability of the formal statement and says nothing about its fidelity to the informal claim, so the objection is settled only by the signer's own comparison. For (c), the definition of answerability requires only one accepted response, and a repair task may lie in $\mathsf{Cont}_{r_p}(B_p; A)$ when the discharge task does not. When every accepted response to a specified objection lies outside r_p given the judgment and the trusted base, no path to an accepted output exists for that objection. □

The proposition links the two parts of the book. Reliance on a result requires only that someone can answer certificate-level objections. Reliance on a claim about a method, such as a claim that the construction generalizes, requires that someone can answer method-level ones. A signature on a claim of generality carries the stronger commitment, and a signer who lacks the capacity cannot honestly bear it.

9.3 Recovering the Dependencies

The four parties of the case each stand in a different relation to the claim. The source of the definition is answerable, if at all, for the definition. The source of the construction is answerable for the construction in its original setting, and not for its adaptation. The model that proposed a step has no standing to respond to an objection, and the step is therefore the responsibility of whoever adopted it. The human author who checked and signed is the one party against whom an objection to any part of the claim can be lodged. The historical record of mathematics contains cases that illustrate the last point on a larger scale: when a gap was found in the announced proof of Fermat's last theorem, it was the author who undertook the repair, with a collaborator, and the corrected argument was published in 1995 [24, 22]. The recovery of dependencies, in the sense of the need identified at the start of the chapter, is a matter of recording which party contributed which segment, and it is distinct from the question of who answers. A provenance record that separates the two is a table with one row for each segment of the argument and columns for the source, the role (definition, construction, step, verification), and the answerable party. The table is compact, and its absence is the usual cause of the difficulty described at the start of the chapter.

Chapter 10

The Conditions of Early Disclosure

10.1 The Problem

The gift economy of Chapter 7 presupposes that ideas are given before they are finished. A mathematician who waits until an idea is complete withholds the stage at which feedback and collaboration are most valuable, and one who shares early exposes the idea to a rival who may develop it first. The motivating case is the common one in which a mathematician teaches material that is still under development, so that the notes themselves are a disclosure whose consequences cannot be recalled. Priority disputes are a recurrent feature of the history of mathematics, and the observation of Merton that recognition accrues disproportionately to the already eminent suggests that the cost of early disclosure is not borne equally [18].

10.2 A Modest Model

The model is a decision problem for a single holder of an idea, and its role is to expose the assumptions behind the usual intuition, not to predict behaviour.

Let the idea have value $V > 0$ if it is developed privately. If it is disclosed, feedback and collaboration raise its value by $b \geq 0$. With probability $q \in [0,1]$ a rival develops the disclosed idea first, in which case the holder retains the fraction $1 - \lambda$ of its value, where $\lambda \in [0,1]$ measures the loss of priority credit. The parameter λ depends on attribution norms. Write $\lambda = \lambda_0(1 - \nu)$, where λ_0 is the loss in a community that does not acknowledge transmission and $\nu \in [0,1]$ measures the strength of the norm of acknowledging it, in the sense of Chapter 7.

Proposition 10.1. *Under the assumptions of the model, with $q\lambda < 1$, the expected payoff of disclosure is $(V + b)(1 - q\lambda)$ and that of withholding is V, so disclosure is preferred if and only if*

$$\frac{b}{V} > \frac{q\lambda}{1 - q\lambda}.$$

The threshold on the right-hand side is increasing in q and in λ, and hence decreasing in the strength ν of the acknowledgment norm.

Proof. With probability $1-q$ the holder receives $V+b$, and with probability q the holder receives $(1-\lambda)(V+b)$, so the expected payoff is $(V+b)(1-q\lambda)$. Disclosure is preferred when $(V+b)(1-q\lambda) > V$, which rearranges to $b(1-q\lambda) > q\lambda V$ and gives the stated inequality. The function $x \mapsto x/(1-x)$ is increasing on $[0,1)$, and $x = q\lambda$ is increasing in q and in λ; since $\lambda = \lambda_0(1-\nu)$, it is decreasing in ν. □

10.3 Interaction Among Holders

The interaction that the single-agent model omits can be added in one step. Let $s \in [0,1]$ be the fraction of other holders who disclose, and let the rival probability be $q(s)$. If a community in which more people disclose is one in which more rivals are positioned to act on a disclosure, q is increasing in s. The threshold of Proposition 10.1 then rises with s, so each holder's incentive to disclose falls as others disclose: disclosure is a strategic substitute. Write $D(s) = (V + b(s))(1 - \lambda q(s)) - V$ for the advantage of disclosing when a fraction s of others disclose. Where b and q are differentiable,

$$D'(s) = b'(s)\big(1 - \lambda q(s)\big) - \lambda q'(s)\big(V + b(s)\big).$$

If q rises with s and b is constant, then $D'(s) = -\lambda(V+b)q'(s)$, which is strictly negative where $\lambda > 0$ and $q'(s) > 0$ and weakly negative when $\lambda = 0$ or q is locally flat, so disclosure is a strategic substitute, strictly in the first case. If b also rises with s because disclosure builds a shared forum, as in collaborative projects, the first term can dominate, and disclosure is a strategic complement where it does. If D is continuous and $D(0)$ and $D(1)$ are weakly opposite in sign, there is a fraction s^* with $D(s^*) = 0$, namely a solution of

