

**Opportunity-Cost Literacy and Ethical AI Decision-Making: A Framework for Jesuit  
Business Education**

**Or**

**Does AI Cost More Than It Saves? Teaching Opportunity-Cost Literacy for Strategic  
Decision-Making in Jesuit Business Education**

Maurice N. Emelu, John Carroll University

Tina Facca-Miess, John Carroll University

## Abstract

Generative and agentic artificial intelligence are increasingly presented to business schools as tools for efficiency, innovation, productivity, and competitive advantage. Yet a central question remains underdeveloped in business education: Does AI cost more than it saves? This paper proposes opportunity-cost literacy as a pedagogical framework for ethical and strategic AI decision-making. Opportunity-cost literacy is the cultivated capacity to evaluate not only what AI saves in time, labor, and cost, but also what it may displace in human judgment, creativity, accountability, relational trust, skill formation, institutional memory, and responsibility for the common good. Drawing on Catholic Social Teaching, Pope Francis's critique of the technocratic paradigm, Pope Leo XIV's *Magnifica Humanitas*, AI ethics scholarship, cognitive load theory, craft theory, virtue ethics, and Ignatian pedagogy, the paper develops a teachable model for business education, the AI Opportunity-Cost Decision Lab. It advances a pedagogical tool for the classroom where students compare human-led, AI-assisted, hybrid, and agentic-AI-delegated workflows; log visible and hidden costs; assess stakeholder and mission implications; and produce a strategic discernment memo. The model is informed by exploratory classroom applications in media writing and media production. These examples function as pedagogical illustrations rather than controlled empirical evidence. Although developed for Jesuit business education, the framework is adaptable to secular and other mission-driven business schools seeking to prepare students for responsible AI use.

**Opportunity-Cost Literacy and Ethical AI Decision-Making: A Framework for Jesuit  
Business Education**

**Or**

**Does AI Cost More Than It Saves? Teaching Opportunity-Cost Literacy for Strategic  
Decision-Making in Jesuit Business Education**

**Introduction: The Missing Question in AI Business Education**

Generative artificial intelligence has moved rapidly into business education and organizational life. Students use AI to draft reports, summarize readings, generate business ideas, prepare presentations, conduct preliminary market analysis, produce marketing content, and support strategic planning. Organizations use AI to automate routine work, personalize communication, model customer behavior, accelerate decision processes, and reduce costs. Increasingly, the conversation is also shifting toward agentic AI, systems that can plan, sequence tasks, call tools, coordinate subtasks, retrieve information, and complete more complex workflows with less continuous human prompting. Business schools therefore face a practical and urgent question: How can students learn to use AI effectively?

The question is important, but incomplete. AI literacy often focuses on what AI is, how it works, how to prompt effectively, how to identify bias or hallucination, and how to use tools responsibly. These competencies are necessary in today's digital media landscape. Yet they do not fully address the strategic and moral question business leaders face. When AI appears to save time, labor, and cost, students and organizations still have to ask: What else might it displace? Put more positively, what is gained, what is lost, and how can opportunity-cost assessment help users benefit from AI while avoiding trade-offs that may weaken the long-term health of persons and organizations?

This paper argues that business education needs AI opportunity-cost literacy. In this paper, opportunity cost is not used only in the economic sense of the next-best alternative forgone. It is expanded into a formative and ethical category: the human, relational, institutional, and moral capacities that may receive less practice or less attention when AI mediates or performs a task. Opportunity-cost literacy is therefore the cultivated capacity to evaluate visible gains alongside hidden costs, including weakened judgment, reduced skill development, diminished creativity, overreliance on automated outputs, loss of institutional memory, erosion of relational trust, and weakened responsibility for stakeholders.

The rise of agentic AI makes this need more urgent. With basic generative AI, the comparison is often between human-led work and AI-assisted drafting or analysis. With agentic AI, the comparison expands to AI-delegated work. The question is no longer only whether AI helps produce an output more quickly. The deeper question is what forms of judgment, responsibility, skill, attention, organizational learning, and stakeholder accountability are transferred away from the human agent when an AI system performs a larger portion of the task. The argument is especially relevant to Jesuit business education because Jesuit schools do not educate students merely for technical competence or professional advancement. They seek to form leaders capable of discernment, justice, responsibility, service, and care for the whole person (*cura personalis*). What Jesuit education seeks is not simply more, but better: *magis* understood as a deeper commitment to human flourishing, discernment, and the common good. The AI question therefore belongs within the broader mission of business leadership formation. Students need to ask not only, "Can AI do this task?" but also, "Is AI appropriate for this task?" "What human capacity may receive less practice if AI does it?" "Who benefits?" "Who bears the

hidden cost?" "What decisions did the AI system make along the way?" and "Does this use of AI serve human dignity and the common good?"

