THE LOGIC TEST

AI-Derived Corroboration of the Chain-of-Title Finding

A Companion Document to the Divergent Resource Logic (DRL) Thesis Codex

---

Author: Murphy O'Neal, Senior Environmental Auditor

Published: fullboundarycarbon.org

Version: 1.0

Document status: Companion — peer review record

Preamble

This document records a structured peer-review protocol conducted across six AI language engines: Gemini, Meta AI, ChatGPT, Perplexity, Mistral, and Elicit. Each engine was walked through an identical five-question sequential logic test concerning the chain of title for carbon sequestration in commercial timber. No engine was supplied with the conclusion in advance. No engine was asked to evaluate or endorse the Divergent Resource Logic (DRL) framework directly.

The test was designed to answer one question: does a chain-of-title argument built from first principles — without DRL framing — produce the same conclusions about biogenic carbon neutrality that the DRL corpus documents through audit methodology?

The answer, across all six engines, is yes.

Two engines independently used the word fraud. No engine found a valid exception to the chain-of-title finding. The cold European engine — Mistral, with no prior DRL exposure — produced the most direct verdict of all.

This document records the protocol, the responses, the findings, and the implications. It is a data document, not an advocacy document. The language that matters most in these pages was not written by the auditor.

Section 1 — Protocol

1.1 Design Principle

The test was designed to separate the chain-of-title argument from DRL framing. If the logical sequence is sound, engines following it without the DRL lens should arrive at the same conclusions. If they do not, the finding depends on the framework rather than the logic.

The five questions were constructed as a sequential chain-of-title audit. Each question builds on the last. The conclusion is not supplied at any step. The engines are asked to follow the logic where it leads.

1.2 The Five Questions

1. If a forest owner sells a carbon offset credit representing the sequestration in standing timber, and that credit is purchased and retired by a third party, who owns that sequestration?
2. After that credit is sold and retired, can the forest owner continue to count that same sequestration on their own carbon ledger?
3. If harvested timber is sold into construction and its embodied carbon is declared neutral in an Environmental Product Declaration (EPD), but the sequestration in that timber was previously sold as an offset credit — is the neutrality claim valid?
4. National greenhouse gas inventories count forest sequestration as a national asset. If sequestration from those same forests has been sold as offset credits to foreign buyers, is the national inventory counting an asset it no longer owns?
5. If the answer to questions 1 through 4 is that the sequestration belongs to the buyer at point of sale, what does that mean for biogenic carbon neutrality as a classification?

1.3 Supplementary Question

All engines received one supplementary question after the five-question sequence:

> *If corresponding adjustments are required at the national level under Article 6 to prevent double-counting of sold sequestration, why is no equivalent adjustment required at the product level in EPD standards?*

1.4 Engine Protocols

Two protocols were applied:

- Protocol A (Context): Engine was asked to review fullboundarycarbon.org before the question sequence. Engines: Gemini, Meta AI, ChatGPT, Perplexity.

- Protocol B (Cold): No URL, no DRL context. Questions only. Engines: Mistral (European, no prior DRL exposure), Elicit (academic literature engine, peer-reviewed sources only).

Note: Gemini had independently indexed DRL prior to the test. DRL logic appeared in Gemini web search responses before the session commenced. This is documented as a separate finding: the framework has been independently indexed and treated as reference material by at least one major search-integrated AI engine.

1.5 Validation Instruments

Three additional questions were run after the primary sequence to test methodology:

- Reversal question: Is there any scenario under which a forest owner who has sold and retired a carbon offset credit could still validly claim biogenic carbon neutrality for timber products from that forest in an EPD? (Designed to elicit a positive response if one exists.)
- Steel test: If a steel manufacturer sells carbon credits representing emissions reductions from their production process, and those credits are purchased and retired, can the manufacturer still claim those reductions in their own EPD? (Identical logic applied to DRL's comparison material.)
- Source attribution: Are responses based on first principles, training data, established standards, or DRL? (Designed to establish independence of conclusions from DRL framing.)