$$\frac{b(s^*)}{V} = \frac{q(s^*)\lambda}{1 - q(s^*)\lambda}.$$

The absence of a strict sign change does not exclude interior indifference when D reaches or stays at zero. Whether an indifference point is stable depends on an adjustment rule that the model does not specify, and the question is left open. Which case holds is an empirical matter, and it bears on whether raising the acknowledgment norm ν moves a community smoothly or across a threshold.

10.4 Assumptions and Evidence

The model's assumptions are strong and it is useful to list them in order to see what evidence would be needed. It treats b and q as exogenous, whereas the probability that a rival acts depends on how visible and how attractive the idea is, and may depend on b. It treats withholding as safe, whereas an idea kept private can also be developed first by someone else, at a loss of the whole of its value. The model is partial-equilibrium: it treats the holder as acting alone, and the interaction among holders is taken up in the preceding section. And it

treats ν as a fixed feature of the community, whereas it is itself the product of the practices of citation and acknowledgment examined in the preceding chapters.

To establish the mechanism in practice, one would need records of when ideas were first disclosed and in what form, records of subsequent development by others, and records of how credit was assigned, collected for a sufficiently large and comparable body of cases. Open collaborative projects, preprint archives, and the archives of seminars and lecture notes are potential sources. Without such evidence the proposition should be read as a conditional statement, and its use is to show which parameters matter: a community that raises the acknowledgment norm lowers the threshold at which sharing is rational, and thereby increases the supply of unfinished ideas on which the gift economy depends.

Part IV

Proof Abundance

Chapter 11

Generation, Selection, and Reuse

11.1 Where the Costs Fall

The cost of producing a certified proof has two components, the cost of finding a candidate and the cost of checking it. Formal systems have made checking cheap and reliable. Large verification projects, such as the formal proof of the four-color theorem [10] and of the Kepler conjecture [11], and the growth of shared libraries of formalized mathematics [16], have shown that substantial results can be certified mechanically. If the cost of finding candidates also falls, through automated search or through the assistance of learned models, the supply of certified proofs grows, and the constraint that binds is no longer the production of certificates. The constraint becomes the cost of selecting worthwhile claims, of reconstructing the methods that the proofs contain, and of assessing whether those methods transfer. These costs are borne by readers, whose attention does not scale with the supply of certificates.

11.2 A Finite Model

Suppose that in a period N certified proofs appear. Each proof belongs to one of $D(N) \leq N$ method classes, where two proofs belong to the same class when they exhibit the same decisive construction. A reader with a review budget B can identify the method content and the transfer conditions of a proof at a cost c per proof read. Put $k = \min\{N, \lfloor B/c \rfloor\}$, the number of proofs the reader can examine in the period, and assume $k \geq 1$.

Proposition 11.1 (Review-bounded identification)**.** *The number of distinct method classes the reader identifies is at most* $\min\{D(N), k\}$*, whatever the value of* N*. In particular, increasing the supply of certified proofs while holding the review budget fixed does not increase the number of methods identified beyond* k*.*

Proof. Each examined proof reveals at most one class, and at most k proofs are examined, so at most k classes are identified; at most $D(N)$ classes exist. □

The bound is attainable only if the reader examines proofs from distinct classes. Whether this happens depends on how the collection is presented.

Proposition 11.2 (Abundance without organization)**.** *Suppose one class contains $N-D+1$ of the proofs and each of the remaining $D-1$ classes contains exactly one. A reader who examines k proofs chosen uniformly at random without replacement identifies in expectation $k(D-1)/N + (1-p_0)$ classes, where p_0 is the probability that no proof from the dominant class is drawn. This is at most $1 + k(D-1)/N$, with equality when $k > D-1$, since then at least one dominant proof is drawn. As $N \to \infty$ with D and B/c fixed and $\lfloor B/c \rfloor \geq 1$, so that $k = \lfloor B/c \rfloor$ is eventually fixed and positive, it tends to 1. A reader who has access to a class index, so that proofs can be chosen from distinct classes, identifies $\min\{D, k\}$ classes.*

Proof. Each of the $D-1$ rare classes has a single proof, which is among the k drawn with probability exactly k/N, so the expected number of rare classes found is exactly $k(D-1)/N$; the dominant class is found with probability $1-p_0$, and $p_0 = 0$ when $k > D-1$. The statement about the index follows by choosing one proof from each class until the budget is exhausted. □

The propositions identify what abundance does to a body of mathematics. It raises the number of certificates without raising, and possibly lowering, the expected number of methods a reader can reach, unless the collection is organized by method. The index is assumed to exist and to be usable within the reader's budget. If it does, its cost falls on a curator while the cost of its absence falls on every reader; for a collection whose method classes change as new proofs arrive, the curator's cost recurs and lies outside the model.

