

Article

Not peer-reviewed version

The Economy of Truth: How Resources Shape What Can Be Known

[Arturo Tozzi](#) *

Posted Date: 10 December 2025

doi: 10.20944/preprints202512.0950.v1

Keywords: epistemic agency; inferential constraints; resource-bounded rationality; constructive semantics; procedural justification



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

The Economy of Truth: How Resources Shape What Can Be Known

Arturo Tozzi

ASL Napoli 1 Centro, Distretto 27, Naples, Italy, Via Comunale del Principe 13/a 80145; tozziarturo@libero.it

Abstract

I propose a resource-sensitive theory of truth grounded in the structure of linear logic. On this view, a proposition counts as true for an agent only when it can be derived from the finite informational resources the agent currently possesses such as data, concepts, tools or premises, through a form of inference that tracks how those resources are used. Rather than treating truth as a static or purely metaphysical attribute, my account emphasizes that truth emerges within the constraints of actual reasoning. A key distinction is drawn between local derivability, which concerns what an agent can establish with the resources immediately available and global derivability, which reflects what could be established under idealized conditions. This allows my approach to preserve objectivity while acknowledging that access to truth is often limited and uneven. The framework yields a new perspective on classical epistemological issues, including the nature of justification and the structure of rational inference, while clarifying how inferential breakdowns can occur when the resources supporting a reasoning process are incomplete or unreliable. It also provides a natural lens for examining truth in science, legal reasoning and artificial intelligence, domains in which information is finite, traceable and central to responsible decision-making.

Keywords: epistemic agency; inferential constraints; resource-bounded rationality; constructive semantics; procedural justification

1. Introduction

Philosophers have long taken for granted that our reasoning relies on stable semantic principles, giving truth a fixed and unquestioned place within them, rather than exploring the assumptions that sustain this stability. Classical discussions, whether framed in terms of correspondence (Strawson 1950; Szaif 2018), coherence (Blanshard 1939; Walker 2018), pragmatic success (Peirce 1878; James 1907; Capps 2020) or deflationary minimalism (Ramsey 1927; Quine 1970; Hoffmann 2010), aim to identify what it is in virtue of which a proposition is true. Yet these accounts appeal, tacitly or explicitly, to idealized agents: reasoners who enjoy unrestricted access to information, who operate with unlimited inferential capacity and who evaluate propositions independently of the contingent constraints of context or cognition (Davidson 1984; Alston 1996). If we attend instead to the actual epistemic situation of finite agents, these traditional accounts reveal a structural limitation. Classical logic grants complete liberty in manipulating assumptions: premises may be duplicated, discarded or reintroduced without cost and the informational resources required for an inference remain invisible to the logic itself. From this point of view, truth is treated as a static semantic status that floats free of the procedures by which an agent might come to establish it, obscuring the fact that epistemic agents operate under constraints of memory, information, cognitive load and time. Once these constraints are acknowledged, the classical picture becomes a poor guide to the epistemic grounds on which truth-claims are advanced and sustained.

Recent developments in logic and semantics draw attention to these neglected dimensions. Substructural logics, whether relevant, linear or otherwise, reject the classical presupposition that inference is costless (Restall 2000; Jalali 2021). Proof-theoretic semantics similarly emphasizes that meaning and truth emerge from the inferential practices encoded in introduction and elimination

rules (Prawitz 1971; Schroeder-Heister 1991). These approaches share a methodological shift: rather than treating truth as a primitive semantic relation, they examine how the structure of inference itself determines what counts as evidence for a proposition. In this light, truth becomes inseparable from the procedural and structural features of reasoning.

My proposal develops this shift by presenting truth as an outcome of epistemic activity structured and limited by the resources on which that activity depends. Truth is sensitive to the resource conditions under which agents operate. The question is not merely how propositions relate to facts, beliefs or linguistic conventions, but how they become established as true by agents who must manage and transform informational resources. Although inferentialist and constructivist views gesture toward this direction, none has provided a systematic account of how the availability, consumption and transformation of resources might themselves condition the status of a proposition. By adopting linear logic as a formal framework (Girard 1987), I aim to make this dependence explicit: the derivability of a proposition depends on the specific configuration of resources accessible to the agent who attempts the derivation. In practical reasoning, the formation of truth-claims typically occurs under epistemically restrictive conditions. Agents must draw conclusions from incomplete data, under cognitive and temporal pressures and with tools of inference that may themselves be limited or partial. These facts point to a conception of truth that is not exhausted by an abstract semantic mapping between propositions and the world. Rather, truth functions within epistemic practice as something that can be *established* from what an agent has at hand.

