

Generative AI and the Allocation of Epistemic Responsibility in Undergraduate Finance Education

Zihao Ning^{1, a}, Jun Cui^{2, b, *}

¹ School of Modern Finance, Jiaxing Nanhu University, Jiaxing, 314001 Zhejiang, China

² SolBridge International School of Business, Woosong University, Daejeon, 34613, Republic of Korea

^a nzh950721@jxnhu.edu.cn, ^b jcui228@student.solbridge.ac.kr

Abstract: Generative artificial intelligence can produce credible financial analyses while leaving students unable to justify the reasoning those analyses contain. This conceptual paper explains when teacher competence in integrating generative AI can support independent financial understanding. Through theory adaptation, it connects technological pedagogical content knowledge with cognitive engagement and cognitive offloading, using epistemic cognition to specify what learners remain responsible for knowing. The central construct is the instructional allocation of epistemic responsibility: the distribution of obligations to select assumptions, justify inferences, verify evidence, and revise conclusions among students, teachers, and AI-supported activities. The paper argues that competence contributes to learning when it creates feasible opportunities for students to exercise these obligations. Assistance can release capacity for conceptual work, but it can also substitute for the very reasoning students need to acquire. Five propositions explain these competing pathways and their dependence on prior knowledge, access to verification, and assessment incentives. A four-configuration typology and finance-specific thought experiments distinguish productive delegation from dependence, including cases in which accurate AI output still undermines opportunities to learn. The contribution is a conditional account of how professional knowledge becomes educationally consequential through task design. It also establishes why assisted product quality, independent understanding, and justified judgment under uncertainty must remain distinct educational outcomes. This paper develops theoretical arguments without collecting or analyzing primary data.

Keywords: Generative artificial intelligence, Finance education, Teacher competence, Epistemic responsibility, Cognitive engagement, Cognitive offloading.

1. Introduction

A student can submit a coherent company valuation without being able to explain why its cash flows and discount rate belong together. Generative artificial intelligence makes this possibility educationally consequential because it can supply both the calculation and the accompanying rationale. The resulting document may exhibit features conventionally associated with understanding even when the learner has not constructed the relevant relationships. For finance educators, the problem extends beyond detecting incorrect answers. It concerns whether the work required to produce a defensible analysis remains work that students undertake.

Recent educational experiments give the distinction between learning and performance a concrete technological setting. Bastani et al. [1] report that unrestricted AI assistance improved mathematics practice performance but reduced subsequent unassisted performance relative to a control condition; instructional safeguards mitigated the harm without establishing a corresponding positive examination effect. Conversely, Kestin et al. [2] report stronger immediate learning outcomes with a deliberately designed AI tutor than with an active classroom comparison in college physics. These studies concern different interventions and settings. They neither establish a general effect of AI nor test a mechanism in undergraduate finance, but they ask instructional configuration a theoretically significant question.

Teacher knowledge offers an important starting point. Technological pedagogical content knowledge, or TPACK, locates effective technology integration in the relationship among disciplinary understanding, pedagogy, and technological knowledge [3]. Intelligent-TPACK extends

attention to teachers' professional knowledge for ethically integrating AI-based tools [4]. These frameworks help specify the knowledge needed for appropriate instructional decisions. Translating them into an explanation of student learning, however, requires an additional account of how those decisions change what students actually do. Knowledge of an AI system does not establish that learners practice the reasoning a course aims to develop.

This paper asks: How does the instructional allocation of epistemic responsibility explain when teachers' generative AI integration competence supports independent financial understanding? Epistemic responsibility refers here to an assigned obligation to provide and evaluate reasons for a knowledge claim. Its instructional allocation concerns which reasoning decisions students must undertake and defend, which resources support them, and which decisions may be delegated. A requirement to submit an investment recommendation allocates little responsibility if an AI system can supply the recommendation and every justification needed for assessment. Requiring students to defend the recommendation after a consequential assumption change allocates responsibility differently.

We argue that competence becomes educationally consequential through this allocation. The same technology can reduce routine burdens while preserving conceptual decisions or completing those decisions before learners confront them. The first arrangement can create opportunities for constructive engagement; the second can produce successful submissions alongside dependence on generated reasoning. Which pathway becomes more plausible depends on whether the required work is feasible for the learner, whether claims can be checked, and whether assessment

rewards the exercise of judgment. The argument therefore does not presume a uniformly positive relationship between competence, AI use, and learning.

