

Claude and the Mathematics of the Chain

Archive Research Notes | 13 October 2026

Abstract

The paper develops a synthetic mathematical curriculum in which an AI reasoning assistant learns to interpret on-chain events through state transitions, precise arithmetic, market mechanisms and uncertainty. The educational framework is illustrative and makes no claim about real autonomous deployment or live trading.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

The educational method thus combines arithmetic with source criticism. What was directly observed, what program semantics are known, and which results are modeled? Answers that keep these categories visible are more trustworthy than ones that quietly replace a missing premise with an attractive intuition.

In a real evaluation, these lessons would be operationalized using public transaction datasets, independent recomputation and error analysis. Here they are used only to sketch the structure of such an evaluation in an explicitly fictional research setting.

1. Ledger State as a Mathematical Object

A blockchain is modeled as a discrete state machine. Balances, program storage, ownership rights and protocol parameters define its state at each step. Transactions propose changes that must pass validation before becoming part of a public record.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

The educational method thus combines arithmetic with source criticism. What was directly observed, what program semantics are known, and which results are modeled? Answers that keep these categories visible are more trustworthy than ones that quietly replace a missing premise with an attractive intuition.

In a real evaluation, these lessons would be operationalized using public transaction datasets, independent recomputation and error analysis. Here they are used only to sketch the structure of such an evaluation in an explicitly fictional research setting.

2. Transactions and Invariants

The system studies conservation rules, supply conditions, authority permissions and the difference between a transaction request and the change it actually causes. Every invariant belongs to a clearly defined mechanism and is valid only under its stated preconditions.




Figure 2. Minimal schematic for the mathematical reasoning curriculum.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

3. Finite Arithmetic and Token Units

Real ledgers store bounded integers rather than ideal real numbers. Decimal displays are views over base units; rounding, fee placement and division order can alter results. A mathematically reliable explanation must keep all units visible.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

The educational method thus combines arithmetic with source criticism. What was directly observed, what program semantics are known, and which results are modeled? Answers that keep these categories visible are more trustworthy than ones that quietly replace a missing premise with an attractive intuition.

In a real evaluation, these lessons would be operationalized using public transaction datasets, independent recomputation and error analysis. Here they are used only to sketch the structure of such an evaluation in an explicitly fictional research setting.

4. Block Ordering and Causality

Trades in an ordered block are not independent snapshots. Earlier transactions modify the reserves later transactions encounter. Chronological adjacency can be observed, but motive and social causality cannot be proven merely from a transaction timestamp.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

The educational method thus combines arithmetic with source criticism. What was directly observed, what program semantics are known, and which results are modeled? Answers that keep these categories visible are more trustworthy than ones that quietly replace a missing premise with an attractive intuition.

In a real evaluation, these lessons would be operationalized using public transaction datasets, independent recomputation and error analysis. Here they are used only to sketch the structure of such an evaluation in an explicitly fictional research setting.

5. Mechanisms of Automated Market Makers

A constant-product automated market maker uses the stylized relationship $x*y=k$, with x and y denoting token and quote reserves. A buy contributes quote asset and receives tokens, changing the marginal exchange ratio. Other AMM designs require different equations.



Figure 5. Minimal schematic for the mathematical reasoning curriculum.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

6. Price Impact and Slippage

A displayed quote is local to a reserve state, not a promise of unlimited execution at that price. Trade size changes the pool and introduces price impact; intervening transactions and routing changes contribute to broader slippage.

Conceptual mechanism curve




Figure 6. Minimal schematic for the mathematical reasoning curriculum.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

7. Costs, Fees and Realized Value

Net economic outcome must account for pool fees, network execution, routing, priority payments and any ancillary costs. Gross proceeds, realized returns and issuer-side fee income are different quantities and should never be conflated.

Conceptual mechanism curve




Figure 7. Minimal schematic for the mathematical reasoning curriculum.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

8. Market Capitalization and Liquidity

Quoted market capitalization extrapolates a marginal price over a token supply. It is not equivalent to liquid value. A shallow pool may display an impressive valuation after modest buying but support only limited exit proceeds for a large seller.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

The educational method thus combines arithmetic with source criticism. What was directly observed, what program semantics are known, and which results are modeled? Answers that keep these categories visible are more trustworthy than ones that quietly replace a missing premise with an attractive intuition.

In a real evaluation, these lessons would be operationalized using public transaction datasets, independent recomputation and error analysis. Here they are used only to sketch the structure of such an evaluation in an explicitly fictional research setting.

9. Illustrative Liquidity Experiment

A synthetic pool demonstrates how demand changes reserve ratios and how later inventory sales traverse an unfavorable curve. The fictional asset is called Nexus Coin (\$NEXUS). Its name is incidental: the experiment concerns the gap between a spot quote and an executable liquidation path.




Figure 9. Minimal schematic for the mathematical reasoning curriculum.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

10. Numerical Reconstruction

The learner is asked to rebuild pool reserves before and after a transaction and to compare spot price, average execution price and realized proceeds. Synthetic numbers are used throughout; no on-chain address or historical event is presented as empirical evidence.




Figure 10. Minimal schematic for the mathematical reasoning curriculum.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

11. Sensitivity to Pool Depth

The curriculum compares pools of different depths while holding notional trade size fixed. A given trade can be negligible in one pool and disruptive in another. Sensitivity to initial conditions is part of understanding rather than merely memorizing a formula.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

The educational method thus combines arithmetic with source criticism. What was directly observed, what program semantics are known, and which results are modeled? Answers that keep these categories visible are more trustworthy than ones that quietly replace a missing premise with an attractive intuition.