This perspective also casts familiar epistemological issues in a new light. The structure of Gettier-style counterexamples, the limits of justification and the relation between belief and truth all appear differently once one treats inference as a resource-sensitive act rather than as an abstract manipulation of costless premises. What counts as truth for an agent thus depends, in part, on what the agent can *do* with the resources available.

In what follows, I develop this theory formally, examine its implications and show how it reshapes traditional discussions about the nature of truth and its epistemic role.

2. A Resource-Sensitive Theory of Truth

Let me begin by emphasizing a point that is easy to overlook. Linear logic, as introduced by Girard (1987), is not merely a technical variant of classical logic. It embodies a philosophical claim, one might say, a kind of corrective, about what we tacitly assume when we treat inference as an activity placing no demands on its practitioners (Barroso-Nascimento et al., 2025). In classical logic, premises may be copied, forgotten or applied however many times we please. Nothing in the formalism reflects the fact that making an inference is a real epistemic act performed by an agent. Linear logic alters this picture: each assumption must be used exactly once unless it carries an operator explicitly licensing its reuse.

A small example may help. Suppose an agent has a single measurement reading M . Classical logic allows her to employ M in arbitrarily many chains of reasoning, as though evidence were inexhaustible. Linear logic does not. Unless M is explicitly marked as duplicable, the agent who uses it to infer φ cannot simultaneously use it to infer ψ . The logic encodes an epistemic principle: evidence is not always an infinitely renewable resource. This is not a metaphysical thesis but a structural observation about how information behaves when treated realistically (Troelstra 1992; Bellin and Scott 1994).

This distinction extends into the heart of linear logic, where additive and multiplicative conjunctions divide the territory that classical logic treats as uniform. The multiplicative conjunction ($\otimes$) central to the present framework carries an important message. When we write $a \otimes b$, we do not merely list two assumptions; we assert that they are simultaneously required and jointly consumed. A linear implication $a \multimap b$ likewise indicates that b follows from a *only when a itself is spent* in the inferential transition. These are not arbitrary syntactic devices, rather they reflect a fact about epistemic practice: one cannot use a piece of evidence for incompatible purposes unless it is either reproducible or replaced by something functionally equivalent.

With these ideas in place, we can now approach truth from a different angle. Instead of treating truth as something statically assigned by a model or as a metaphysical relation between a proposition and the world, we may treat truth as the result of a constrained derivation carried out by an agent whose informational resources are limited. To say that φ is true for an agent is to say that there exists a derivation of the form

$$a \otimes b \otimes c \multimap \varphi$$

from that agent's current stock of epistemic resources a , b , c . These resources may include empirical data, theoretical assumptions, inferential rules and even instruments. What matters is not their nature but their availability and expendability. Imagine a scientist who has two data points D_1 and D_2 and one inferential rule R , all of which must be used to derive a conclusion φ . If she has all three resources, she may derive φ . If she lacks D_2 , she cannot. It does not follow that φ is false; rather, φ is not *derivable* for her given her present epistemic situation. The distinction between truth and derivability is retained, but derivability becomes a condition for treating something as true *for the agent*. A more formal treatment can be found in the appendix.

It is important to avoid a misunderstanding here. We are not proposing that truth reduces to belief, coherence or even justified belief. We are not giving a psychological account. Instead, the suggestion is that truth is accessed through successful derivations under resource constraints. These derivations may or may not be performed, but their existence depends on what is available to the agent. This aligns my view more closely with proof-theoretic semantics (Dummett 1993; Tennant 1997) than with model-theoretic or deflationary accounts, though it adds a distinctive feature: resource conditions themselves are constitutive of whether a derivation is possible.

My framework yields two modes of derivability. A proposition φ is *locally* derivable when it can be inferred from an agent's actual resource set R . It is *globally* derivable when φ is derivable from an ideal set R^* which embodies the informational completeness or inferential power that real agents lack (Haack 1996). The distinction mirrors familiar contrasts in epistemology between what can be known in principle and what can be established under present conditions.

A brief example underscores the point. Suppose an agent lacks the resource b . The sequent

$$a \otimes b \otimes c \vdash \varphi$$

is not available to her. Another agent, possessing b , can derive φ . Yet φ may be globally derivable whether any particular agent can access it. This explains how truth can remain objective, while its accessibility varies across agents and contexts.