The paper makes a focused contribution to theories of teaching with AI. It introduces an intermediate instructional construct between teacher capability and student activity, explains how offloading can support or displace disciplinary reasoning, and specifies why apparently demanding AI activities can fail when responsibility is assigned without resources for exercising it. Finance supplies a revealing setting because mathematical consistency, uncertain assumptions, and evidential justification must be considered together. The resulting model is intended for undergraduate courses that teach financial reasoning through AI-assisted activities; it is not a theory of all AI adoption or of professional investment performance.

2. Conceptual Approach and Scope

This conceptual theory adaptation article uses an explicit, problem-centered research design. Each theoretical resource serves a defined explanatory function, and the steps of inference are made visible. TPACK supplies the focal account of professional knowledge. ICAP supplies distinctions among forms of learning activity, while cognitive offloading explains how external assistance can alter the work a learner performs. Epistemic cognition identifies the object of that work: claims, reasons, sources, and standards of justification. Cognitive load and constructive alignment provide boundary arguments about feasibility and incentives. The article uses a selective engagement with these literatures rather than a systematic review or a synthesis of effect sizes.

Theory selection follows explanatory function. A technology-acceptance account would help explain willingness to use a system, but acceptance is not the outcome in question. A broad account of motivation would explain effort but would not, by itself, distinguish effort spent improving an AI-produced document from effort spent evaluating its financial reasoning. TPACK alone identifies relevant knowledge without specifying this allocation of work. ICAP alone distinguishes engagement modes without locating the instructional decisions that make them possible. Bringing these perspectives into dialogue allows the model to address a narrower problem than a general inventory of benefits and risks.

The analysis proceeds by distinguishing teacher capacity, designed obligations, enacted activity, and learning outcomes. We then derive alternative consequences of delegating work to AI and examine limiting cases in which one proposed condition is absent. The finance scenarios are analytical thought experiments constructed to expose these differences; they are not classroom observations, interviews, or intervention evidence. Published empirical studies establish the plausibility of the motivating problem and inform theoretical premises. They do not validate the propositions developed here.

The focal unit is the instructional episode: a bounded task, its available resources, and the conditions under which students must account for their decisions. Teacher competence is a capacity that can shape multiple episodes. Responsibility is allocated in the episode's design, engagement is enacted by individual learners or collaborating students, and durable understanding develops across episodes. Keeping these levels distinct prevents the model from treating a teacher characteristic as if it were a student cognitive state.

A well-designed task can be bypassed, and students can undertake productive reasoning in a poorly designed task.

Three scope assumptions delimit the argument. The course must have an identifiable financial reasoning objective; students must encounter the AI-supported activity; and teachers must have some discretion over tasks, guidance, or assessment. The model does not apply directly when AI is used only to prepare administrative materials or when a platform fixes every instructional decision. It also assumes that independent understanding remains an educational aim. Where the sole objective is assisted production, the normative case for preserving particular forms of independent reasoning requires separate justification.

3. Theoretical Foundations and Construct Distinctions

3.1. Teacher Competence as a Capacity for Instructional Judgment

Generative AI pedagogical integration competence is the capacity to select, evaluate, and organize AI-supported activity in relation to disciplinary learning objectives and learners' needs. Mishra and Koehler [3] establish why technological knowledge should be considered in relation to pedagogy and content. Celik [4] extends this professional-knowledge problem to AI and ethical evaluation. These sources support a situated conception of competence; neither supplies a validated finance-specific construct for the model proposed here.

We specify three areas of judgment. Instructional alignment concerns whether assistance serves the intended financial concept. Disciplinary evaluation concerns whether generated claims, assumptions, and sources are suitable for the task. Learning orchestration concerns the sequencing of assistance, student decisions, and feedback. These are conceptual domains, not empirically established dimensions of a scale. The definition excludes usage frequency, self-confidence, and favorable student outcomes. A teacher can be competent despite limited opportunity to use AI, while a frequent user can organize activities that leave little financial reasoning to students.

Competence is also distinct from its enactment. A teacher may recognize that students should compare competing assumptions but lack the time or institutional discretion to change a standardized assessment. Another may adopt a well-designed shared activity without possessing all the knowledge needed to create it independently. Thus, competent teachers are not the only possible source of productive task design, and competence cannot guarantee its realization. The theoretical claim concerns a capacity to make appropriate decisions under enabling conditions, not an invariant personal effect on learning.