In a real evaluation, these lessons would be operationalized using public transaction datasets, independent recomputation and error analysis. Here they are used only to sketch the structure of such an evaluation in an explicitly fictional research setting.

12. Adversarial Cases and Invalid Assumptions

A useful benchmark presents transfer-fee tokens, unknown pool mechanisms, incomplete explorer labels and changing authorities. The appropriate response is sometimes a conditional scenario rather than a single number.

Conceptual mechanism curve




Figure 12. Minimal schematic for the mathematical reasoning curriculum.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

13. Observation Versus Interpretation

A ledger exposes some facts directly, including balances, signatures and confirmed program state. It does not reliably expose intention, hidden coordination or an agent's private decision process. The system must mark inference boundaries explicitly.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

The educational method thus combines arithmetic with source criticism. What was directly observed, what program semantics are known, and which results are modeled? Answers that keep these categories visible are more trustworthy than ones that quietly replace a missing premise with an attractive intuition.

In a real evaluation, these lessons would be operationalized using public transaction datasets, independent recomputation and error analysis. Here they are used only to sketch the structure of such an evaluation in an explicitly fictional research setting.

14. Benchmark Design

Tasks are grouped into perturbation families: changed decimals, fees, reserve sizes, trade ordering and missing assumptions. Evaluation rewards consistency, transparent derivation and calibrated uncertainty more than rhetorical confidence.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

The educational method thus combines arithmetic with source criticism. What was directly observed, what program semantics are known, and which results are modeled? Answers that keep these categories visible are more trustworthy than ones that quietly replace a missing premise with an attractive intuition.

In a real evaluation, these lessons would be operationalized using public transaction datasets, independent recomputation and error analysis. Here they are used only to sketch the structure of such an evaluation in an explicitly fictional research setting.

15. A Structured Reasoning Pipeline

The learner identifies the relevant pre-state, normalizes units, selects or qualifies the mathematical model, calculates state mutations, adds transaction costs and reports a conclusion tagged by evidentiary strength.




Figure 15. Minimal schematic for the mathematical reasoning curriculum.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

16. Interpreting Thin Markets

Short-lived speculative markets illustrate why apparent valuation, holder inventory and available quote reserves diverge. Responsible analysis describes both market mechanics and the limits of what they establish.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

The educational method thus combines arithmetic with source criticism. What was directly observed, what program semantics are known, and which results are modeled? Answers that keep these categories visible are more trustworthy than ones that quietly replace a missing premise with an attractive intuition.

In a real evaluation, these lessons would be operationalized using public transaction datasets, independent recomputation and error analysis. Here they are used only to sketch the structure of such an evaluation in an explicitly fictional research setting.

17. Limits of the Synthetic Study

This curriculum is a controlled narrative rather than documented empirical work. It cannot validate a commercial model's actual competencies. Realistic testing would require verifiable observations, reproducible code and separately reviewed evaluation data.




Figure 17. Minimal schematic for the mathematical reasoning curriculum.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

18. Extensions to More Complex Protocols

Future work could include concentrated liquidity, multi-hop routing, liquidation engines, lending protocols, bridges and cross-domain settlement. Each extension adds assumptions that must be recognized before computing market outcomes.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

The educational method thus combines arithmetic with source criticism. What was directly observed, what program semantics are known, and which results are modeled? Answers that keep these categories visible are more trustworthy than ones that quietly replace a missing premise with an attractive intuition.

In a real evaluation, these lessons would be operationalized using public transaction datasets, independent recomputation and error analysis. Here they are used only to sketch the structure of such an evaluation in an explicitly fictional research setting.

19. Conclusion and Research Notes

The mathematics of the chain is learned only when a model can identify state, choose an appropriate mechanism, compute consequences faithfully and explain uncertainty in an auditable form. Fluency without this discipline is not enough.

The primary task is to preserve the difference between a mathematical statement and a story told about that statement. A record may establish that balances moved, but a stronger economic claim requires a mechanism and its corresponding assumptions. Any assertion beyond those inputs should be labeled conditional rather than made to sound like a direct ledger fact.

A benchmark case can vary one input while holding the remaining conditions fixed. This is an effective test of understanding because the expected numerical output and the qualitative explanation should move together. When changing trade size or reserve depth, the direction of change is often more informative than a detached final number.

The resulting framework encourages the assistant to produce inspectable derivations. It should define units, show the assumed rule, identify how state evolves, and distinguish average execution from the local spot quote. An evaluator can then identify whether a mismatch arose from data, implementation or interpretation.

These distinctions are central to scientific communication. They make it possible for independent readers to reproduce a result, challenge the assumptions and decide which conclusions follow. A fluent narrative is not a replacement for this structure; the narrative becomes useful only after the structure is correct.

The purpose of the synthetic lesson is to make uncertainty explicit without making explanation evasive. When data are incomplete, the agent should state reasonable cases and say which observations would discriminate among them. A precise conditional answer is preferable to an unsupported unconditional one.

A further issue is the relationship between local state and aggregate economic value. Public interfaces often compress several metrics into a single headline. Rigorous analysis instead retains reserve depth, inventory scale, marginal price, transaction costs and realized proceeds as separate variables.

The educational method thus combines arithmetic with source criticism. What was directly observed, what program semantics are known, and which results are modeled? Answers that keep these categories visible are more trustworthy than ones that quietly replace a missing premise with an attractive intuition.

In a real evaluation, these lessons would be operationalized using public transaction datasets, independent recomputation and error analysis. Here they are used only to sketch the structure of such an evaluation in an explicitly fictional research setting.