The contribution of this paper is pedagogical and conceptual. Conceptually, it reframes opportunity cost as a category for ethical AI decision-making in business education. Pedagogically, it offers a concrete classroom model—the AI Opportunity-Cost Decision Lab—that faculty can adapt in courses on management, marketing, business ethics, analytics, entrepreneurship, leadership, communication strategy, organizational behavior, public relations, nonprofit management, and media production. The lab asks students to test AI claims through workflow comparison, cost logging, stakeholder analysis, and a strategic discernment memo. This paper builds on three strands of prior scholarship while developing a new business education application. First, formation-based AI ethics argues that AI systems are not merely tools but formative environments that shape patterns of attention, judgment, responsibility, and delegation over time; it also identifies opportunity cost as an underdeveloped ethical category in AI ecosystems (Emelu, 2026a). Second, a study of Pope Leo XIV's discourse on AI and digital media shows how recent papal teaching frames AI through human dignity, relational anthropology, moral discernment, and ethical formation (Emelu, 2026b). Third, the OAEPA model—Observation, Analysis, Evaluation, Production, and Assessment—offers a pedagogical strategy for teaching digital and AI-shaped learners through attentiveness, reflection, creative agency, and formative assessment (Emelu, 2025).

This paper focuses primarily on the LLM-use and application layers of AI, where generative and agentic systems produce content, support decisions, create products, or deliver consumer and organizational services. Opportunity-cost literacy can also apply to the model-development layer, where questions arise about training data, design values, alignment choices,

labor, safety, environmental costs, and accountability. Those questions are important, but they fall outside the scope of this paper. The energy layer, chip and computing layer, cloud infrastructure layer, and model-development layer shape the conditions under which AI applications exist, but the present argument concerns the application layer, where AI meets actual human life in education, business communication, healthcare, finance, media, family life, faith formation, pastoral care, and organizational decision-making. Opportunity-cost literacy, as developed here, asks what happens when students, professionals, and organizations use AI to draft, analyze, recommend, plan, communicate, produce, or partially delegate work.

The paper proceeds by first explaining why AI disruption requires business education to move beyond tool-focused AI literacy. It then situates the argument within Catholic Social Teaching, develops opportunity-cost literacy as an ethical and strategic decision-making capacity, connects the framework to formation and virtue, adapts Ignatian pedagogy and the OAEPA model to AI decision-making, presents the AI Opportunity-Cost Decision Lab and exploratory classroom illustrations, offers a worked example, and concludes with broader applications for business leadership formation.

### **AI Disruption and the Limits of Tool-Based AI Literacy**

AI literacy has become an essential part of business education. Students need to understand algorithmic systems, data, automation, generative outputs, prompting, human oversight, privacy risk, bias, and hallucination. Scholarship on AI literacy identifies competencies such as understanding what AI can and cannot do, recognizing how AI systems affect decision-making, and evaluating when human judgment remains necessary (Long & Magerko, 2020; Ng, 2021). These competencies are important, but they often remain tool-centered.

A tool-centered approach asks whether students can use AI effectively. A strategic approach asks whether AI improves the decision. An ethical approach asks whether AI use respects persons, stakeholders, and institutions. Opportunity-cost literacy integrates these concerns. It asks whether a particular use of AI strengthens or weakens the human and organizational capacities that good business judgment requires.

This distinction is important because AI's visible benefits often appear immediately. A student can produce a polished market summary in minutes; a manager can generate a performance review draft quickly; a marketing team can produce campaign copy at scale. An AI agent may go further by gathering background material, drafting a memo, creating a task list, generating outreach language, and recommending next steps. The visible gain is speed and expanded capability. The hidden costs are slower and less visible: fewer opportunities to practice analysis, overconfidence in plausible outputs, weakened writing ability, reduced stakeholder attentiveness, loss of contextual nuance, and a subtle shift from judgment to automated suggestion. These costs matter because the health of any organization depends on more than automated output. Innovation and trust often emerge in the gray areas where judgment, experience, and human attention are decisive.

AI adoption in business and professional education therefore requires judgment about trade-offs. Behavioral economics shows that human beings often misjudge costs, especially when immediate gains are obvious and deferred losses are harder to see (Kahneman, 1973; Kahneman & Tversky, 1979). Cognitive load theory similarly shows that tools can reduce cognitive burden in some contexts while increasing it in others through troubleshooting, revision, verification, and error correction (Sweller, 1988; Paas et al., 2003; Mayer, 2020). In AI use, the

apparent savings of automation may be offset by prompting, re-prompting, checking, editing, correcting, explaining, and defending the output.

Agentic AI complicates this further. Because AI agents can complete multi-step workflows, they may reduce some forms of manual labor while increasing hidden process costs. These costs include reduced process visibility, difficulty tracing how a recommendation emerged, greater verification burden, dependence on assumptions embedded in the agent's workflow, and blurred accountability when an agent makes intermediate choices. In such cases, the opportunity cost may lie less in the quality of a single output and more in the loss of procedural understanding, reduced human oversight, and weakened formation of professional judgment.

Business education needs to make these costs visible through disciplined comparison rather than novelty or convenience alone. AI can save time, improve access to information, support creativity, and help people work more effectively in many contexts. Yet students still need to ask: Does AI actually save time in this case? Does it improve quality? Does it strengthen or weaken skill development? Does it reduce cognitive burden, or merely relocate it? Does it support stakeholder trust, or undermine it? Can the student explain how the AI system reached its output? Does the use of an AI agent strengthen human judgment by freeing people for higher-order discernment, or weaken judgment by removing the formative struggle through which expertise develops?