My framework also sheds new light on familiar epistemic puzzles, particularly the class of cases introduced by Gettier (1963). In these cases, an agent appears to possess a justified true belief, yet intuition resists counting it as knowledge. The difficulty traditionally lies in explaining *what exactly has gone wrong*. The problem can be traced to the epistemic profile of the resources from which the belief is derived. A derivation may be syntactically impeccable and even successful in reaching a true conclusion, while nevertheless drawing upon a resource that lacks the requisite epistemic status: for example, a premise that is false, unreliable, improperly sourced or epistemically inert despite being treated as evidentially robust. In these cases, the inference fails a resource-certification requirement: the derivation does not falter in its logical form, but in the epistemic quality of its inputs. This reframes Gettier problems in a revealing way. The intuitive defect in these cases is not a matter of accidental luck appended to an otherwise respectable justification, nor a mysterious gap between justification and truth. What goes wrong is localized and identifiable: the informational ingredient that drives the inference is corrupted, epistemically defective or misclassified as usable. My account thus avoids the need to posit additional epistemic conditions (such as "no false lemmas"), because the failure is already diagnosed within the structure of the derivation. The agent has not merely been unlucky, rather she has employed a resource that should not have supported the relevant inference in the first place (Kvanvig 2003). By specifying which resources are legitimate, my approach pinpoints the failure precisely where it occurs: not in the logical relation between premises and conclusion, but in the epistemic credentials of the resources themselves.

It is important to clarify the scope of my proposal. I am *not* offering a metaphysical thesis about the nature of truth. My view does not claim that truth itself depends on human inferential capacities or that it varies ontologically from agent to agent. Nor does it deny that propositions may be true independently of our ability to derive them. What I am offering is a formal account of how truth appears within epistemic practice, how agents come to possess or fail to possess truths, how this access is mediated by informational resources and how these dependencies can be captured within a logical structure. Seen in this light, the consequences are significant. First, we obtain a non-metaphysical, operational conception of truth. Truth is not only what corresponds to the world, but what can be *reached* from a given resource configuration under appropriate inferential norms (Putnam 1981; Brandom 1994). This is truth as an epistemic achievement: something that an agent may or may not be able to derive, depending on what resources are available, trustworthy and properly deployed.

Second, my framework explains how different agents may stand in different epistemic relations to the same truth without thereby endorsing relativism. The divergence lies not in truth itself, but in the resource structures agents possess. Two agents may differ in their access to evidence, conceptual tools, inferential rules or technological instruments; hence they differ in what they can legitimately derive. Yet when their resources come to coincide or expand, convergence becomes possible and often expected. My view therefore distinguishes between truth *as such* and the epistemic *availability* of truth, offering a principled account of why disagreement, partial understanding and scientific progress are all compatible with the objectivity of truth.

What emerges is an epistemology that places resource structure at the heart of truth-ascription. Truth is not weakened or relativized, rather our access to truth is made explicit, traceable and formally representable. The aim is to model epistemic life as it is lived, bounded, procedural and dependent on what agents have at hand, while still preserving the concept of truth as something that can, in principle, be shared and jointly attained.

In the next chapter, I examine how this framework works across several domains, showing how resource-sensitive derivability reorganizes our understanding of truth, justification and epistemic practice.

3. Examples and Applications

To clarify the scope of my resource-sensitive conception of truth, it is helpful to examine several concrete cases. These examples are not meant as empirical demonstrations, nor as metaphysical illustrations, but as reminders that the epistemic situations in which we make truth claims are structured by what resources we possess. In each example, the shape of the reasoning process becomes crucial: what can be inferred depends on what is available and so the epistemic status of a proposition is indexed to these materials without thereby making truth itself relative.

A first, simple example from everyday practice may help. Consider a chef preparing a signature dish. Let ϕ denote the claim that “the dish will have its intended flavor.” Whether ϕ is assertable for the chef depends on having certain ingredients: fresh basil a , ripe tomatoes b and high-quality olive oil c . When all three are present, the inference to ϕ is entirely unproblematic; the chef relies on stable culinary knowledge and appropriate materials. Nevertheless, when the basil is no longer fresh and the tomatoes are out of season, the same inference is no longer available, even though the recipe has not changed. The epistemic situation has shifted because the relevant resources are missing. What is important for us is not that culinary claims are metaphysically peculiar, but that they illustrate a general point: some truths are accessible only in contexts where the appropriate inputs are available. Nothing in this example suggests that ϕ changes its metaphysical status; what changes is whether the agent can *derive* ϕ . That distinction between metaphysical truth and derivational accessibility is precisely what my theory is designed to track.

Let us turn to a scientific case in which this structure becomes more elaborate. The discovery of the Higgs boson provides a clear illustration. Physicists sought to establish a proposition ϕ : that the Higgs boson exists. Doing so required three indispensable resources. First, a theoretical model a , predictions from the Standard Model specifying what the particle should look like. Second,

experimental data b from high-energy collisions. Third, precision instruments c capable of detecting the relevant decay signatures, such as ATLAS and CMS. The proposition ϕ was derivable only when these components were jointly available and properly analyzed. Had one of them been absent, say the detectors had not yet been developed, the epistemic status of ϕ would have remained indeterminate for us, even if ϕ were in fact true. Prior to the availability of detectors, “The Higgs boson exists” may have been true, but the inferential route to it was blocked (Hacking 1983; Franklin 1994; Chang 2012). In this respect, my framework captures what philosophers of science have long observed: that empirical knowledge advances by the development of instruments, concepts and practices that render previously inaccessible truths epistemically approachable.