3.2. The Instructional Allocation of Epistemic Responsibility

The instructional allocation of epistemic responsibility denotes how a task assigns obligations to make and defend judgments about knowledge claims. Chinn et al. [5] distinguish epistemic aims, sources and justification, and processes for attaining epistemic ends. Drawing on that account, we locate instructional responsibility in the obligation to connect a financial conclusion to reasons that can be examined. This is a task-level arrangement, rather than a claim that students possess a stable moral or psychological

trait.

Responsibility is allocated to students when successful completion requires them to choose an assumption, explain why an inference follows, check a source or model relationship, or reconsider a conclusion when conditions change. It is supported when students have access to the concepts, information, feedback, and time needed to discharge that obligation. It is displaced when an external resource can supply the relevant decision and its apparent justification without requiring the learner to evaluate either. These distinctions concern educational design; they do not assign legal accountability to an AI system.

The construct adds specificity beyond cognitive engagement. A teacher can require an oral defense, provide reference information, and introduce an unannounced change in assumptions before any student responds. That design allocates responsibility regardless of whether a particular student reasons well, imitates an explanation, or refuses to participate. Engagement concerns the response. Likewise, responsibility differs from autonomy: a highly structured task can require substantial justification, whereas unrestricted choice among AI tools can leave every substantive judgment delegated. It differs from workload because extensive formatting or prompt revision can demand effort without requiring a financial inference.

A critical distinction separates nominals from feasible responsibility. An instruction to verify all generated content is nominal if students cannot identify relevant sources or lack the knowledge required to evaluate them. Feasible responsibility couples an obligation with an accessible basis for judgment. It does not ensure success, but it makes substantive engagement possible. The theory predicts consequences from this coupling rather than from the number of verification instructions appearing on an assignment sheet.

Engagement and the location of cognitive work

The ICAP framework distinguishes passive reception, active manipulation, constructive generation, and interactive development of ideas [6]. Its value here is to separate visible activity from contributions that extend the information provided. A student who repeatedly requests rewrites may be active without generating a new financial inference. A student explains why two valuations differ and contributes reasoning beyond reproducing either output. Interactive engagement requires substantive development through participants' contributions; the existence of a chat interface does not establish that process.

We use these distinctions as an analytic guide, not as an automatic classification of technologies. A chatbot exchange can elicit constructive student contributions, but conversational turns alone do not demonstrate the jointly generative interaction described in ICAP. The relevant question is where the inferential work occurs. If the system supplies each counterargument and the student merely selects one, the transcript can look intellectually rich while revealing little student construction. In contrast, a brief student explanation can expose a meaningful connection between a model and its assumptions.

Cognitive offloading supplies the complementary perspective. Risko and Gilbert [7] explain how people use external resources to reduce internal processing demands. Applied to instruction, reducing demand has no fixed educational meaning. Delegating a calculation already mastered may make room for interpreting its implications. Delegating the calculation when its structure is the learning

objective may remove essential practice. We therefore distinguish supportive offloading, which releases capacity for the focal learning objective, from substitutive offloading, which transfers the focal reasoning itself. These labels are relational: they depend on the objective and the learner's existing knowledge, not on an operation's permanent classification as routine or complex.

Automation and augmentation can occur within the same educational task. Some operations may be delegated while human activity is extended elsewhere. The temporal stakes matter: assistance can improve the present product while changing the opportunities through which future independent capability develops. Whether augmentation has occurred must therefore be judged against a specified learner capability and time horizon.

3.3. Outcomes and Disciplinary Scope

Independent financial understanding is the capacity to explain and apply relevant concepts when generative assistance does not supply the reasoning. It need not mean working without calculators, reference tables, or accessibility accommodation. The restriction concerns delegation of the target inference. Retention concerns the persistence of this understanding; transfer concerns its use when a problem's surface features or substantive assumptions change. Assisted product quality is a separate outcome describing what the student and available tools can produce together.

Justified financial judgment adds an evaluative dimension. Students must distinguish internal consistency from the plausibility of assumptions and from the adequacy of evidence. Risko and Gilbert's [8] account of valuation emphasizes matching cash flow and discount rates. A technically consistent model can nevertheless depend on contestable forecasts. For the purposes of this article, financial understanding includes recognizing both kinds of issues, while justified judgment includes explaining why a conclusion is warranted and how uncertainty limits it. A subsequently favorable market outcome would not, by itself, establish the quality of the original reasoning.