These questions do not begin from suspicion of AI. They begin from responsible adoption. In many cases, AI may offer clear and meaningful benefits. Yet grounded answers are rarely automatic. Even the claim that AI saves time deserves scrutiny. A skilled writer responding to a customer's query after a recent boardroom decision may rely on discretion,

context, institutional memory, and awareness of what has or has not yet been made public. An agentic AI email manager may produce a faster response, but that speed must be weighed against possible costs to judgment, trust, confidentiality, and later repair.

### **Catholic Social Teaching, Integral Human Development, and AI**

A Jesuit business education framework for AI decision-making needs grounding not only in general AI ethics, but also in the moral tradition that gives Jesuit education its distinctive character. Catholic Social Teaching evaluates economic, technological, and organizational life according to the dignity of the human person, the dignity of work, solidarity, subsidiarity, the common good, justice, the preferential option for the poor, and integral human development. Catholic Social Teaching has long responded to economic and technological disruption by asking what such disruption does to persons, workers, families, communities, and social life. From Leo XIII's *Rerum Novarum* through Paul VI's *Populorum Progressio*, John Paul II's *Laborem Exercens*, Benedict XVI's *Caritas in Veritate*, and Pope Francis's *Laudato Si'* and *Fratelli Tutti*, the tradition consistently resists reducing economic and technological development to productivity, efficiency, or growth alone (Leo XIII, 1891; Paul VI, 1967; John Paul II, 1981; Benedict XVI, 2009; Francis, 2015, 2020). Development, in this tradition, concerns the whole person and every person.

Pope Francis's critique of the technocratic paradigm is especially important for AI. In *Laudato Si'*, Francis argues that modern technological power can narrow human reason by treating reality as raw material for control, efficiency, extraction, and manipulation (Francis, 2015, nos. 106–114). Far from being an anti-technology or progress critique, the papal position warns against allowing technical power to become the measure of all value. In his 2024 G7 address on artificial intelligence, Francis applied this concern directly to AI, insisting that

technological development remain accountable to human dignity, freedom, responsibility, and political discernment (Francis, 2024).

Building on that tradition, Pope Leo XIV's *Magnifica Humanitas* develops a more direct path for the AI age. The encyclical places AI within the long arc of Catholic Social Teaching and frames it as one of the defining moral and social questions of the present moment. AI differs from many earlier technologies because it increasingly mediates cognition, attention, judgment, communication, professional identity, relational life, and moral responsibility. For business education, this means AI cannot be evaluated only by productivity, competitiveness, innovation, or efficiency. It must also be evaluated by how it shapes persons, organizations, and communities.

The central contribution of *Magnifica Humanitas* for this paper is its insistence that AI must be evaluated according to the safeguarding and integral development of the human person. Integral development means more than economic growth, technical capacity, or institutional efficiency. It concerns the whole person and the whole community: material well-being, moral agency, social participation, relational life, meaningful work, spiritual depth, and openness to truth and communion. Applied to AI, the evaluative question is not simply, "What does this technology optimize?" but "Does this technology serve the full dignity and vocation of the human person?"

This point connects directly a prior analysis of Pope Leo XIV's early discourse on AI and digital media, which identifies ethical formation grounded in anthropology and the primacy of relational anthropology as central to the Pope's emerging AI thought (Emelu, 2026b). That analysis argues that Leo XIV's AI discourse evaluates technology in relation to human dignity, relational intelligence, moral discernment, and formation-based ethics (Emelu, 2026b). In the

present paper, that insight is applied to business education: AI use in organizations must be evaluated not only by output quality, but also by how it shapes persons, relationships, institutions, and communities.

Catholic Social Teaching helps name the hidden costs that opportunity-cost literacy seeks to make visible. The dignity of work reminds business educators that work is not merely a means of production but a site of human development, participation, responsibility, and social contribution. The common good requires leaders to ask how AI decisions affect the conditions that allow persons and communities to flourish. Solidarity requires attention to those who bear the hidden costs of automation. Subsidiarity cautions against removing meaningful agency from persons and communities when decisions can and rightly remain closer to those affected. The preferential option for the poor asks who is excluded, exploited, rendered invisible, or made more vulnerable by AI adoption.

These concerns become even sharper with agentic AI. When an AI system does more than produce a response—when it coordinates steps, chooses sources, sequences actions, and recommends decisions—the question of subsidiarity and human agency becomes more pressing. Who remains close enough to the decision to understand it? Who has the power to intervene? Who can explain the steps taken? Who is responsible when a delegated AI workflow produces harm? For Catholic Social Teaching and Jesuit business education, these questions are not peripheral. They are central to safeguarding the human person in technological systems. In this respect, Catholic Social Teaching strengthens the theoretical foundation of opportunity-cost literacy while keeping the paper focused on business education. AI must be evaluated not only by accuracy, speed, compliance, or return on investment, but also by what it does to human judgment, work, trust, responsibility, and the common good.

## **Opportunity-Cost Literacy as a Business Education Framework**

Opportunity cost traditionally refers to the value of the next best alternative forgone when a decision is made (Becker, 1965). In this paper, the concept is extended beyond economics. It refers to the human, organizational, moral, relational, and strategic capacities that may receive less practice, less attention, or less institutional support when AI systems substitute for, mediate, or reorganize human work.