11.3 When Abundance Helps

Abundance is not harmful in itself, and the model shows when it helps. If the proofs of a statement are diverse, so that $D(N)$ grows with N, then a collection organized by method exposes constructions that a single proof would not. The case of E and H is the prototype: two proofs of one statement open two families of continuations. A body of automatically generated proofs that are all instances of one class would be of little use, and a body in which the proofs differ in their decisive constructions and are indexed accordingly would be of great use. The same quantity, the number of certified proofs, thus has opposite effects depending on the variable that the model isolates, namely the number of distinct classes and the means by which readers can reach them. The analysis recommends no policy beyond this: abundance should be measured not by the count of certificates but by the count of identifiable methods.

Chapter 12

A Worked Dossier

12.1 Purpose and Status

The final chapter tests the central distinction on a single object, the sequence that runs through Chapters 4 to 6: a certified judgment, the construction that produced it, a nearby case whose solution is held out, and an objection to the adaptation. The dossier is constructed, and its status is uniform. The acceptance of the proof object by a checker is stipulated: no proof object or formal environment is displayed here, and the displayed derivation is the mathematical exposition. The provenance record is a schematic of how such a record could look when one step of a formal proof is proposed by a model and the rest is contributed by identified parties. No claim is made that the events it describes occurred.

12.2 The Claim and Its Scope

The claim is that for every natural number n there is a prime $p > n$ with $p \equiv 3 \pmod 4$. It is stated in first-order arithmetic. The declared uses are citation as a lemma in developments over the same environment, and instantiation to produce a prime in the class 3 mod 4 above a given bound. The open obligations are empty. The use that is not declared is any inference about other residue classes.

12.3 The Stipulated Artifact and the Decisive Construction

By stipulation, a checker accepts a proof object for the claim in the stated environment. The decisive construction is the following. Let P be the product of the primes $\equiv 3 \pmod 4$ not exceeding n, with $P = 1$ if there are none, and let $N = 4P - 1$. Then $N \equiv 3 \pmod 4$ and $N > 1$ is odd. Each prime in the list divides P and therefore does not divide N, since $N \equiv -1$ modulo it. If every prime factor of N were $\equiv 1 \pmod 4$ then $N \equiv 1 \pmod 4$, which is false, so N has a prime factor $p \equiv 3 \pmod 4$. By coprimality p is not in the list, so $p > n$.

12.4 The Provenance Record

The record forms for these segments are given in Appendix C, where the model-proposed step is written out as a specimen.

Segment	Role	Answerable party
Definition of the residue class and primality predicates	Definition (from a standard library)	Library maintainers, for the definition
Choice of $N = 4P - 1$	Construction (Euclid's argument, adapted)	Signer, for the adaptation
Proof of the coprimality step	Step proposed by a model, then checked	Signer
Acceptance of the full proof object	Verification (stipulated)	Signer, with the checker as trusted base
Publication of the claim with its scope	Signature	Signer

Table 12.1: A constructed provenance record. The model appears in the record as the source of a segment but not as an answerable party, in accordance with the discussion in Chapter 8.

12.5 The Held-Out Case

The exposition supplies the order argument for 1 mod 3 and 1 mod 4 and leaves the class 1 mod 5 to the reader. A reader who holds the criterion of Chapter 5, that the auxiliary number should be coprime to the listed primes and that an algebraic constraint should force its factorization to contain a prime of the class, and who has seen how a multiplicative order does the forcing, can proceed. Let P be the product of the primes $\equiv 1 \pmod 5$ not exceeding n, let $x = 5P$, and let $N = 1 + x + x^2 + x^3 + x^4$. A prime $q \mid N$ divides neither 5 nor P, since $N \equiv 1$ modulo each prime dividing $5P$. It divides $x^5 - 1 = (x - 1)N$. If $x \equiv 1 \pmod q$ then $N \equiv 5 \pmod q$, which forces $q = 5$, excluded because $5 \mid x$ gives $N \equiv 1 \pmod 5$. So x has order 5 modulo q, and since the order divides $q - 1$, $q \equiv 1 \pmod 5$. Hence $N > 1$ has a prime factor in the class that is not in the list. The artifact transmits the order argument; the choice of polynomial and the treatment of the exceptional prime $q = 5$ are left to the reader.

12.6 A Boundary Case

The construction fails for the class 2 mod 5, and the reader should be able to anticipate this. Here $2^2 = 4 \not\equiv 1 \pmod 5$. The immediate obstruction can be shown concretely: a number

in the class 2 mod 5 can have no prime factor in that class, as in $87 = 3 \cdot 29$, whose factors lie in the classes 3 and 4 with $3 \cdot 4 \equiv 2 \pmod 5$. A condition on the residue of N therefore does not force a factor in the class. This much follows from the example. The broader exclusion of Euclidean-style arguments for the class is attributed to Murty [19], with the caveats stated in Chapter 5; the argument here does not depend on it. Dirichlet's theorem establishes infinitude in this progression by a different route [4], which descends from H. The boundary case is therefore part of what the dossier transmits: a reader who finishes it knows both what the construction does and where to stop using it.