These features make finance an analytically useful setting without making the mechanism exclusive to finance. Many disciplines combine formal techniques and interpretive judgment. Finance makes their interaction conspicuous because numerical precision can coexist with considerable uncertainty about inputs and because a polished recommendation can conceal that uncertainty. The model should travel to other settings only where the relevant epistemic obligations and verification resources can be specified. It does not assume that finance students or institutions form a homogeneous population.

4. Conditional Model of AI Supported Learning

4.1. How Competence Becomes Consequential Through Task Design

The proposed sequence connects professional judgment to feasible responsibility, feasible responsibility to student engagement, and engagement to opportunities for independent understanding. Each connection is contingent. Competence can help a teacher identify which inference students need to practice, recognize where AI can complete it for them, and organize a task that requires their own evaluation. Its contribution lies in the alignment among these

decisions. Technical skills that only accelerate delivery of complete solutions can improve materials without activating this pathway.

Consider a teacher designing a valuation comparison. Selecting two outputs because they look different does not establish a useful learning contrast. Disciplinary judgment is needed to identify a difference that changes the financial conclusion. Pedagogical judgment is needed to make that difference accessible to the students' level. Knowledge of AI behavior is needed to anticipate that students may request a complete comparative explanation and to arrange a subsequent decision that cannot be discharged by submitting it. These judgments concern the design's responsiveness to delegation, rather than a general preference for making assignments difficult.

Proposition 1. Where teachers have discretion over instructional design, their generative AI integration competence is expected to support independent financial understanding to the extent that it is enacted in tasks assigning students feasible responsibility for the target reasoning.

This proposition specifies a pathway rather than complete mediation. Competence may also improve, for example accuracy or feedback clarity. Nor does it define competence by favorable outcomes: an instructor's decisions can be assessed for their disciplinary and pedagogical appropriateness before learning is observed. The proposition would lose explanatory value if feasible responsibility merely meant any design followed by good learning. It must instead be identifiable in advance through required judgments and the resources available to support them.

4.2. Supportive and Substitutive Offloading

Assistance can alter learning opportunities in opposite directions. Under supportive offloading, AI handles operations peripheral to the current objective, leaving students more capacity to examine the target relationship. Under substitutive offloading, the system completes that relationship and supplies a rationale that the task permits students to accept. Both pathways can improve the submitted product. Only the first preserves the specified opportunity for students to practice the focal reasoning, and even that opportunity requires uptake.

The distinction becomes especially important when generated output is correct. Correctness can protect students from misinformation, but it does not establish that they understand why an answer holds. A flawless explanation that is copied without evaluation can remove the same opportunity for conceptual work as a flawed explanation. Conversely, a deliberately prepared error can become useful material if students can diagnose it and receive appropriate feedback. Accuracy and allocation of responsibility are therefore different dimensions of design.

Proposition 2. For a given financial learning objective, AI assistance is more likely to support independent understanding when it offloads peripheral operations while retaining student responsibility for the target inference than when it supplies that inference without requiring student evaluation, even if both arrangements produce accurate submissions.

The learning rationale is that students need opportunities to generate connections they must later reconstruct. In ICAP terms, receiving a rationale and constructing or critically extending its inferential link offers different engagement opportunities [6]. Receiving a rationale can still provide

useful instruction; generating or critically reconstructing its connection creates a different opportunity. The proposition does not imply that every solution must be discovered unaided or that worked examples have no educational value.

4.3. Feasibility and The Limits of Assigned Responsibility

Allocating more responsibility can fail when the demands exceed the learner's resources. Cognitive load theory cautions that demanding problem solving can consume capacity without producing the intended schema acquisition [9]. The guidance learners need also varies with prior knowledge, so support appropriate for novices may become redundant for more knowledgeable learners. These considerations preclude a simple prescription to maximize independent critique from the beginning of a course.

For a novice, comparing two complete enterprise valuations may require simultaneously understanding cash-flow definitions, financing claims, terminal assumptions, and discounting. A general instruction to challenge AI could induce superficial disagreement or another request for AI-generated criticism. The teacher can instead establish one relevant relationship through a checked worked example, ask students to identify it in a second example, and subsequently require explanation under changed conditions. The responsibility becomes feasible through sequencing and feedback. For advanced students, retaining the same step-by-step support could unnecessarily constrain analysis.

Proposition 3. Increasing student responsibility for evaluating AI-generated financial reasoning is expected to support constructive engagement only when prerequisite knowledge, task complexity, and instructional guidance make that responsibility feasible; otherwise, additional evaluative demands can increase dependence on external explanations.