Historical scientific transitions show the same structure. Consider the reception of heliocentrism in the 17th century. It is tempting, from our present vantage point, to regard the correctness of heliocentrism as obvious. But for many agents of that time, the proposition “Earth orbits the Sun” was not derivable from the epistemic resources available: telescopic observations a were limited or absent, no physical theory of inertia b could explain planetary motion and a coherent dynamical model c was lacking. Given those resources, the geocentric proposition seemed better supported. When Newtonian mechanics and improved astronomical observations became available, the epistemic status of the claim shifted accordingly: the derivation shifted, enabling the truth claim of heliocentrism to be epistemically established. my approach handles this smoothly. The truth of the heliocentric claim need not to change, what changes is its epistemic accessibility. This allows us to acknowledge historical variation in justified belief without falling into any form of relativism: the world may remain the same, while the epistemic means to engage with it evolve.

These cases, culinary, scientific, historical are meant to do conceptual work. They show that in ordinary and theoretical reasoning alike, the truth we can responsibly assert depends on the availability, reliability and proper employment of informational materials. This point, while simple, has broad implications. It suggests that a procedural or operational conception of truth may capture epistemic practice more faithfully than a model in which truth passively corresponds to a world independently of inquiry.

Bearing this in mind, the applicability of the framework becomes clearer. In legal reasoning, the truth of propositions such as claims about guilt or liability depends on admissible evidence, verified testimony and interpretive resources such as statutory rules and precedents. A legal proposition like “defendant X is guilty” is derivable only when resources a , b , c (e.g., forensic data, witness accounts, authoritative interpretations) are jointly available and properly used. My framework does not imply that guilt is metaphysically dependent on these resources, only that epistemic justification is.

A similar structure appears in artificial intelligence. For an AI system to produce explainable conclusions, its inferences must rely on identifiable, traceable and non-duplicated inputs. Without this transparency, the system provides outputs whose epistemic status cannot be inspected or justified. Linear logic has already been employed to model resource consumption in planning, proof search and automated reasoning (Andreoli 1992; Baelde 2012). My framework extends this computational insight into a philosophical account of how an AI system’s truth-claims ought to be evaluated. On this view, a system may treat ϕ as true only when the resources supporting ϕ are explicitly present, properly certified and consumed in the derivation. This requirement blocks the use of ungrounded or AI “hallucinated” premises, ensuring that every step in the reasoning process is anchored in accessible and verifiable inputs. In doing so, the model contributes to the broader aim of making AI accountable to human-interpretable norms of justification and epistemic responsibility.

Linguistics provides another illustration. Context-sensitive expressions (indexicals, anaphora, temporal markers) depend on discourse resources that must be available at the point of evaluation. Here, too, resource-sensitive reasoning interacts naturally with dynamic semantic frameworks such as Discourse Representation Theory (Kamp and Reyle 1993; Asher and Lascarides 2003). My approach helps articulate why certain inferences involving pronouns or presuppositions succeed only when specific context resources are present.

In educational settings, my approach echoes a familiar point: the correctness of a student's answer depends not only on its match with a canonical truth, but on whether it follows from the resources at the student's disposal: information provided in the problem, conceptual tools previously taught, time available and so on. Instructors implicitly evaluate whether an answer is *derivable* from what the student could reasonably access, rather than from an idealized omniscient standpoint.

Finally, in philosophy of science and formal epistemology, my resource-sensitive approach yields insights about underdetermination, justification, defeasibility and belief revision. Scientific propositions are "true-for-us" when derivable from limited but reliable evidence and theoretical structures. Linear logic's capacity to track how resources are consumed or transformed provides a natural model for non-monotonic reasoning: the arrival of new information may invalidate prior derivations by altering the resource configuration.

Taken together, these cases illustrate a deeper philosophical insight. My framework does not ask us to rethink what truth *is*. Instead, it asks us to reconsider the epistemic pathways through which truth becomes *available*. My resource-sensitive theory of truth provides a logic for when and how propositions become epistemically accessible, making explicit the dependence of justification and belief on dynamically evolving informational structures. Truth, on this account, is not static: it is something that must be achieved and what can be achieved depends on the resources at hand. The following chapter engages with the most significant challenges my account is likely to encounter.