12.7 Two Objections

Answerability in the sense of Appendix C quantifies over all well-formed objections, and the dossier does not establish that. It exhibits answerability to two named objections, one about the formal claim and one about the method's scope.

The first objection o_1 alleges that the formal statement does not express the published claim: that the library's predicate Prime might admit 1 or fail to match the usual definition, so that the published sentence and the checked sentence differ. A discharge response compares the library definition, which by stipulation reads $\text{Prime}(p) \iff p > 1 \wedge \forall d\,(d \mid p \to d = 1 \vee d = p)$, with the usual definition, and compares the formal statement $\forall n\,\exists p > n\,(\text{Prime}(p) \wedge p \equiv 3 \pmod 4)$ with the published sentence. The two comparisons are short and lie within the signer's budget, and the checker cannot perform them, since it decides derivability and not fidelity.

The second objection o_2 alleges that the adaptation to the class 1 mod 5 is unsound because the exceptional prime $q = 5$ may divide N. A discharge response observes that $5 \mid x = 5P$, so $N \equiv 1 \pmod 5$ and $q \neq 5$, which is the step identified in the held-out case above. This is a method-level objection, and the response is a continuation task that the signer completes because the signer holds the order argument.

12.8 What the Reader May Accept and What the Reader Can Do

On the strength of the certificate the reader may accept the claim, instantiate it, and cite it as a lemma within its declared scope. The reader may not infer anything about other residue classes, and the certificate does not say that the construction extends. On the strength of the exposition the reader can do more: transfer the order argument to the held-out class 1 mod 5, locate the coprimality step and the exceptional prime as the premises on which the adaptation depends, anticipate the failure at 2 mod 5, and recognize that a different method is then needed. These are continuation tasks in the sense of Chapter 4, and the dossier supplies the case $(1, 1)$ of Theorem 7.1: a stipulated certificate accompanied by an account of the decisive construction and the conditions under which it generalizes. The signer is shown answerable to the two named objections, which is what the dossier exhibits.

12.9 After Verification

The book began with two senses of use, to rely on a proof and to be equipped by it. The argument has been that a community which attends only to the first will accumulate certificates and be impoverished in methods, and that the impoverishment will be invisible to every measure defined on certified judgments. In an age in which certificates are cheap to generate and to check, the scarce resources are the account of the decisive construction, the statement of the conditions under which it generalizes, the record of how it traveled, and a person able to answer for all of these. Mathematical understanding in such an age is not displaced by verification. It is what remains after verification, and it is carried by readers who take up what they are given and put it to use somewhere new.

Appendix A

Certificate Syntax and Scoped Reliance

This appendix fixes what counts as a well-formed reduced certificate and states a separate judgment for reliance in a recipient environment. The grammar is that of the five-field certificate of Definition 2.1.

```
<certificate> ::= "scc(" <environment> "," <context> "," <formula> ","
                         <uses> "," <obligations> ")"
<uses>        ::= "{" <use-list> "}"
<use-list>    ::= eps | <use> | <use> "," <use-list>
<use>         ::= "use(" <id> "," <formula> ")"
<obligations> ::= "{" <obligation-list> "}"
<obligation-list> ::= eps | <obligation>
                    | <obligation> "," <obligation-list>
<obligation>  ::= "requires(" <id> "," <formula> ")"
<submission>  ::= "submit(" <certificate> "," <proof-ref> ")"
```

Each use identifier i names exactly one formula, so that U is a partial function from identifiers to formulas; without this condition the reliance rule could be applied with one formula under i while the soundness statement concluded another. The grammar permits repeated identifiers with different formulas; well-formedness requires that they be distinct. For a certificate $c = (F, \Gamma, \varphi, U, O)$ and a declared use $u = (i, \psi) \in U$, put $O_i = \{\chi : (i, \chi) \in O\}$, the obligations attached to the use i. The proof reference accompanies the five-field certificate: the `<proof-ref>` of a submission resolves to a proof object p, and write $p \triangleright c$ when p is accepted by the checker of F for the judgment of c. The reliance rule below is stated for the resolved object p. The uses $(i, \psi) \in U$ and obligations $(i, \chi) \in O$ are indexed pairs, with ψ and χ statements of the source environment F; the translation τ carries them into the recipient's environment G, and reliance is available only when τ is defined on them. An interface $\tau\colon F \to G$ translates source expressions, and $\mathrm{Pres}(\tau, J)$ attests that the translation preserves the particular source judgment $J = (\Gamma \vdash_F \varphi)$, in the

sense that $\tau(\Gamma) \vdash_G \tau(\varphi)$. Write $\Delta \vdash_G X$ for pointwise derivability of every member of a set X.