This broader understanding draws directly on formation-based AI ethics, which identifies opportunity cost as a neglected ethical category in AI ecosystems (Emelu, 2026a). That framework argues that AI systems shape patterns of attention, delegation, responsibility, and judgment over time. The present paper adapts that claim to business education by proposing opportunity-cost literacy as a teachable decision-making capacity.

Opportunity-cost literacy trains students to ask seven questions:

1. What does AI save in time, cost, labor, or cognitive effort?
2. What new work does AI create through prompting, checking, revising, correcting, or explaining?
3. What human skill may receive less practice if AI performs this task?
4. What forms of judgment, interpretation, or stakeholder understanding may be weakened?
5. What relational or trust-based costs may arise if stakeholders know or suspect AI was used?
6. What organizational habits may form through repeated reliance on AI?
7. Does this AI use serve human dignity, responsible stewardship, integral human development, and the common good?

Agentic AI requires an expanded version of this questioning. Because agentic systems can complete larger portions of a workflow, students must also ask: What decisions did the AI agent make along the way? Which assumptions did it embed in the process? Which sources, tools, or criteria did it select? Which steps did the human user no longer practice or understand? What parts of the process became less visible? Who remains accountable if the agent's work is inaccurate, biased, incomplete, or misaligned with organizational mission? Does the agentic workflow free the human person for higher-order judgment, or does it weaken judgment by reducing the human role to acceptance, correction, or after-the-fact approval?

This framework shifts AI education from tool adoption to strategic discernment. Students learn that a tool can be useful and still costly. Saving time does not automatically improve a decision. Efficiency can support mission, but it can also displace the slow practices through which mission is understood and embodied. In agentic AI contexts, delegation may make a process appear seamless while hiding assumptions, trade-offs, and intermediate judgments embedded in the workflow.

Opportunity-cost literacy is therefore not anti-technology. It is a disciplined way of choosing technology well. It helps students avoid both naïve enthusiasm and fearful rejection by asking whether a particular use of AI, in a particular context and for a particular stakeholder purpose, strengthens or weakens the organization's mission and moral responsibilities.

### **Formation, Virtue, and Human Judgment**

The theoretical foundation for opportunity-cost literacy also draws from virtue ethics and formation-centered approaches to technology. Much contemporary AI ethics has focused on principles such as fairness, transparency, accountability, privacy, explainability, beneficence, nonmaleficence, and justice. These principles are important. Comparative analyses show both the

proliferation of AI ethics frameworks and convergence around a smaller set of high-level commitments (Floridi & Cowls, 2019; Jobin et al., 2019). However, principles alone do not guarantee ethical practice. Ethical guidelines do not ensure that designers, managers, students, or organizations have the moral capacity to interpret and enact them well.

Mittelstadt (2019) identifies this gap in critiques of AI principlism, arguing that principles alone cannot guarantee ethical AI because they often lack professional norms, accountability mechanisms, and institutional structures needed for implementation. Formation-based AI ethics responds by asking how human agents and institutions are formed to recognize responsibility, exercise judgment, and attend to hidden costs across the AI lifecycle (Emelu, 2026a).

Aristotelian virtue ethics offers an important foundation here. For Aristotle, moral excellence develops through habituation and practical wisdom (Aristotle, 1999). MacIntyre (2007) extends this account by situating moral agency within practices, traditions, and internal goods. Applied to business education, this means students do not develop judgment simply by learning rules. They develop judgment by practicing discernment within concrete situations, weighing competing goods, receiving feedback, and forming habits of responsible action.

Technology studies support this concern. McLuhan (1964) argued that media shape perception and social organization beyond the content they transmit. Postman (1992) warned that technologies carry implicit philosophies and can reorder culture around efficiency, control, and calculation. Floridi (2014) describes digital technologies as reshaping the infosphere in which human identity and agency unfold. Vallor (2016) argues that emerging technologies require renewed attention to technomoral virtues, including honesty, humility, justice, courage, empathy, care, civility, flexibility, perspective, magnanimity, and practical wisdom.

For business education, the implication is direct: AI does not merely help students complete work. Repeated use of AI shapes how students approach work. If students habitually ask AI to generate ideas before they have struggled with a problem, they may lose opportunities for creative and analytical formation. If managers habitually ask AI to rank candidates, summarize employees, or recommend decisions, they may gradually weaken their own capacities for discernment. If organizations treat AI as a substitute for difficult human conversations, they may reduce some efficiency costs while increasing moral and relational costs.

Agentic AI intensifies this formative concern. When AI systems can perform multi-step tasks, users may lose practice not only in producing outputs, but also in understanding processes. The opportunity cost may include reduced procedural knowledge, weaker capacity to diagnose failure, and diminished confidence in defending a decision. A leader who relies on an AI agent to produce a market-entry recommendation may receive a coherent report while having little understanding of the assumptions, sources, exclusions, or trade-offs that shaped it. The output may look professional, but the leader's judgment may have received less formation. The goal of opportunity-cost literacy is to make these formative effects visible and teachable.