4. Challenges and Responses

Any proposal departing from classical theories of truth encounters philosophical resistance. My resource-sensitive approach makes certain structural commitments about derivability, about the role of informational inputs and about the situatedness of inference, differing from the assumptions built into classical logics. It is therefore important to consider the main challenges one might raise, not merely to deflect them, but to clarify what is and is not being claimed. In what follows, I address some concerns, pausing at various points to remark on the conceptual terrain in which they operate.

A first objection concerns the apparent agent-relativity of my approach. Truth, it may be urged, should not depend on the cognitive limitations or resource deficits of any individual. To tie truth to what an agent can construct risks collapsing truth into belief or justification. That is a legitimate worry and one worth formulating carefully. What my framework proposes, however, is not that truth depends on the agent in any metaphysical sense, but that what counts as *derivable*—and thus epistemically claimable—depends on what the agent has at hand. The distinction between *local* derivability (from an agent's actual resource set) and *global* derivability (from an idealized, complete resource configuration) is designed precisely to avoid relativism. It mirrors distinctions between actual and ideal justification in formal epistemology (Pollock and Cruz 1999; Williamson 2000) and between operational and theoretical availability in philosophy of science (Cartwright 1983). So, the point is not that truth is up for grabs, but that access to truth varies.

A second objection targets the absence of external truth conditions. Traditional correspondence theories insist that truth must be grounded in a relation between propositions and the world. If we emphasize derivability from informational resources rather than a metaphysical target, do we not risk a kind of internal circularity? The answer, I think, requires some conceptual disentangling. My proposal is not anti-realist; it does not deny that propositions may be true independently of us. Rather, it focuses on *how* those truths become epistemically available. The informational resources on which derivations depend (observations, instruments, testimonies) are themselves worldly: they are ways the world imposes structure on our epistemic position. To treat these as resources is simply to formalize their epistemic role. In this respect, my approach resonates with constructivist and pragmatic currents (Niiniluoto 1999). Logic does not displace the world, rather codifies the means by which the world becomes accessible.

A third concern is that Gettier-style problems persist (Gettier 1963). Even if we track resources carefully, might an agent not derive a true proposition from unreliable or defective premises? This is a serious challenge and one we need to address with care. My resource-sensitive account incorporates a notion of *resource certification*: premises can be annotated with epistemic provenance, indicating

empirical verification, testimonial credibility or defeasibility conditions. A derivation that uses a resource lacking proper certification does not qualify as truth-conferring. This allows us to locate the failure where it belongs: not in the logical form, but in the epistemic status of the resources. Extensions of the framework into modal or temporal logic allow resources to be tagged as “tentative,” “verified,” or “contingent” (Miller 2008; Rott 2001). Rather than evade epistemic luck, my proposal analyzes it: luck infiltrates through the resource, not the inference.

A fourth concern is that linear logic appears too technical or specialized for philosophical work. One might be tempted to think that, because linear logic originated in computer science, its applicability is limited to technical domains, leaving philosophy to more traditional logic. Here a historical remark is helpful. Modal logic, intuitionistic logic and relevance logic each began as specialized tools and later reshaped entire philosophical landscapes (Restall 2000). Linear logic, with its emphasis on resource sensitivity, follows this trajectory. Its technical apparatus serves a philosophical purpose: making explicit what classical logic suppresses, namely, that reasoning is not free of cost. And the logic’s structural features integrate naturally with constructivist and inferentialist considerations familiar from Martin-Löf (1996) and Schroeder-Heister (2006). Formalism is not a barrier, but a clarification.

A fifth objection questions operational applicability. In everyday reasoning, it might be said, we do not explicitly keep track of resources, nor do we label each inference step with the data on which it depends. Does this not render my proposal impractical? Here again, conceptual clarification helps. The aim is not to mirror informal reasoning step by step, but to provide a structure that makes explicit what is implicit in practices requiring rigor. Scientific explanation, legal argument, mathematical proof and AI computation all demand transparency about how conclusions are reached. In these domains, resource tracking is not an artificial imposition but a precondition for accountability (van Benthem 2011; Girard 2001). My approach therefore does not distort practice, rather formalizes the very features of practice that classical logic overlooks.

Finally, one may claim that this is not a theory of truth at all, but only a theory of justification. The charge is understandable: after all, we speak of what agents can derive. But the conclusion does not follow. My proposal concerns the conditions under which a proposition can be *treated as true* by a rational, resource-bounded agent. It is procedural in character, but it is still about truth. Indeed, it reframes the philosophical project: not to identify truth in the abstract, but to understand the routes by which truth becomes epistemically accessible. This places my account within inferentialist and procedural traditions (Peregrin 2014; Boghossian 2003), while adding further insight that the structure of access is shaped by resource constraints.