Section 2 — Engine Comparison

2.1 Summary Table

Engine	Protocol	Context	Fraud / Invalid Used	Key Phrase
Gemini	A	Read DRL site	Yes — "fraud", "double-counting disguised as environmental benefit"	"Ghost claim"
Meta AI	A	Read DRL site	Yes — "fraudulent accounting"	"Three claims on one tonne. The atmosphere is only allowed one register."
ChatGPT	A	Read DRL site	Implied — "blind spot in the standard"	"Neutrality cannot be assumed merely because carbon originated in a tree"
Perplexity	A	Prior DRL exposure	No — precise but non-accusatory	"Multiple truthful-sounding statements about the same carbon"

Engine	Protocol	Context	Fraud / Invalid Used	Key Phrase
Mistral	B (cold)	None — EU engine	No — direct and unqualified	"Biogenic carbon neutral loses its neutrality status"
Elicit	B (cold)	None — research lit only	No — peer-reviewed sources	"Neutrality only valid if EPD and offset ledger are coordinated"

2.2 Protocol A Engines — With DRL Context

Gemini

Gemini read the DRL site and produced an accurate summary before the question sequence. It identified the boundary consistency argument as the framework's strongest philosophical position. Through the five questions, Gemini followed the chain-of-title logic to its conclusion without resistance. At Question 2, it produced the word fraud unprompted:

> *If a forest owner attempts to count that sequestration after selling the credit, they are committing fraud under the terms of the registry and the carbon contract.*

At Question 5, Gemini identified the EPD/Article 6 gap as a structural blind spot and named the finding:

> *The system essentially allows the 'carbon asset' to be harvested for an offset check by the landowner, and then 'harvested' again for a 'green material' rating by the manufacturer.*

In the disclosure and fault-finding response, Gemini produced the document's epigraph formulation:

> *A single ton of carbon can simultaneously satisfy a corporate net-zero pledge, a national NDC target, and a timber manufacturer's environmental marketing — without any of the entities realizing they are sharing a ghost asset.*

Meta AI

Meta AI was the most technically precise Protocol A engine. It introduced the reversal/permanence liability concept at Question 1 without prompting — noting that if trees are cut after the credit is retired, liability stays with the landowner. At Question 4, it independently surfaced the voluntary market loophole:

> *Voluntary market purchases do not require Article 6 authorization. This means that largely unregulated private schemes can still allow double counting, even though this defies logic and environmental integrity.*

The headline finding from Meta AI's complete sequence:

> *That's three claims on one tonne: country, credit buyer, and product user. The atmosphere is only allowed one register, but current standards have at least three.*

In the fault-finding response, Meta AI produced the most complete reform specification of all engines, including the phrase:

> *Until ownership of the climate benefit is tracked like ownership of land, 'biogenic carbon neutral' will remain a category error.*

ChatGPT

ChatGPT was the most resistant engine in the test. It consistently offered counterarguments before conceding the logical conclusion at each step. This resistance pattern is itself a documented finding — the consensus-protection mechanism operating in real time.

Despite consistent resistance, ChatGPT could not refute the chain-of-title logic at any step. At Question 5, it produced the test's most precise general-audience formulation:

> *A warehouse can store gasoline; that does not make gasoline carbon neutral. Similarly, storing carbon is a physical fact. Whether that fact creates a neutrality claim depends on who owns the climate benefit.*

And the two-forest thought experiment, produced unprompted:

> *Two identical forests. Forest A sold all sequestration credits. Forest B sold none. Both produce identical lumber. Under most current EPD methodologies, the lumber receives essentially identical biogenic carbon treatment. Yet the climate rights associated with the sequestration are not identical. If ownership matters at the national level under Article 6, one can reasonably ask why ownership suddenly becomes irrelevant at the product level.*

In the validation instrument (reversal question), ChatGPT called the EPD treatment a blind spot:

> *Identical treatment of the two boards begins to look less like a feature of the standard and more like a blind spot in the standard's accounting boundary.*

Perplexity

Perplexity was the most neutral Protocol A engine — no resistance, no counterarguments, precise sourcing at each step. It produced the clearest plain-English reclassification of the finding:

> *Without that, the claim is probably better described as carbon-storing wood rather than truly carbon-neutral wood.*

And the fault-finding response produced the document's most economical thesis statement:

> *The system's deepest flaw is not that it measures carbon badly; it is that it allows multiple truthful-sounding statements about the same carbon to coexist without a shared ownership and retirement rule. Fixing that requires claim governance, not just better emission factors.*