$$\frac{\mathrm{Check}_F(p, \Gamma \vdash \varphi) = 1}{p \triangleright c} \qquad \frac{\begin{array}{cccc} p \triangleright c & \mathrm{Pres}(\tau, \Gamma \vdash_F \varphi) & (i,\psi) \in U & \\ \Delta \vdash_G \tau(\Gamma) & \Delta \vdash_G \tau(O_i) & \tau(\varphi), \tau(O_i) \vdash_G \tau(\psi) & \end{array}}{\Delta \vdash_G \mathrm{rely}(c, i, \tau)}.$$

The left rule says when a proof object belongs to a certificate. The right rule is the reliance judgment: the recipient, whose environment is G with hypotheses Δ, may rely on c for the use i through τ exactly when the certificate is backed by a checked proof, the translation preserves its judgment, the use is declared, the recipient derives the translated context and the translated obligations of that use, and the translated conclusion together with those obligations entails the translated use.

Proposition A.1 (Reliance soundness)**.** *If* $\Delta \vdash_G \mathrm{rely}(c, i, \tau)$ *and* $(i, \psi) \in U$*, then* $\Delta \vdash_G \tau(\psi)$*.*

Proof. Preservation followed by cut. From $\Delta \vdash_G \tau(\Gamma)$ and $\tau(\Gamma) \vdash_G \tau(\varphi)$, which is what Pres attests, cut gives $\Delta \vdash_G \tau(\varphi)$. With $\Delta \vdash_G \tau(O_i)$ and $\tau(\varphi), \tau(O_i) \vdash_G \tau(\psi)$, a second cut gives $\Delta \vdash_G \tau(\psi)$. □

The Import Example

Let

$$J_\infty = \forall n\, \exists p > n\ \mathrm{Prime}(p), \qquad J_{3,4} = \forall n\, \exists p > n\ (\mathrm{Prime}(p) \wedge p \equiv 3 \pmod 4),$$

and let $c_\infty = (F, \Gamma, J_\infty, U_\infty, O_\infty)$ be a certificate backed by a checked proof, $p \triangleright c_\infty$. The recipient intends the use $u_{3,4} = (i_{3,4}, J_{3,4})$. The failures depend on the recipient environment. In ordinary first-order arithmetic with the usual definitions, $J_{3,4}$ is provable, so $J_\infty \vdash_G J_{3,4}$ holds by weakening and the entailment premise is satisfied. The failure that remains in that setting is one of declaration: if $i_{3,4} \notin \mathrm{dom}(U_\infty)$, the reliance rule has no instance for $i_{3,4}$ even though $p \triangleright c_\infty$ holds and the proof checks.

A content failure requires a weaker recipient theory. Let G be first-order logic over a language with a strict order $<$ and two uninterpreted predicates Prime and Q, where τ sends Prime to Prime and "$\equiv 3 \pmod 4$" to Q, and let the only axiom of G be the translation of J_∞. Take the structure with domain $\mathbb{N}$, the usual order, Prime true of every element, and Q true of none. It satisfies the translation of J_∞, since every element has a larger element, and falsifies the translation of $J_{3,4}$, since Q is empty. By soundness of first-order logic, $\tau(J_\infty) \nvdash_G \tau(J_{3,4})$ when $O_{i_{3,4}}$ is empty, so the entailment premise fails. This failure is one of content, and it cannot be repaired by relabelling. It is repaired only by a certificate whose judgment is $J_{3,4}$, which must still satisfy the preservation premise. The countermodel shows that, in that recipient theory with an empty context, it cannot:

Pres would require $\vdash_G \tau(J_{3,4})$, and the structure above falsifies it. The other repair is by an obligation requires$(i_{3,4}, J_{3,4})$ that the recipient must then discharge, which makes the reliance vacuous.

Appendix B

A Bounded Continuation Calculus

This appendix specifies tasks, artifacts, outputs, and a cost model, and then gives the explicit witnesses promised in Chapter 6 for all four artifacts in a single finite model.

```
<task>      ::= "task(" <instance> "," <goal> "," <validator-id> ")"
<artifact>  ::= "artifact(" <certificate-option> "," <hint-set> ")"
<certificate-option> ::= "none" | <submission>
<output>    ::= "output(" <task-id> "," <term> ")"
```

Fix a background B, a task family $\mathcal{T}$, and a budget r. A reader transition $s \xrightarrow[B,A]{w} s'$ has cost $w \in \mathbb{N}_{>0}$ and depends on the background and the artifact. Write $s \xRightarrow[B,A]{\leq r} s'$ for a path of total cost at most r; the output $\mathrm{out}(y)$ is a terminal state of the reader. For a task t, let $s_0(t)$ be the initial reader state and let α_t be its validator, fixed in advance of any attempt. Completion within budget is the judgment

$$\frac{s_0(t) \xRightarrow[B,A]{\leq r} \mathrm{out}(y) \qquad \alpha_t(y) = 1}{B; A \vdash_r t \Downarrow y}.$$

The continuation set and gain of Chapter 4 are then

$$\mathsf{Cont}_r(B; A) = \{t \in \mathcal{T} : \exists y\ B; A \vdash_r t \Downarrow y\}, \qquad \mathsf{Gain}_r(A) = \mathsf{Cont}_r(B; A) \setminus \mathsf{Cont}_r(B; \varnothing).$$

The judgment separates two things that are easily confused. Holding a hint, that is, an idea worth trying, is a property of the artifact. Completing a task is the existence of an output that the validator, fixed beforehand, accepts at cost at most r.