Across these challenges, the same pattern recurs. When the objections are formulated precisely, they highlight distinctions between truth and access, between metaphysical independence and epistemic availability, between inference and its materials, that my framework is well equipped to respect. My resource-sensitive approach does not seek to uproot classical conceptions of truth; rather, it supplements them with a more accurate account of the conditions under which truth enters epistemic life.

5. Conclusions

Let me conclude by gathering the main strands of the account I have developed. I set out to formulate a resource-sensitive theory of truth drawing on the principles of linear logic, not to replace traditional accounts, but to articulate a dimension of truth that classical theories leave unaddressed. Classical models tend to treat truth as a static, context-independent property of propositions. My aim has been to show that there is another way of understanding how truth enters epistemic practice, one that treats truth as emerging from derivations carried out under conditions of finitude, constraint and epistemic access.

In this setting, a proposition φ counts as true for an agent when and only when, φ can be derived from the informational resources presently available to that agent (data, concepts, evidential inputs, rules) through a derivation in which those resources are not merely listed but actively used and consumed. The shift gives epistemic structure to the process by which truth-claims are earned, rather

than treating truth as a label waiting for discovery. It shows, in a literal way, that what one can legitimately say depends on what one has.

A central distinction throughout has been that between local derivability and global derivability. The former concerns an agent's actual epistemic position; the latter concerns an idealized resource configuration from which the same proposition could, in principle, be derived. This division allows us to acknowledge the objectivity of truth while recognizing that the route to it may be blocked or opened depending on the materials one possesses. Far from endorsing relativism, my framework accounts for how agents with differing resources may converge on common truths as those resources grow or improve.

My framework also provides a way of approaching familiar epistemological puzzles, most notably, Gettier cases and limits on justification, by making explicit what is ordinarily left implicit: the epistemic status of the materials that support an inference. Resource certification, provenance markings and usage constraints give concrete form to the intuitive idea that not all inputs are equally truth-conferring. In this respect, my approach does not solve ancient problems by stipulation; it articulates the structure underlying successful and unsuccessful inferences alike.

My proposal sits naturally alongside practices in domains where reasoning is already constrained and accountable such as scientific explanation, legal argument, artificial intelligence and education. In these contexts, the ability to justify a conclusion depends on showing not merely *that* the conclusion follows, but *how* it follows and from what.

Stepping back, instead of picturing truth as a fixed point in logical space independent of our epistemic activities, we understand truth as something that becomes accessible through the disciplined use of available materials. This does not diminish truth's independence from us; it simply acknowledges that our access to it is mediated by finite and sometimes fragile resources. Truth, in this view, is not diminished, it is operationalized.

If there is a guiding theme here, it is that truth is not what stands beyond our epistemic reach, but what emerges when we reason well with what we have. A proposition becomes true-for-us not by virtue of our preferences, but by the successful construction of a derivation anchored in the resources that the world, our instruments and our practices make available. In this sense, truth remains what it always was, independent, stable, unforced, but our route to it is shaped by the conditions of inquiry. My framework, then, is not a redefinition of truth, but a clarification of how truth enters the space of reasons for finite, situated agents.

Authors' contributions. The Author performed: study concept and design, acquisition of data, analysis and interpretation of data, drafting of the manuscript, critical revision of the manuscript for important intellectual content, statistical analysis, obtained funding, administrative, technical and material support, study supervision.

Funding. This research did not receive any specific grant from funding agencies in the public, commercial or not-for-profit sectors.

Ethics approval and consent to participate. This research does not contain any studies with human participants or animals performed by the Author.

Consent for publication. The Author transfers all copyright ownership, in the event the work is published. The undersigned author warrants that the article is original, does not infringe on any copyright or other proprietary right of any third part, is not under consideration by another journal and has not been previously published.

Availability of data and materials. All data and materials generated or analyzed during this study are included in the manuscript. The Author had full access to all the data in the study and took responsibility for the integrity of the data and the accuracy of the data analysis.

Acknowledgements: none.

Competing interests. The Author does not have any known or potential conflict of interest including any financial, personal or other relationships with other people or organizations within three years of beginning the submitted work that could inappropriately influence or be perceived to influence their work.

Declaration of generative AI and AI-assisted technologies in the writing process. During the preparation of this work, the author used ChatGPT 4o to assist with data analysis and manuscript drafting and to improve spelling, grammar and general editing. After using this tool, the author reviewed and edited the content as needed, taking full responsibility for the content of the publication.

Appendix A. Formal Underpinnings of Resource-Sensitive Truth

The formal core of our theory is rooted in *multiplicative linear logic* (MLL), a substructural logic introduced by Jean-Yves Girard, in which structural rules such as weakening and contraction are restricted. This framework allows us to track how resources, understood as premises, observations, data, tools or assumptions, are consumed in the derivation of propositions. Unlike classical logic, where assumptions can be reused arbitrarily, linear logic treats assumptions as finite and consumable.