This proposition allows a nonmonotonic relationship between assigned responsibility and productive engagement. The point at which demands become counterproductive varies with students and tasks; the model does not posit a universal threshold. It also permits assistance to change across an instructional sequence. A generated explanation can initially support entry into a problem, while later tasks require students to reconstruct its logic. The theoretical requirement is to increase capability to exercise the target judgment, rather than immediate withdrawal of every form of help.

4.4. Verification Resources and Justified Engagement

Constructive contributions are not necessarily well justified. Students can generate elaborate explanations from invented figures, inaccurate premises, or mutually reinforcing AI outputs. Requiring more explanation can then strengthen a mistaken account. What matters is whether claims encounter a standard of evaluation external to the unsupported claim itself. In a finance task, that standard may be a checked source document, a clearly stated modeling convention, an independently performed calculation, or an explicitly hypothetical case with stipulated assumptions.

Verification is therefore broader than checking whether an answer agrees with another model. Several generated responses may share unsupported premises. Nor does a source's existence establish its relevance: a financial figure must match the intended period, unit, and entity. Students need a basis for distinguishing a factual error from reasonable disagreement about uncertain assumptions. An activity that

penalizes every departure from one model answer can suppress justified judgment even while making verification appear straightforward.

Proposition 4. The contribution of constructive engagement with AI-generated finance analyses to justified judgment is expected to weaken when students lack accessible sources or disciplinary criteria for evaluating the premises on which their explanations depend.

This is a boundary condition on the quality of engagement, not an argument that every task requires authentic corporate data. An explicitly hypothetical case can support sound reasoning when students know which quantities are stipulated and which judgments remain open. A hypothetical forecast presented as verified historical information cannot serve the same purpose. The key issue is whether the status of the material and the standards of appraisal are available to the learner.

4.5. Assessment Incentives and the Enactment of Responsibility

Students can bypass carefully designed opportunities when assessment rewards only the final product. Constructive alignment directs attention to consistency among intended outcomes, teaching activity, and assessment [10]. In AI-assisted work, this alignment must account for a changed route to completion: a student can obtain an acceptable rationale without having performed the reasoning. A rubric that awards points for the presence of a justification may therefore encourage the appearance of understanding rather than its exercise.

Assessment makes responsibility consequential when students must respond to a new condition, explain a chosen assumption, or defend an interpretation using accessible evidence. The purpose is to elicit reasoning that cannot be discharged solely through the original generated submission. Clear expectations and feasible demands matter because accountability can otherwise encourage memorized scripts, anxiety, or concealment. A process log is similarly insufficient if it records only activity volume or can be generated retrospectively.

Proposition 5. Holding the assigned reasoning obligations and available support constant, students are more likely to enact those obligations when assessment requires them to demonstrate understanding under changed conditions than when assessment rewards only the quality of an AI-assisted final product.

The proposition concerns incentives for uptake rather than the elimination of all strategic behavior. It predicts that changing assessment can alter the pathway even when the teacher, tool, and assignment topic remain constant. It also implies that low immediate product quality need not indicate poor learning opportunities: students who expose uncertainty and revise an initially weak explanation may undertake more consequential conceptual work than those who submit a polished analysis immediately. Product evaluation and evidence of understanding should inform one another without being collapsed into a single indicator.

5. Configurations and Finance Thought Experiments

5.1. Four Instructional Configurations

The model can be summarized through two analytically independent dimensions: how much target reasoning is

delegated to AI and how strongly students are required and supported to evaluate that reasoning. Extensive delegation can coexist with substantive student evaluation when generated analyses become objects of inquiry. Limited delegation can coexist with weak evaluation when students perform procedures mechanically. Table 1 presents ideal types, not categories recovered from observations or a validated classification instrument.

Table 1. Configurations of AI Delegation and Student Evaluation

AI delegation	Weak evaluation	Required and supported evaluation
Limited	Procedural completion. Students perform operations without explaining their significance.	Guided construction. Students perform target operations and explain them with calibrated support.
Extensive	Substitutive delegation. Students receive analyses and justifications without evaluating them.	Evaluative delegation. Students scrutinize generated analyses and defend judgments about adequacy.

Note: These ideal types are defined relative to the focal learning objective. They do not imply observed differences in learning outcomes.