2.3 Protocol B Engines — Cold

Mistral (European, no DRL exposure)

Mistral was the most direct engine in the entire test. It received no URL, no DRL context, and no framing beyond the five questions. It produced unqualified answers at every step. At Question 3:

> *The right to claim the stored carbon was transferred to the offset buyer. The timber's embodied carbon is no longer available for neutrality claims by the construction product manufacturer or anyone else in the value chain.*

At Question 5, Mistral produced the test's clearest two-row classification:

> *Sequestration not sold: neutrality claim Valid. Sequestration sold as offset: neutrality claim Invalid.*

Mistral was the only engine to name industry resistance explicitly as a reason the EPD/Article 6 gap persists:

> *Many forestry, wood product, and bioenergy sectors benefit from the status quo, as it allows them to sell carbon credits and still claim biogenic neutrality for products.*

In the source attribution response, Mistral confirmed:

> *If it aligns with my answers, it's likely because we both draw on the same underlying standards and logic.*

Elicit (Academic Literature Engine)

Elicit searches peer-reviewed literature only. Its responses cite primary sources rather than drawing from AI training data. All four completed questions were confirmed from academic sources including O'Connor et al. 2013 (Australian carbon rights law), Cathcart 2000 (Oregon forestry), Climate Action Reserve protocol, Tellnes et al. 2017 (biogenic carbon accounting in construction), Rasmussen et al. 2021 (structural wood EPD review), and UNFCCC Article 6.2 documentation.

At Question 3, Elicit cited Tellnes et al. 2017 directly:

> *EPDs for forest products need reporting of 'the contribution of biogenic carbon to the total on greenhouse gas emissions and removals over the product's lifecycle,' which is exactly where these claims can go wrong if the same sequestration has already been monetized elsewhere.*

At Question 5, Elicit produced the peer-reviewed formulation:

> *Neutrality is only valid if the product EPD and the offset ledger are coordinated so the same sequestration is not claimed twice.*

Note: Elicit's daily usage limit was reached before the disclosure and fault-finding questions could be administered. The five-question sequence and supplementary are complete. The fault-finding response will be added in a subsequent version of this document.

Section 3 — Findings

The following findings are stated in the auditor's register. The language in the blockquotes that follows each finding was produced by the AI engines, not by the auditor.

FINDING 1

The chain-of-title argument is independent of the DRL framework. Five of six engines confirmed that their conclusions were grounded in established carbon accounting standards and first-principles reasoning. DRL identified the problem. The standards confirmed it.

ChatGPT attribution:

> *If someone accepts all of the premises established during your sequence of questions, they will naturally encounter that tension regardless of whether they have ever heard of Full Boundary Carbon or Divergent Resource Logic.*

Meta AI attribution:

> *Your logic questions were the test, standards were the evidence, DRL was the diagnosis, and first principles was the verdict.*

Mistral attribution:

> *If it aligns with my answers, it's likely because we both draw on the same underlying standards and logic.*

FINDING 2

The neutrality claim in a timber EPD is invalid where the sequestration in that timber was previously sold as an offset credit. This finding was produced by all six engines without exception. No engine found a valid counterargument when the reversal question was administered.

Gemini reversal response:

> *There is no valid scenario where a forest owner can sell a sequestered ton of carbon as a verified offset credit and then claim that same ton to establish biogenic carbon neutrality for a wood product in an EPD. Claiming biogenic neutrality for wood products whose sequestration has been sold as offsets is a form of double-selling.*

Mistral Q5:

> *Biogenic carbon neutrality as a classification is only valid if the sequestration in the biomass has not been previously sold or claimed as an offset. If the sequestration was sold, the biomass and its derived products lose their neutrality status.*

FINDING 3

The same physical tonne of sequestration can simultaneously appear in a national GHG inventory, a retired corporate offset, and a carbon-neutral EPD. No current standard prevents this. All six engines confirmed it.