### **Ignatian Pedagogy and AI Decision-Making**

Ignatian pedagogy offers a practical educational structure for opportunity-cost literacy. Its core movements—context, experience, reflection, action, and evaluation—invite learners to examine reality concretely, reflect on meaning, and act responsibly (International Commission on the Apostolate of Jesuit Education, 1993). In a business education setting, this tradition can help students approach AI not only as a tool for task completion but as an occasion for discernment.

The OAEPA strategy--Observation, Analysis, Evaluation, Production, and Assessment--can be adapted for AI decision-making in business education. The author previously developed OAEPA as a pedagogical strategy for teaching theology to digital and AI-shaped learners, connecting it to Ignatian discovery, cura personalis, reflection, action, and formative assessment (Emelu, 2025). In that model, observation begins with attentiveness to context and experience; analysis deepens understanding; evaluation invites personal and moral integration; production moves learning into creative action; and assessment functions as a formative process of continued growth (Emelu, 2025).

In the present paper, OAEPA is adapted to business education. Observation asks students to attend to the organizational context, stakeholder realities, and pressures shaping AI adoption. Analysis asks them to compare AI's strategic benefits, risks, and alternatives, including the difference between assistance and delegation. Evaluation asks them to weigh moral, relational, and mission-related trade-offs. Production asks them to create a decision artifact, such as an AI-use recommendation memo or responsible AI policy. Assessment asks them to reflect on what the process revealed about their own assumptions, judgment, and leadership formation. In this model, AI decision-making becomes a formative practice. Students do not merely learn to use a tool; they learn to evaluate a tool in light of persons, purpose, and consequences. In agentic AI contexts, they also learn to evaluate whether a workflow remains sufficiently transparent, accountable, and formative for the human agent.

### **The AI Opportunity-Cost Decision Lab**

The AI Opportunity-Cost Decision Lab operationalizes opportunity-cost literacy for business and professional education. It can be used in courses in management, marketing,

business ethics, analytics, entrepreneurship, organizational communication, leadership, strategic communication, nonprofit management, public relations, and media production.

The lab asks students to compare human-led, AI-assisted, hybrid, and, where appropriate, agentic-AI-delegated approaches to a specific professional task. The goal is not to prove that AI is helpful or harmful in general. The goal is to teach students how to evaluate AI use with strategic and moral discipline. Students learn to treat AI use as a decision requiring evidence, judgment, stakeholder awareness, and reflection.

### **Exploratory Classroom Applications**

The framework proposed here is informed by exploratory classroom applications in two communication courses: a media writing course and a media production course. These applications were not designed as controlled empirical studies and did not include experimental controls, pre/post measures, or formal assessment instruments sufficient to determine statistical precision. They function here as pedagogical illustrations that helped clarify the need for opportunity-cost literacy and informed the development of the framework proposed in this paper. In the media writing course, students used generative AI tools to draft public relations materials, including a standard news release and a reaction statement. The assignment asked students to compare AI-generated drafts with professional expectations for accuracy, tone, newsworthiness, audience awareness, organizational voice, and strategic clarity. Student responses varied. Some found AI useful for structure, headline options, and initial phrasing. Others expressed less confidence in submitting a final copy produced primarily by a large language model because the drafts often required substantial revision for specificity, tone, organizational context, and professional judgment. The exercise helped students see that AI could accelerate a first draft but

could not replace the writer's responsibility for accuracy, credibility, audience sensitivity, and strategic intent.

A second exploratory application occurred in a media production course. Students were asked to use generative AI tools, including ChatGPT and Gemini, to script a documentary spin from a series of raw documentary footage based on their descriptions of the project's possible direction. The students had already watched the raw footage and had contextual knowledge of the story, characters, tone, and available material. After several attempts and close to an hour of prompting and revising AI outputs, some students concluded that the tools did not meet their needs. They recognized that they could likely have written a more precise and usable script in less than an hour because they had already engaged the footage directly. The exercise made visible an important opportunity cost: AI did not simply save time; it shifted labor into prompting, correction, interpretation, and frustration while producing outputs that lacked the specificity and narrative judgment students had developed through direct viewing.

These classroom experiences revealed different levels of student assessment. Some students saw AI as useful for early ideation, structure, and language variation. Others recognized that AI-generated materials required extensive human correction before they could meet professional standards. Across both courses, the most valuable learning was not whether AI "worked" or "failed" in a general sense, but whether students learned to ask more precise questions: Did AI save time after revision was included? Did the output reflect the assignment's context? Did the tool weaken or support their professional judgment? Did the AI-generated draft carry the right voice, tone, and ethical responsibility? Could they defend the final product as their own professional work?

Because these applications were exploratory, they should not be presented as empirical evidence of AI effectiveness or inefficiency. Their value lies in their pedagogical significance. They suggest that students need structured methods for evaluating AI use beyond convenience, novelty, or surface-level productivity. They also support the central conceptual claim of this paper: AI literacy in business and professional education must include opportunity-cost literacy, the ability to measure and reflect on what AI saves, what it adds, what it displaces, and what human capacities must remain central.