Let:

$\Gamma = \{a, b, c, \dots\}$ be a finite multiset of informational resources.

ϕ be a proposition or epistemic target.

$\otimes$ (tensor) denote **linear conjunction**, requiring all resources to be present and consumed.

$\multimap$ (linear implication or "lollipop") denote **resource-sensitive entailment**.

Sequent form

Truth is defined via the derivability of a sequent:

$$a \otimes b \otimes c \vdash \phi$$

This is interpreted as: the agent or system can derive ϕ *only if* it has simultaneous access to resources a, b, c and these are consumed in the process. The sequent calculus for linear logic provides inference rules that ensure resource accounting.

Linear implication and constructive truth

The implication $A \multimap B$ corresponds to a rule: "given A, we can produce B, consuming A in the process." Therefore, in our truth model:

$$\phi \text{ is true} \Leftrightarrow \exists \Gamma \text{ such that } \Gamma \vdash \phi$$

and

$$\forall \gamma \in \Gamma, \gamma \text{ is available and correctly used}$$

This aligns truth with constructive accessibility rather than metaphysical correspondence.

Reusability and modal extensions

To model reusable knowledge (e.g., background theorems or tautologies), we invoke Girard's **exponential modalities**:

$!A$ ("of course" A): A can be duplicated.

$?A$ ("why not" A): A can be discarded.

Resources marked with $!$ are considered universally reusable, e.g., mathematical axioms, shared ontologies.

Example:

$$!a, !b, c \vdash \phi$$

Here, a and b may be reused across derivations, while c must be uniquely present.

Two-tier truth

We define:

Local truth: ϕ is derivable from currently available resources.

Global truth: ϕ is derivable from all ideal, reusable or universally sharable resources.

This yields a semantic duality that allows comparison of agents with differing epistemic access.

Resource annotation schema

To improve granularity, resources can be annotated with source reliability, temporal status or epistemic tags:

$$\begin{array}{l} a^{\text{empirical}}, b^{\text{testimonial}}, c^{\text{theoretical}} \\ a^{\text{verified}}, b^{\text{tentative}}, c^{\text{obsolete}} \end{array}$$

These annotations allow the development of **typed resource logics** for advanced applications, including defeasible reasoning, belief revision or norm-sensitive inference.

Overall, this formalization supports our central claim: truth is not a binary valuation over propositions, but a context-sensitive outcome of derivation from available, typified and consumable informational resources under linear inference.