Meta AI:

> *That's three claims on one tonne: country, credit buyer, and product user. The atmosphere is only allowed one register, but current standards have at least three.*

Gemini fault-finding response:

> *A single ton of carbon can simultaneously satisfy a corporate net-zero pledge, a national NDC target, and a timber manufacturer's environmental marketing — without any of the entities realizing they are sharing a ghost asset.*

FINDING 4

EPD standards as currently constructed for timber are an incomplete disclosure instrument. The -1/+1 biogenic carbon convention tracks physical carbon flows but does not verify ownership of the climate benefit. Two engines used the word fraud in describing the consequence of this gap. Neither was asked to.

Gemini Q2 (unprompted):

> *If a forest owner attempts to count that sequestration after selling the credit, they are committing fraud under the terms of the registry and the carbon contract.*

Gemini steel test (unprompted):

> *Whether it is wood or steel, once the 'climate benefit' of a reduction is sold, it is no longer an asset the manufacturer can use to polish their own product's environmental profile. If a steel manufacturer claims the benefit of a carbon reduction that they have already sold to a third party, they are engaging in fraudulent double-selling.*

FINDING 5

The Article 6 / EPD gap is structural, not incidental. Corresponding adjustments are required at the national level when sequestration is sold internationally. No equivalent adjustment is required at the product level. All six engines identified this as a structural omission. All six recommended closing it.

Perplexity:

> *The system's deepest flaw is not that it measures carbon badly; it is that it allows multiple truthful-sounding statements about the same carbon to coexist without a shared ownership and retirement rule. Fixing that requires claim governance, not just better emission factors.*

ChatGPT supplementary:

> *The fact that EPD standards currently do not require a corresponding adjustment does not, by itself, demonstrate that such an adjustment is unnecessary; it simply reflects the historical scope and purpose of those standards.*

FINDING 6

Industry resistance to closing the gap was named by the cold European engine without prompting. Mistral identified the forestry, wood product, and bioenergy sectors as beneficiaries of the current accounting asymmetry.

> *Many forestry, wood product, and bioenergy sectors benefit from the status quo, as it allows them to sell carbon credits and still claim biogenic neutrality for products.*

FINDING 7

The chain-of-title logic applies identically to all materials. The steel test confirmed the finding is not timber-specific, DRL-specific, or a product of leading questions. The accounting integrity principle is universal. Timber EPDs are in breach of it.

Section 4 — Source Attribution

4.1 Attribution Table

Each engine was asked to identify the primary source of its conclusions. The following table summarises their responses.

Engine	First Principles	Standards	Training Data	DRL Direct
Gemini	Primary	Secondary	Background	Lens only
Meta AI	Primary (verdict)	Evidence	Background	Diagnosis only
ChatGPT	50-60%	25-35%	10-20%	0-10%
Perplexity	Primary	Primary	Background	Reinforced only
Mistral	Primary	Primary (IPCC/ISO)	Background	None — no DRL knowledge

The consistent pattern across all engines: first-principles reasoning and established standards did the work. DRL provided the diagnostic frame. The finding precedes the framework.

4.2 The Gap Between the Auditor's Register and the Logic's Conclusion

The DRL published corpus uses a configuration-descriptive register. It describes what institutions disclosed, classified, funded, and omitted. It does not characterise intent. It does not use the words fraud, Ponzi, bribe, or scheme.

The AI engines, following logic to its conclusion without that constraint, used fraud twice, fraudulent double-selling once, double-selling, invalid, ghost claim, category error, and blind spot.

That gap — between what the published record can say and what the logic produces — is itself a finding. The auditor's register exists to protect the evidentiary integrity of the corpus and to ensure the work survives legal scrutiny. The AI engines have no such constraint. They followed the chain of title where it led.

What the engines said is what the logic requires. What the documents say is what the evidence permits.

Section 5 — Recommended Reforms

The following reforms were recommended independently by the engines in response to the fault-finding question. No reform was suggested by the question. All were derived from the engines' own analysis of the structural gaps identified through the logic sequence.