### Operationalization Chart

<b>Phase</b>	<b>Student Task</b>	<b>Evidence Collected</b>	<b>Opportunity-Cost Question</b>	<b>Example from Exploratory Applications</b>
1. Task Selection	Choose a business, communication, or production task where AI might appear useful.	Task description and expected outcome.	Is this a task where AI may genuinely add value, or is AI being used because it is available?	Students selected PR news releases, reaction statements, and documentary spin scripting tasks.
2. AI Use Hypothesis	Predict whether AI will save time, improve quality, or create hidden costs.	Short pre-use hypothesis.	What do students expect AI to save, and what might it displace?	Students expected AI to help structure PR drafts and generate documentary script options quickly.
3. Workflow Comparison	Complete or evaluate the task through human-led, AI-assisted, hybrid, or agentic-AI-delegated workflows.	Drafts, prompt logs, time logs, revision notes, workflow notes.	Does AI actually improve the workflow after revision, verification, and process review are included?	In the documentary scripting task, some students spent close to an hour prompting ChatGPT and Gemini and concluded they could have written a stronger script themselves.
4. Cost Logging	Record prompting time, revision time, errors, loss of nuance, frustration, and	Cost log or reflection sheet.	What hidden labor did AI create?	PR students found that AI helped with structure but required multiple revisions for tone, context, and organizational voice.

<b>Phase</b>	<b>Student Task</b>	<b>Evidence Collected</b>	<b>Opportunity-Cost Question</b>	<b>Example from Exploratory Applications</b>
5. Agentic AI Delegation Review	confidence level. If using an AI agent, identify autonomous steps, tool calls, intermediate choices, and assumptions.	Delegation map, process notes, verification checklist.	What did the AI agent decide, select, sequence, or obscure along the way?	Future applications may ask students to compare AI-assisted drafting with agentic workflows that gather background, draft materials, and recommend next steps.
6. Quality and Judgment Review	Evaluate final output against professional standards.	Rubric-based or reflective evaluation.	Did AI improve quality, or did it produce a polished but less contextually responsible draft?	Students were less confident submitting final PR copy produced mainly by an LLM because it lacked specificity and professional judgment.
7. Stakeholder and Ethical Analysis	Identify who is affected by the output and what responsibilities remain human.	Stakeholder notes and ethical reflection.	Who bears the risk if the AI output is inaccurate, generic, misleading, or tone-deaf?	In PR writing, students recognized that credibility, accuracy, audience sensitivity, and organizational trust remain the writer's responsibility.
8. Strategic Discernment Memo	Recommend whether, when, and how AI should be used for the task.	Final memo.	Under what conditions is AI appropriate, and what human capacities must remain central?	Students concluded that AI may support ideation or structure, but human judgment must lead when context, voice, narrative coherence, and trust matter.
9. Formative Assessment	Reflect on what the exercise revealed about AI, skill, judgment, and professional formation.	Reflection or class discussion.	What kind of professional habit does this AI use form?	Students learned to ask whether AI saved time, shifted labor, weakened skill practice, or supported better professional decision-making.

## **Applying the Lab**

The operationalization chart summarizes the full sequence of the AI Opportunity-Cost Decision Lab. Faculty can adapt the lab to a single class session, a multi-week assignment, or a course project. The essential movement is simple: students select a professional task, make an AI-use hypothesis, compare workflows, log visible and hidden costs, assess stakeholder implications, and produce a strategic discernment memo.

The lab works best when students complete at least two workflows: one AI-assisted and one human-led or hybrid. Where agentic AI is available and appropriate, faculty can add a fourth workflow: AI-delegated work. Students then compare time spent, revision burden, factual accuracy, contextual fit, tone, stakeholder sensitivity, confidence in the final product, process visibility, delegation depth, verification burden, accountability clarity, and the degree to which human judgment remained central. This comparison prevents AI use from being evaluated by speed alone.

The strategic discernment memo is the culminating artifact. It asks students to answer four questions: Is AI appropriate for this task? Under what conditions? What safeguards are necessary? What human capacities must remain central? In agentic AI workflows, students add two further questions: What parts of the process did the AI agent perform autonomously, and can the human user explain and defend those steps? These questions require students to move from tool use to leadership judgment. The final recommendation may endorse AI use, reject it, or propose a hybrid workflow.

Students can then develop personal or organizational heuristics for future AI use in those kinds of tasks. These heuristics should remain provisional rather than fixed. When a new version

of a relevant tool is released, users or organizations can retest the original judgment and revise the guideline as needed. Once formalized, such heuristics provide clearer guidance for tool selection and responsible use.

For example, if a task requires empathy, users may decide to begin with human understanding before turning to AI drafting. If the output affects a person's opportunity, dignity, employment, reputation, or public trust, they may require human deliberation and documented accountability. If prompting, correction, and verification take more time than human drafting, they may conclude that AI has not saved time. If an AI agent completes a workflow whose intermediate steps cannot be explained or audited, they may require process review before deployment. If context, voice, narrative coherence, or trust matter, they may determine that human judgment must lead and AI may only assist.

### **Worked Example: Donor Communication Plan for a Nonprofit**

To illustrate the Decision Lab in a business education context, consider a business communication or nonprofit management course. Students are asked to prepare a one-page donor communication plan for a regional nonprofit seeking to re-engage lapsed donors after a difficult fundraising year.

#### **Task**

Develop a donor communication plan that includes a core message, audience insight, communication channels, sample opening paragraph, and ethical considerations.