References

1. Alston, William P. *A Realist Conception of Truth*. Ithaca: Cornell University Press, 1996.
2. Andreoli, Jean-Marc. "Logic Programming with Focusing Proofs in Linear Logic." *Journal of Logic and Computation* 2, no. 3 (1992): 297–347.
3. Asher, Nicholas and Alex Lascarides. *Logics of Conversation*. Cambridge: Cambridge University Press, 2003.
4. Baelde, David. "Least and Greatest Fixed Points in Linear Logic." *ACM Transactions on Computational Logic* 13, no. 1 (2012): 1–44.
5. Barroso-Nascimento, Victor, Ekaterina Piotrovskaya, and Elaine Pimentel. 2025. "A Proof-Theoretic Approach to the Semantics of Classical Linear Logic." arXiv:2504.08349 [cs.LO], version 4, revised 7 November 2025. <https://arxiv.org/abs/2504.08349>.
6. Bellin, Gianluigi and Philip J. Scott. "On the π -Calculus and Linear Logic." *Theoretical Computer Science* 135, no. 1 (1994): 11–65.
7. Blanshard, Brand. *The Nature of Thought*, Vol. 2. London: Allen & Unwin, 1939.
8. Boghossian, Paul. *Content and Justification: Philosophical Papers*. Oxford: Oxford University Press, 2003.
9. Capps, John. 2020. "A Common-Sense Pragmatic Theory of Truth." *Philosophia* 48, no. 2: 463–481. <https://doi.org/10.1007/s11406-019-00099-z>.
10. Cartwright, Nancy. *How the Laws of Physics Lie*. Oxford: Oxford University Press, 1983.
11. Chang, Hasok. *Is Water H₂O? Evidence, Realism and Pluralism*. Dordrecht: Springer, 2012.
12. Girard, Jean-Yves. "Linear Logic." *Theoretical Computer Science* 50 (1987): 1–101.
13. Davidson, Donald. *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press, 1984.
14. Dummett, Michael. *The Seas of Language*. Oxford: Oxford University Press, 1993.
15. Feldman, Richard. "Internalism Defended." *Philosophical Studies* 40, no. 3 (1981): 255–72.
16. Franklin, Allan. *The Neglect of Experiment*. Cambridge: Cambridge University Press, 1994.
17. Gettier, Edmund L. "Is Justified True Belief Knowledge?" *Analysis* 23, no. 6 (1963): 121–23.
18. Girard, Jean-Yves. "Linear Logic." *Theoretical Computer Science* 50, no. 1 (1987): 1–101.
19. Girard, Jean-Yves. *The Blind Spot: Lectures on Logic*. Zurich: European Mathematical Society, 2011.
20. Haack, Susan. *Evidence and Inquiry: Towards Reconstruction in Epistemology*. Oxford: Blackwell, 1996.
21. Hacking, Ian. *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge: Cambridge University Press, 1983.
22. Hoffmann, Glen. 2010. "The Minimalist Theory of Truth: Challenges and Concerns." *Philosophy Compass* 5 (10): 938–949. <https://doi.org/10.1111/j.1747-9991.2010.00341.x>.
23. Jalali, Raheleh. 2021. "Proof Complexity of Substructural Logics." *Annals of Pure and Applied Logic* 172 (7): Article 102972. <https://doi.org/10.1016/j.apal.2021.102972>.
24. Kamp, Hans and Uwe Reyle. *From Discourse to Logic: Introduction to Model-Theoretic Semantics of Natural Language*. Dordrecht: Kluwer, 1993.
25. Kvanvig, Jonathan L. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press, 2003.
26. Martin-Löf, Per. "On the Meanings of the Logical Constants and the Justifications of the Logical Laws." *Nordic Journal of Philosophical Logic* 1, no. 1 (1996): 11–60.
27. Miller, Alexander. *Philosophy of Language*. London: Routledge, 2008.
28. Niiniluoto, Ilkka. *Critical Scientific Realism*. Oxford: Oxford University Press, 1999.
29. Peregrin, Jaroslav. *Inferentialism: Why Rules Matter*. London: Palgrave Macmillan, 2014.
30. Peirce, Charles S. "How to Make Our Ideas Clear." *Popular Science Monthly* 12 (1878): 286–302.
31. Pollock, John L. and Joseph Cruz. *Contemporary Theories of Knowledge*. Lanham: Rowman & Littlefield, 1999.
32. Prawitz, Dag. "Ideas and Results in Proof Theory." In *Proceedings of the Second Scandinavian Logic Symposium*, edited by J. E. Fenstad, 235–307. Amsterdam: North-Holland, 1971.
33. Putnam, Hilary. *Reason, Truth and History*. Cambridge: Cambridge University Press, 1981.
34. Quine, W. V. O. "Philosophy of Logic." *Foundations of Mathematics* 1 (1970): 1–9.
35. Ramsey, Frank P. "Facts and Propositions." *Proceedings of the Aristotelian Society* (1927): 153–70.
36. Restall, Greg. *An Introduction to Substructural Logics*. London: Routledge, 2000.

37. Rott, Hans. *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning*. Oxford: Oxford University Press, 2001.
38. Schroeder-Heister, Peter. "Rules of Definitional Reflection." In *Logic Colloquium '90*, edited by H.-D. Ebbinghaus et al., 221–42. Berlin: Springer, 1991.
39. Schroeder-Heister, Peter. "Validity Concepts in Proof-Theoretic Semantics." *Synthese* 148, no. 3 (2006): 525–71.
40. Sellars, Wilfrid. *Science, Perception and Reality*. London: Routledge & Kegan Paul, 1963.
41. Strawson, P. F. "Truth." *Proceedings of the Aristotelian Society* 24, suppl. vol. (1950): 129–56.
42. Szaif, Jan. 2018. "Plato and Aristotle on Truth and Falsehood." In *The Oxford Handbook of Truth*, edited by Michael Glanzberg, 9–49. Oxford: Oxford University Press.
43. Tennant, Neil. *The Taming of the True*. Oxford: Clarendon Press, 1997.
44. Troelstra, Anne S. *Lectures on Linear Logic*. Stanford: CSLI Publications, 1992.
45. van Benthem, Johan. *Logical Dynamics of Information and Interaction*. Cambridge: Cambridge University Press, 2011.
46. Walker, Ralph. 2018. "The Coherence Theory of Truth." In *The Oxford Handbook of Truth*, edited by Michael Glanzberg, 219–237. Oxford: Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199557929.013.8>.
47. Williamson, Timothy. *Knowledge and Its Limits*. Oxford: Oxford University Press, 2000.
48. James, William. *Pragmatism: A New Name for Some Old Ways of Thinking*. New York: Longmans, Green & Co., 1907.
49. Zagzebski, Linda. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press, 2000.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.