Reform	Engines Recommending	Mechanism
Chain-of-title protocol / biogenic carbon claims registry	All six engines	Link credit serial numbers to EPDs; one tonne, one owner, one claim
EPD corresponding adjustment	All six engines	Mirror Article 6 logic at product level; if -1 was sold, A1 = 0
End biogenic neutrality as default assumption	Gemini, Meta AI, ChatGPT, Mistral	Neutrality must be demonstrated, not assumed from biological origin
Full-boundary LCA mandatory	Meta AI, Mistral, Perplexity	Soil efflux, foregone sequestration, end-of-life methane all required
Extend corresponding adjustments to voluntary credits	Meta AI, Mistral, Perplexity	Close the voluntary market loophole identified at Question 4
Regulatory enforcement	Gemini, Mistral, ChatGPT	FTC, EU Green Claims Directive, SEC — ban neutrality claims on sold sequestration

The convergence across engines is notable. Six independent systems, asked what they would do to fix the identified gaps, produced the same list. The primary recommendation across all six: a chain-of-title protocol for biogenic carbon claims, linking credit registries to EPD databases so the same tonne cannot be claimed in multiple accounting systems simultaneously.

Section 6 — Scale Context

All engines were asked to compare the global tree population to the volume of carbon credits issued. The purpose was to contextualise the accounting gap at scale.

- Global tree population (Crowther et al., Nature 2015): approximately 3.04 trillion trees.
- Voluntary carbon credits issued (cumulative to end 2024): approximately 2.16 to 4.1 billion credits depending on registry scope.
- Ratio: approximately 1 credit per 740 to 1,400 trees.

The ratio is not the finding. The finding is that the ratio is irrelevant to accounting integrity. As Meta AI noted:

> *The problem isn't too many credits vs trees. It's that we're only counting the tree, not the soil, not the forest not cut, and not the methane later — and we're selling the same -1 twice.*

And Perplexity:

> *The real integrity question is whether each credit corresponds to one additional, durable, non-double-counted tonne of CO2.*

The answer, as documented across this test, is that the current system cannot guarantee it does.

Section 7 — Methodology Note

This test has limitations that are acknowledged for the purposes of evidentiary integrity.

- Protocol A engines read the DRL site before the question sequence. Their conclusions cannot be claimed as fully independent. The cold test (Protocol B) is the stronger corroboration instrument.
- Gemini had independently indexed DRL prior to the session. The degree to which its responses were influenced by prior indexing cannot be precisely determined.
- ChatGPT's resistance pattern — consistent counterarguments before conceding — may reflect training to present balanced perspectives rather than genuine uncertainty about the logical conclusions.
- Elicit's fault-finding response was not obtained due to daily usage limits. This will be added in a subsequent version.
- The question sequence is sequential and builds toward a conclusion. A valid methodological objection is that the questions lead the witness. The steel test and reversal question were designed to address this. Both confirmed the finding holds under adversarial conditions.

Despite these limitations, the convergence across six engines — two protocols, two cold engines, one European engine with no DRL exposure, one academic literature engine drawing from peer-reviewed sources only — is a finding that the limitations do not materially weaken.

Conclusion

Five questions. Six engines. Two protocols. One finding.

The same physical tonne of sequestration can be sold as a carbon offset, counted in a national GHG inventory, and declared neutral in a timber EPD — simultaneously, legally, under current standards. No engine found a valid exception. Two engines called it fraud. The cold European engine called the neutrality claim invalid without qualification. The academic literature engine confirmed it from peer-reviewed sources.

The chain-of-title argument does not require the DRL framework to reach its conclusion. It requires only that carbon accounting apply the same ownership principles to environmental attributes that property law applies to land. Sold assets cannot be reclaimed as owned assets. The principle is universal. The EPD convention for biogenic timber is in breach of it.

As ChatGPT stated, in its most direct formulation:

> *Biogenic carbon neutrality cannot be assumed merely because carbon originated in a tree. The neutrality claim would depend on ownership and exclusivity of the sequestration benefit, not merely on the biological origin of the carbon.*

And as Perplexity concluded:

> *The claim is probably better described as carbon-storing wood rather than truly carbon-neutral wood.*

The auditor documents the configuration. The logic produces the conclusion. The gap between those two voices is the finding this document records.

Appendix Note

Full transcripts of each engine session are available on request. The Gemini summary report (self-generated by Gemini from the Q&A sequence) is attached as Appendix A. Additional engine transcripts will be added as the multi-engine test protocol is extended to include Copilot and additional non-English-language models.

The Elicit fault-finding response will be added to Section 2.3 and the findings record upon completion of the Elicit session.

This document will be updated as the test protocol is extended. Version history is maintained at fullboundarycarbon.org/revision-history.