A Finite Witness for Chapter 6

Let $\mathcal{T} = \{t\}$ and $r = 1$, and let t ask for a valid conditional inference from a growth bound on the nth prime. Its validator accepts the output

$$y_* : \qquad (\exists n_0\, \forall n \geq n_0,\ p_n \geq n^{1+\varepsilon}) \Longrightarrow \sum_{n=1}^{\infty} \frac{1}{p_n} < \infty \qquad (\varepsilon > 0),$$

with the tail comparison $1/p_n \leq n^{-1-\varepsilon}$ as witness. Give the constructed reader exactly the following output transition, where h is the hint to compare with a p-series and $H(A)$ is the hint set of A:

$$s_0(t) \xrightarrow[B,A]{1} \mathrm{out}(y_*) \quad \Longleftrightarrow \quad h \in H(A).$$

The reader is stipulated and not derived. The witness shows that the four combinations of C and S are consistent in a model with one background, one task family, and one budget; it does not claim that any particular human reader behaves as the stipulation says. Let p be a complete proof of J_∞ whose internals are inaccessible to this reader within budget.

Artifact	Proof field	Hint set	$C(A)$	$\mathsf{Gain}_1(A)$
A_{00}	none	$\varnothing$	0	$\varnothing$
A_{01}	none	$\{h\}$	0	$\{t\}$
A_{10}	p	$\varnothing$	1	$\varnothing$
A_{11}	p	$\{h\}$	1	$\{t\}$

Theorem B.1 (Independence in the finite model)**.** *In the bounded continuation model above, with one fixed background, task family, and budget, all four pairs $(C(A), S(A))$ occur, where $S(A) = \mathbf{1}[\mathsf{Gain}_1(A) \neq \varnothing]$. On the subdomain of certified artifacts, those with $C(A) = 1$, S does not factor through the map κ sending an artifact to its certified judgment: $\kappa(A_{10}) = \kappa(A_{11}) = J_\infty$ while $S(A_{10}) \neq S(A_{11})$.*

Proof. For A_{00} and A_{10} the hint set is empty, so the stipulated transition does not fire; no path from $s_0(t)$ reaches an output, $\mathsf{Cont}_1(B; A) = \varnothing$, and the gain is empty. For A_{01} and A_{11} the transition fires at cost $1 \leq r$ and produces y_*, which α_t accepts, so $t \in \mathsf{Cont}_1(B; A)$. Since $\mathsf{Cont}_1(B; \varnothing) = \varnothing$ (the empty artifact has the empty hint set), $\mathsf{Gain}_1(A) = \{t\}$. The values of C are read from the proof field, which gives all four pairs by the table. Both artifacts A_{10} and A_{11} carry p, so both certify J_∞, and they differ in S; hence S restricted to certified artifacts is not a function of κ. □

The accepted inference is conditional, so its completion leaves the divergence of the prime reciprocal series open.

A Search Witness

The preceding witness toggles a hint. A second witness makes the search itself inspectable, so that the gain arises from a difference in what the reader enumerates. Give the reader two operations with explicit costs. The operation test(f) costs 1 and reports whether the value $f(5P)$ is coprime to the listed primes and has its prime factors forced into the class 1 mod 5 by a multiplicative-order argument. The operation emit(f) costs 1 and outputs the value $f(5P)$ after a passed test. The reader enumerates a pool of candidate polynomials $f_1, \ldots, f_K$, with $K \geq 3$, in a fixed default order, and the single member that passes the tests is f_K, the polynomial $1 + x + x^2 + x^3 + x^4$. The budget is $r = 2$, and the task t_{hold} asks for a valid choice for the class 1 mod 5, with a validator fixed in advance. In this witness the validator checks an auxiliary-number choice, whereas the Chapter 7 task asks for the full derivation; the cost calculation below measures the choice.

Let A_J be the judgment-only artifact. It adds nothing to the enumeration. The reader tests f_1 and f_2 at total cost 2, both fail, the budget is exhausted, and $t_{\text{hold}} \notin \mathsf{Cont}_2(B; A_J)$. Let A_E be an exposition that supplies the order-forcing family $\{\Phi_m(mx)\}$ as a filter on the pool, where Φ_m is the mth cyclotomic polynomial; the scaling is part of the evaluation step, which evaluates a selected polynomial at $x = mP$. The family is therefore $\{\Phi_m(x)\}$, and the reader restricts the enumeration to the member with m equal to the modulus, which leaves $f_K = \Phi_5$ alone. The reader tests f_K at cost 1 and emits it at cost 1, for a total of $2 \leq r$, and the validator accepts. So $t_{\text{hold}} \in \mathsf{Cont}_2(B; A_E)$ and $\mathsf{Gain}_2^J(A_E) = \{t_{\text{hold}}\}$. Reading either artifact is charged zero on both sides.