#### **Student Hypothesis**

Students predict that AI will help generate a first draft quickly but may produce generic language that lacks organizational voice, emotional specificity, and donor trust. If an AI agent is used, students predict that it may accelerate background research and message planning, but may

also introduce hidden assumptions about donor motivation, organizational identity, or audience segmentation.

### Workflow Comparison

<b>Workflow</b>	<b>Time Spent</b>	<b>Strengths</b>	<b>Hidden Costs</b>	<b>Quality Judgment</b>
Human-led draft	45 minutes	More specific, better tone, stronger empathy	Slower; required more initial struggle	Strong relational authenticity
AI-assisted draft	18 minutes prompting + 22 minutes revision	Fast structure; useful wording options	Generic voice; needed fact-checking; overused emotional language	Efficient but required heavy human correction
Hybrid workflow	15 minutes human outline + 12 minutes AI expansion + 18 minutes revision	Balanced structure and human specificity	Still required careful editing and mission alignment	Best overall result
Agentic-AI-delegated workflow	10 minutes setup + 25 minutes review and verification	Could generate donor segments, draft message options, suggest channels, and create next steps	Hidden assumptions; unclear source selection; required careful verification of audience logic and mission fit	Potentially useful for planning support, but not acceptable without human review

### Cost Log

The AI-assisted workflow appeared faster at first, but revision time reduced the advantage. Students noted that AI produced polished but generic donor language. It used phrases such as "your generosity changes lives" without evidence of the nonprofit's actual impact. The human-led workflow produced more authentic language because students had to think about donor trust, organizational mission, and concrete examples. The hybrid workflow worked best because students first clarified the organization's mission, donor concern, and communication purpose before using AI to expand possible wording.

The agentic-AI-delegated workflow adds another layer of analysis. If an AI agent generates audience segments, drafts a communication calendar, suggests donor appeals, and recommends next steps, students must evaluate both the output and the process. They must ask which assumptions about donors guided the segmentation, which communication channels were prioritized, whether the recommendations reflect the nonprofit's mission and capacity, and whether the human team can explain and defend the strategy. The opportunity cost may include reduced understanding of the donor relationship, diminished mission-specific judgment, and increased dependence on an automated planning process.

### **Stakeholder and Mission Analysis**

Students identify three stakeholder groups: lapsed donors, nonprofit staff, and beneficiaries. They conclude that donor communication requires trust and authenticity because donors may feel skeptical after a difficult fundraising year. AI can assist with structure and language variation, and agentic AI may assist with planning and sequencing. However, the core message must come from human knowledge of the organization's mission and actual relationships.

### **Strategic Discernment Memo**

**Recommendation:** Use AI only after the human team has developed the core message, donor insight, and mission-specific examples. AI may assist with alternative subject lines, readability, tone refinement, and communication sequencing, but final language and strategy must be reviewed by staff who know the donor community. The nonprofit must not use AI to fabricate emotional stories, exaggerate impact, automate relational trust, or replace genuine donor relationship-building.

Human capacities that must remain central: empathy, mission knowledge, donor trust, truthfulness, audience understanding, and responsibility for the organization's voice.

### **Learning Outcome**

Students learn that AI can help with efficiency but cannot substitute for relational judgment. The task teaches them to see donor communication not merely as content generation but as an ethical relationship with stakeholders. Agentic AI also shows that delegation carries its own opportunity costs. The more AI systems can do, the more students must ask what humans still need to know, practice, verify, understand, and own.

### **Broader Application Beyond Jesuit Institutions**

Although this framework is informed in part by Ignatian pedagogy, it is adaptable beyond Jesuit institutions. Other business schools, public universities, and nonsectarian professional programs can use opportunity-cost literacy as a human-centered decision-making framework. The language of Catholic Social Teaching can be translated into widely shared educational concerns: human dignity becomes respect for persons; the common good becomes stakeholder well-being and public value; subsidiarity becomes participatory decision-making; solidarity becomes attention to those affected by organizational choices; integral human development becomes holistic professional and civic formation.

In this broader context, opportunity-cost literacy can support business ethics, responsible management education, stakeholder theory, sustainability education, and leadership development. The framework helps students ask whether AI adoption strengthens or weakens judgment, accountability, trust, and organizational responsibility. Thus, while the model is deeply compatible with Jesuit education, it also offers a transferable approach for any business school concerned with responsible AI use.

The framework may be especially useful as agentic AI becomes more common in organizations. Any business school concerned with responsible management education will need to teach students how to evaluate delegation to AI systems. Opportunity-cost literacy provides a way to ask whether agentic AI improves organizational decision-making or hides complex assumptions inside automated workflows.

### **Implications for Business Leadership Formation**

Opportunity-cost literacy contributes to business education in four ways. First, it strengthens strategic judgment. Students learn that AI adoption is not a matter of novelty or convenience. It is a strategic decision involving alternatives, trade-offs, stakeholder effects, and long-term organizational consequences.

Second, it supports ethical formation. Students do not merely learn AI rules; they practice discernment. They learn to recognize hidden costs, weigh competing goods, and defend decisions in light of human dignity and stakeholder responsibility.

Third, it deepens professional responsibility. Students learn that AI does not eliminate human accountability. Even when AI systems generate recommendations, leaders remain responsible for deployment, interpretation, communication, oversight, and correction. This is especially important for agentic AI, where the system may perform intermediate steps that are not immediately visible unless the user is trained to examine them.