The configurations clarify why neither AI restriction nor frequent verification instructions provide a sufficient design principle. Limited delegation protects an opportunity to perform an operation, but it does not ensure that students understand it. Extensive delegation can be compatible with learning when the objective is to evaluate model reasoning, and students actually conduct that evaluation. However, evaluative delegation does not preserve every learning opportunity. A task that develops criticism of generated code may still provide insufficient practice for learning to write the code independently. Configurations must therefore be classified relative to the specified objective.

Movement across configurations can occur within one course. An introductory task might use a checked worked example with limited student decisions. A later task might require learners to inspect an extensive generated analysis. The transition is educationally defensible when the learner's responsibility changes with the objective and accumulated knowledge. It is not defensible merely because later assignments involve longer AI conversations. The model concerns progression in what students can justify, rather than a progression in tool exposure.

5.2. Valuation Consistency and the Availability of Answers

Consider a hypothetical activity in which students compare two AI-generated accounts of a company valuation. One account uses cash flows to equity with a cost of capital intended for the firm; the other pairs cash flows to the firm with that cost of capital. The contrast draws on the cashflow and discount-rate consistency principle described by Risko and Gilbert [8]. The case is stipulated for conceptual analysis and reports no real company's value.

In a substitutive configuration, students ask the system to identify the superior account and submit its explanation. The task can yield a correct answer without revealing whether students understand the different claims on cash flows. In a configuration with feasible responsibility, the teacher provides checked definitions and requires students to explain

who receives each cash flow, identify which rate corresponds to it, and respond independently when the cash-flow definition changes. AI can assist with arithmetic or initial exposition, while the focal matching judgment remains consequential for the learner.

Now suppose both generated accounts are internally consistent. A design focused only on finding an AI error has lost its purpose. A responsibility-based design remains meaningful because students must explain the consistency and determine whether the alternatives answer the same valuation question. This counterfactual exposes a theoretical distinction: the model's learning pathway depends on retained judgment, not on the presence of a machine mistake. It also guards against teaching students that indiscriminate disagreement is evidence of critical thinking.

5.3. Investment Judgment Under Uncertain Assumptions

A second hypothetical activity supplies a checked case description and two defensible forecasts for an investment project. An AI system produces a recommendation under each forecast. The exercise asks students to identify which assumptions account for the difference and how strongly the available evidence favors either one. Because future conditions are uncertain, a single recommendation is not treated as the only acceptable answer. Students must justify the relation between their confidence and the evidence.

With extensive delegation and weak evaluation, students can choose the most persuasive narrative. With extensive delegation and supported evaluation, they can use both narratives as alternatives to examine, identify an unsupported leap, and state which new information would change their judgment. The AI-produced analysis is then an input to evaluation. If the teacher omits the case evidence and asks novices simply to verify the forecasts, the same activity becomes nominally demanding but difficult to evaluate. It can drift toward generated criticism of generated claims. This limiting case illustrates the joint operation of feasibility and verification conditions.

5.4. Analytical Code and Changing Learning Objectives

A third thought experiment concerns AI-generated code for portfolio analysis. If the objective is interpreting how estimated relationships affect portfolio risk, delegating routine code construction may support conceptual work, provided students can examine inputs and identify consequential assumptions. If the objective is independently implementing the calculation, delegating the implementation removes a different opportunity. An identical tool interaction can therefore be supportive in one lesson and substitutive in another.

The distinction also applies to assessment. Executable code and attractive graphs provide evidence about the joint product. A student's ability to predict and explain what changes when an input relationship change provides different evidence about understanding. Neither outcome exhausts professional competence, which may appropriately include productive work with tools. The theoretical point is that an educator must specify the capability being developed before judging whether delegation advances it.

6. Theoretical Implications and Alternative Explanations

6.1. A Mechanism Linking Professional Knowledge and Student Activity

The model develops the teacher-knowledge account by identifying an instructional relation through which knowledge can become consequential. It does not propose another comprehensive list of AI competencies. Its explanatory focus is the allocation of obligations under conditions in which generated systems can perform both a task and its apparent justification. A teacher's capability matters when it changes this allocation in ways suited to learners and objectives. The same capability may have little visible consequence when institutional conditions prevent enactment.

This move also clarifies the place of student agency. Students are not passive recipients of a competence effect. They may accept, reinterpret, or evade assigned obligations; they may also choose to reason more deeply than assessment requires. The model consequently separates opportunities from uptake and leaves room for motivation, confidence, and prior experience to shape enactment. Its claim is that task design changes the conditions of agency, not that teachers determine learners' cognitive states.