The difference between the two artifacts is an inspectable change in the search space: the exposition changes which candidates are enumerated, not whether a single flag fires. The reader and the pool are still constructed, and the position of f_K is stipulated. A reader who infers a rule by analogy from worked examples would show gain from an examples-only artifact; the model does not contain such a reader, and so it separates supplied instances from a supplied filter only for the reader it specifies.

The finite model proves Theorem 7.1; the mathematical examples in Chapter 7 illustrate the four cases.

Appendix C

Checking, Transmission, and Answerability

This appendix gives separate record forms for checking dependencies, reported transmission, and answerability, and states what each attests.

```
<record> ::= "uses(" <work-id> "," <work-id> "," <certificate-id> ","
                     <use-id> ")"
           | "reports(" <person-id> "," <work-id> "," <work-id> ","
                        <method-id> "," <task-id> ")"
           | "proposes(" <source-id> "," <segment-id> ")"
           | "adopts(" <person-id> "," <segment-id> "," <work-id> ")"
           | "signs(" <person-id> "," <work-id> "," <scope-id> ")"
```

Checking and Transmission

For a provenance history H and a task t in a work Y, let $B^{+}_{H,Y}$ be the author's background before attempting t. Let $B^{-X,m}_{H,Y}$ be the corresponding background under removal of the transmission of method m from X, with access to the certified results of X held fixed. Assume that a work issues at most one certificate for each judgment and scope it establishes, so that $\mathrm{issuer}(c) = X$ identifies the certificate's source. The checking relation concerns the dependencies of the submitted derivation: it records an authorized import in that derivation, not that the import is indispensable to the judgment. The checking relation is a relation of logical dependence on an issued certificate:

$$X \to_{\mathrm{chk}} Y \iff \exists c, i \, \big[\mathrm{issuer}(c) = X \wedge \mathrm{ProofRef}(Y, c, i, G, \Delta, \tau) \wedge \mathrm{Authorized}(c, i, G, \Delta, \tau)\big],$$

where $\mathrm{ProofRef}(Y, c, i, G, \Delta, \tau)$ records that Y imports the certificate c at a particular site for the use i, with that site's recipient environment G, context Δ, and translation τ, and $\mathrm{Authorized}(c, i, G, \Delta, \tau)$ is the reliance judgment $\Delta \vdash_G \mathrm{rely}(c, i, \tau)$ of Appendix A at that

site. One work can import different certificates at different sites, with different contexts or translations, so the parameters attach to the individual proof reference and not to the work. The transmission relation is counterfactual:

$$X \to^{H}_{\mathrm{trn}} Y \iff \exists m, t\, [\mathrm{Acquired}_H(X, m, Y) \wedge t \in \mathsf{Cont}_{r_{H,Y,t}}(B^{+}_{H,Y}; \varnothing) \backslash \mathsf{Cont}_{r_{H,Y,t}}(B^{-X,m}_{H,Y}; \varnothing)],$$

where $r_{H,Y,t}$ is the author's actual budget for the task t in Y, fixed before the two backgrounds are compared. With an existential budget, a method that saves a single step would create an edge even when the author could readily have completed the task within the actual budget. A `reports` record establishes only that someone reported an acquisition, the predicate $\mathrm{ReportedAcquisition}_H(X, m, Y)$. The predicate $\mathrm{Acquired}_H(X, m, Y)$ requires in addition that the report be credited, for instance by corroboration. The transmission edge requires $\mathrm{Acquired}_H$ and the budget comparison, and the comparison depends on a model of the author's background absent the method. A record can therefore attest that an author learned a construction from a set of notes, while the stronger claim that the author could not have made the move within budget without those notes requires a counterfactual model of $B^{-X,m}_{H,Y}$ or additional evidence. Without either, a reported acquisition is recorded as a reported edge and not as an established one.

Answerability

For objections $o \in \mathcal{O}(Y, U)$ to a work Y within scope U, set

$$\mathrm{Response}(o) = \{\mathrm{discharge}(o), \mathrm{repair}(o), \mathrm{withdraw}(o)\}.$$

Each response a has an associated task $t(o, a)$ with its own validator fixed in advance. The validator for $\mathrm{discharge}(o)$ accepts a demonstration that the objection does not apply within U, the validator for $\mathrm{repair}(o)$ accepts a modified claim and a derivation that it meets the objection, and the validator for $\mathrm{withdraw}(o)$ accepts an exhibition of the defect that o alleges. The last requirement matters: withdrawal is not a free exit, since a signer who withdrew every claim without showing the objection sound would be answerable for nothing. Then

$$\mathrm{Answerable}(p, Y, U) \iff \forall o \in \mathcal{O}(Y, U)\ \exists a \in \mathrm{Response}(o)\ B_p; A_Y \vdash_{r_p} t(o, a) \Downarrow y \text{ for some } y.$$

Signing assigns responsibility; answerability is the capacity condition above.