Fourth, it aligns AI education with the broader purpose of business education. Students are not only being prepared to use tools; they are being formed to lead organizations. Opportunity-cost literacy helps ensure that technological competence does not come at the expense of human judgment.

For Jesuit business education specifically, the framework translates Catholic Social Teaching and Ignatian pedagogy into an AI-era practice of leadership formation. It asks students to consider whether technological efficiency remains ordered toward human dignity, integral development, and the common good. For broader business education, it offers a practical way to teach responsible AI decision-making without reducing ethics to compliance or AI literacy to tool use.

### **Conclusion**

Generative and agentic AI promise speed, scale, efficiency, and increasingly sophisticated forms of delegated workflow. These benefits matter, but they do not exhaust the ethical and strategic questions business leaders must ask. Business education has a responsibility to form leaders who can evaluate not only what AI saves, but what it costs.

Opportunity-cost literacy offers a practical framework for that task. By teaching students to compare workflows, log hidden costs, assess stakeholder implications, examine delegation depth, and produce strategic discernment memos, business schools can move AI education beyond tool use toward responsible judgment. The exploratory classroom applications discussed in this paper should be read as pedagogical illustrations, not empirical proof of AI's effectiveness or inefficiency. Their value lies in showing how students can test AI's claims rather than assume them.

Agentic AI does not make opportunity-cost literacy obsolete. It clarifies its importance. Generative AI raises the question of assistance. Agentic AI raises the question of delegation. The more AI systems can do, the more business education must teach students to ask what humans still need to know, practice, verify, understand, and own.

For Jesuit business education, this framework is mission-aligned because it grounds AI decision-making in human dignity, integral development, discernment, and the common good. For

broader business education, it offers a transferable model for preparing students to use AI without surrendering the human capacities that ethical leadership requires. The central educational task is not simply to teach students how to use AI, but to form leaders capable of judging when AI use serves persons, organizations, and the common good--and when its hidden costs outweigh its visible savings.

## References

- Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems.
- Aristotle. (1999). *Nicomachean ethics* (T. Irwin, Trans.). Hackett.
- Becker, G. S. (1965). A theory of the allocation of time. *The Economic Journal*, 75(299), 493–517.
- Benedict XVI. (2009). *Caritas in veritate*. Libreria Editrice Vaticana.
- Crawford, M. B. (2009). *Shop class as soulcraft: An inquiry into the value of work*. Penguin Press.
- Emelu, M. N. (2025). Gen Z and Gen Alpha: Designing theology courses for a digital and AI generation. *Jesuit Higher Education*, 14(1), 49–66.
- Emelu, M. N. (2026a). Formation-based AI ethics: Ethical formation, responsibility, and opportunity cost in AI ecosystems. *AI and Ethics*, 6, 282.
- Emelu, M. N. (2026b). Formation at the frontier: Pope Leo XIV discourse on AI and digital media. *Journal of Media and Religion*. Advance online publication.
- Floridi, L. (2014). *The fourth revolution: How the infosphere is reshaping human reality*. Oxford University Press.
- Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1).
- Francis. (2015). *Laudato si': On care for our common home*. Libreria Editrice Vaticana.
- Francis. (2020). *Fratelli tutti: On fraternity and social friendship*. Libreria Editrice Vaticana.

Francis. (2024). Address at the G7 Session on Artificial Intelligence, Borgo Egnazia, Italy, June 14, 2024. Libreria Editrice Vaticana.

International Commission on the Apostolate of Jesuit Education. (1993). *Ignatian pedagogy: A practical approach*. Jesuit Secondary Education Association.

Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence, 1*, 389–399.

John Paul II. (1981). *Laborem exercens*. Libreria Editrice Vaticana.

Kahneman, D. (1973). *Attention and effort*. Prentice-Hall.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*(2), 263–291.

Leo XIII. (1891). *Rerum novarum*. Libreria Editrice Vaticana.

Leo XIV. (2026). *Magnifica humanitas: On safeguarding the human person in the time of artificial intelligence*. Libreria Editrice Vaticana.

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.

MacIntyre, A. (2007). *After virtue* (3rd ed.). University of Notre Dame Press.

Mayer, R. E. (2020). *Multimedia learning* (3rd ed.). Cambridge University Press.

McLuhan, M. (1964). *Understanding media: The extensions of man*. McGraw-Hill.

Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence, 1*, 501–507.

Ng, D. T. K. (2021). What is AI literacy? A review of AI literacy definitions, competencies and pedagogical approaches. *Computers and Education: Artificial Intelligence, 2*, 100041.

- Paas, F., Renkl, A., & Sweller, J. (2003). Cognitive load theory and instructional design. *Educational Psychologist*, 38(1), 1–4.
- Paul VI. (1967). *Populorum progressio*. Libreria Editrice Vaticana.
- Postman, N. (1992). *Technopoly: The surrender of culture to technology*. Knopf.
- Sennett, R. (2008). *The craftsman*. Yale University Press.
- Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2), 257–285.
- Vallor, S. (2016). *Technology and the virtues: A philosophical guide to a future worth wanting*. Oxford University Press.