6.2. Extending the Account of Engagement Under Generative Assistance

The conceptual extension to engagement concerns the ambiguity of observable output when another system can supply the apparent contribution. An explanation appearing in a submission is no longer sufficient to locate constructive activity in the student. This does not invalidate engagement theory. It requires attention to the provenance and enactment of reasoning rather than treating a finished artifact as a transparent record of cognition. The distinction is especially relevant when the AI system supplies critiques, revisions, and reflections that resemble evidence of learning.

The theory also separates the quantity of cognitive work from its epistemic quality. A learner can construct an explanation that rests on a false premise or revise a position for an irrelevant reason. Feasible responsibility must therefore include an accessible basis for appraisal. The resulting contribution is a conditional account of productive engagement: the target inferential work remains with the learner, its demands are manageable, and its premises can be examined. These conditions identify different points of failure rather than a single undifferentiated lack of engagement.

6.3. Rival Accounts and Discriminating Implications

General instructional quality is the strongest rival explanation. Clear teaching, useful feedback, and aligned assessment can improve learning without AI. The present model does not claim that these principles originated with generative technology. Its additional explanatory claim is that AI can redistribute the target inference itself, including the explanation used to infer understanding. If different allocations of this work made no difference once general instructional quality was considered, the model would offer little beyond a contextual restatement of good teaching.

A second rival is an accuracy account in which better learning follows primarily from exposure to more correct

material. Accuracy remains important, but Proposition 2 predicts a difference between equally accurate materials that leave different reasoning obligations to students. Conversely, a difference confined to inaccurate outputs would favor an error-correction explanation over the broader responsibility mechanism. The first thought experiment makes this distinction visible by retaining a meaningful learning task after every generated statement is corrected.

A third rival emphasizes time saved. More efficient assistance may permit additional practice, independent of who performs each inference. The model predicts that time released from peripheral work benefits the focal outcome only when it creates or supports relevant opportunities that students use. If efficiency alone accounted for gains regardless of retained reasoning, the allocation argument would need revision. A novelty account would similarly attribute short-term participation to interest in the tool; sustained understanding after novelty declines would be more consistent with the proposed learning process.

Finally, student preparation can explain both deeper engagement and stronger understanding. The model treats prior knowledge as a feasibility condition rather than assuming that observed engagement caused subsequent capability. Its propositions describe expected patterns under stated comparisons, not causal estimates. Their value lies partly in exposing what would count against the explanation: equally durable independent understanding despite systematic delegation of the target reasoning would challenge the importance assigned to practicing that reasoning, unless another identifiable learning process supplied it.

7. Implications for Pedagogy and Institutional Practice

The framework suggests that teacher development should center on consequential instructional decisions. A useful exercise would ask instructors to identify the target inference in an existing assignment, determine which parts an AI system can complete, and redesign the task so that students must exercise a feasible judgment. Peer review could examine whether the activity provides appropriate evidence and feedback, rather than counting the number of AI tools included. These are deductions from the conceptual model, not demonstrated effects of a training program.

Assessment can distinguish joint production from independent understanding without treating all external assistance as incompatible with learning. An AI-assisted report can be accompanied by a brief explanation of a consequential choice or an individual response to a changed assumption. The independent component should match the stated learning objective and retain appropriate accommodations. For example, an oral defense should not inadvertently make language fluency the dominant criterion when the intended construct is financial reasoning. A short-written explanation or an annotated model may serve the same purpose more appropriately in some settings.

Institutional support matters because feasible responsibility has resource requirements. Shared, checked cases can make verification accessible, and common evaluation criteria can reduce ambiguity about what counts as a warranted financial claim. Tool access and instructional time should be considered when assigning tasks. Students should not be assessed as less capable merely because their available systems offer fewer features or because relevant reference

materials are inaccessible. The allocation of responsibility is educationally defensible only when opportunities to exercise are reasonably supported.

Institutional approaches to human agency and ethical AI integration become concrete in this framework through a specific instructional question: can students examine and defend the reasoning for which they are held responsible? Acknowledging AI assistance supports transparency, but acknowledgment alone does not answer that question. Likewise, counts of AI-enabled courses describe adoption rather than the development of understanding. Implementation decisions should distinguish these objectives rather than treating participation as evidence of educational value.