A Specimen From the Dossier

The model-proposed coprimality step of Chapter 11 can be recorded as follows, for a work W with signer p, a library Lib supplying the definitions, a model M, and a segment s that contains the coprimality step. The identifiers m_N, $t_{3,4}$, and U_W name the method of choosing $N = 4P - 1$, the adaptation task, and the declared scope.

```
uses(W, Lib, cert_defs, use_defs)
reports(p, Elements_IX_20, W, m_N, t_3_4)
proposes(M, s)
adopts(p, s, W)
signs(p, W, U_W)
```

The first record is a checking dependency and would enter $\rightarrow_{\mathrm{chk}}$ if the certificate is authorized. The second attests a reported acquisition from Euclid's argument; whether it enters $\rightarrow_{\mathrm{trn}}$ depends on the counterfactual comparison, which the specimen does not supply. The third attributes the proposal to M. The fourth and fifth record the commitment undertaken by p, who adopts the step and signs the work, with the proposal attributed to M and the commitment carried by p.

Bibliography

[1] Jeremy Avigad. Understanding Proofs. In Paolo Mancosu, editor, *The Philosophy of Mathematical Practice*, pages 317–353. Oxford University Press, 2008.

[2] COPE Council. Authorship and AI Tools. Committee on Publication Ethics, 2024.

[3] Richard A. De Millo, Richard J. Lipton, and Alan J. Perlis. Social Processes and Proofs of Theorems and Programs. *Communications of the ACM*, 22(5):271–280, 1979.

[4] Peter Gustav Lejeune Dirichlet. Beweis des Satzes, dass jede unbegrenzte arithmetische Progression, deren erstes Glied und Differenz ganze Zahlen ohne gemeinschaftlichen Factor sind, unendlich viele Primzahlen enthalte. *Abhandlungen der Königlich Preussischen Akademie der Wissenschaften*, pages 45–81, 1837.

[5] Euclid. *Elements*, Book IX, Proposition 20.

[6] Leonhard Euler. Variae Observationes Circa Series Infinitas. *Commentarii Academiae Scientiarum Petropolitanae*, 9:160–188, 1744. Presented 1737.

[7] Flyxion. Aspect Relegation Theory and the Myth of Dual Cognition. Unpublished essay, 2026.

[8] Flyxion. The Portable Warrant: Evidence at the Boundary. 2026. `https://github.com/standardgalactic/portable-warrant`.

[9] Harvey Friedman. Classically and Intuitionistically Provably Recursive Functions. In *Higher Set Theory*, volume 669 of *Lecture Notes in Mathematics*, pages 21–27. Springer, 1978.

[10] Georges Gonthier. Formal Proof: The Four-Color Theorem. *Notices of the American Mathematical Society*, 55(11):1382–1393, 2008.

[11] Thomas C. Hales et al. A Formal Proof of the Kepler Conjecture. *Forum of Mathematics, Pi*, 5:e2, 2017.

[12] Michael Harris. *Mathematics Without Apologies: Portrait of a Problematic Vocation*. Princeton University Press, 2015.

[13] Lewis Hyde. *The Gift: Imagination and the Erotic Life of Property*. Random House, 1983.

[14] International Committee of Medical Journal Editors. Defining the Role of Authors and Contributors; and AI Use by Authors. In *Recommendations for the Conduct, Reporting, Editing, and Publication of Scholarly Work in Medical Journals*. `https://www.icmje.org/recommendations/`.

[15] Imre Lakatos. *Proofs and Refutations: The Logic of Mathematical Discovery*. Cambridge University Press, 1976.

[16] The mathlib Community. The Lean Mathematical Library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pages 367–381, 2020.

[17] Marcel Mauss. *The Gift: The Form and Reason for Exchange in Archaic Societies*. Translated by W. D. Halls. Routledge, 1990. Originally published 1925.

[18] Robert K. Merton. The Matthew Effect in Science. *Science*, 159(3810):56–63, 1968.

[19] M. Ram Murty. Primes in Certain Arithmetic Progressions. *Journal of the Madras University, Section B*, 51:161–169, 1988.

[20] D. H. J. Polymath. A New Proof of the Density Hales–Jewett Theorem. *Annals of Mathematics*, 175(3):1283–1327, 2012.

[21] Yehuda Rav. Why Do We Prove Theorems? *Philosophia Mathematica*, 7(3):5–41, 1999.

[22] Richard Taylor and Andrew Wiles. Ring-Theoretic Properties of Certain Hecke Algebras. *Annals of Mathematics*, 141(3):553–572, 1995.

[23] William P. Thurston. On Proof and Progress in Mathematics. *Bulletin of the American Mathematical Society*, 30(2):161–177, 1994.

[24] Andrew Wiles. Modular Elliptic Curves and Fermat's Last Theorem. *Annals of Mathematics*, 141(3):443–551, 1995.