8. Limitations and Future Theoretical Development

This paper establishes a conceptual explanation, not the prevalence, magnitude, or causal strength of its pathways. The literature is selected for theoretical relevance and does not support claims of comprehensive coverage. The propositions may need revision as alternative accounts of human learning with generative systems become more precise. Their present contribution is to organize a bounded problem, distinguish constructs, and derive implications that are not reducible to a uniformly beneficial or harmful view of AI.

Several conceptual boundaries require further work. First, feasible responsibility is a relational property: identical instructions can be manageable for one learner and inaccessible to another. It therefore cannot be reduced to assignment wording alone. Second, responsibility in collaborative finance tasks can be distributed among students in ways that conceal individual dependence. A group may produce well-justified work while only one member performs the critical evaluation. Extending the model to groups requires specifying how obligations and opportunities are distributed within them.

Third, the model emphasizes independent understanding because it is an important undergraduate learning aim. Professional capability also includes selecting, supervising, and correcting external tools. Independent and assisted capabilities can develop together, but they should not be presumed interchangeable. A fuller developmental theory would explain when foundational reasoning must be internalized, which operations can remain externally supported, and how learners retain the ability to recognize failures in that support. The present account identifies the problem without establishing a universal curriculum sequence.

The propositions also generate a focused agenda for subsequent inquiry without making new data collection a component of this article. The most informative comparisons would hold the financial content and AI output quality constant while changing who must justify the target inference. Further comparisons could examine whether prerequisite guidance changes the effect of assigned evaluation demands, whether accessible evidence changes the quality of criticism, and whether assessment conditions change student uptake. Such inquiries need not begin with a broad self-report survey; task artifacts, instructional comparisons, and accounts of reasoning could address different parts of the model.

Any later empirical extension should preserve the distinctions developed here. Teacher competence, task-level

responsibility, student enactment, and learning over time are different objects of explanation. Measures that ask teachers whether their AI use improves student learning would collapse them. Likewise, an AI-assisted assignment grade would not establish independent understanding, and one immediate unassisted test would not establish durable transfer. These distinctions clarify the kinds of evidence that could differentiate the proposed explanation from its rivals.

9. Conclusion

Generative AI makes the allocation of reasoning an explicit problem for undergraduate finance education. An accurate and persuasive analysis can be educationally valuable, yet its production does not establish that the student can justify it. This paper explains teacher competence through its enactment in tasks that assign learners feasible responsibility for the financial judgments they are expected to acquire.

The proposed model distinguishes assistance that releases capacity for target reasoning from assistance that substitutes for it. It specifies why prior knowledge, verification resources, and assessment incentives condition the resulting opportunities for engagement and understanding. The central educational criterion is therefore whether students become able to explain, evaluate, and revise the reasoning behind an analysis. Preserving that capability requires attention to the distribution of intellectual work throughout the learning activity, including when AI-generated answers are entirely correct.

References

- [1] Bastani, H., Bastani, O., Sungu, A., et al. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>
- [2] Kestin, G., Miller, K., Klates, A., et al. (2025). AI tutoring outperforms in-class active learning: An RCT introducing a novel research-based design in an authentic educational setting. *Scientific Reports*, 15, 17458. <https://doi.org/10.1038/s41598-025-97652-6>
- [3] Mishra, P., & Koehler, M. J. (2006). Technological pedagogical content knowledge: A framework for teacher knowledge. *Teachers College Record*, 108(6), 1017–1054. <https://doi.org/10.1111/j.1467-9620.2006.00684.x>
- [4] Celik, I. (2023). Towards Intelligent-TPACK: An empirical study on teachers' professional knowledge to ethically integrate artificial intelligence (AI)-based tools into education. *Computers in Human Behavior*, 138, 107468. <https://doi.org/10.1016/j.chb.2022.107468>
- [5] Chinn, C. A., Buckland, L. A., & Samarapungavan, A. (2011). Expanding the dimensions of epistemic cognition: Arguments from philosophy and psychology. *Educational Psychologist*, 46(3), 141–167. <https://doi.org/10.1080/00461520.2011.587722>
- [6] Chi, M. T. H., & Wylie, R. (2014). The ICAP framework: Linking cognitive engagement to active learning outcomes. *Educational Psychologist*, 49(4), 219–243. <https://doi.org/10.1080/00461520.2014.965823>
- [7] Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688. <https://doi.org/10.1016/j.tics.2016.07.002>
- [8] Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4
- [9] Biggs, J. (1996). Enhancing teaching through constructive alignment. *Higher Education*, 32, 347–364. <https://doi.org/10.1007/BF00138871>