THINK WITH AI
IN HEALTHCARE FINANCE

Becoming a More Capable Healthcare Financial Decision Maker

Mike Rockwood & Yvonne Rockwood

THINK WITH AI • QUANTUM CLASSROOM • 2026

So...what are you actually supposed to be learning here?

Healthcare finance, obviously. You'll learn financial statements, costs, revenue, ratios, staffing, break-even analysis, budgeting, reimbursement, capital decisions and a fair amount of terminology that might not exactly make you the life of the party.

And yes, you need to know this stuff. If you're going to lead a healthcare organization, you need enough financial knowledge to understand what your finance people are telling you, ask intelligent questions, recognize when something doesn't look right, and hold people accountable for both the accuracy of the numbers and the interpretation of them.

But five years from now, nobody is going to pay you because you can calculate a current ratio. A spreadsheet can do that. AI can do it before you've finished typing the question. Your value will come from knowing which numbers matter, what they mean, what might be wrong with them, what else they affect, and what your organization should do next.

> **That's a very different capability. We call it judgment. And judgment has always been the real point.**

Judgment Has Always Been the Goal

AI didn't suddenly make judgment important. Healthcare leaders have always had to make decisions with incomplete information, competing priorities, financial constraints, operational realities and consequences for real people. Should we add the service line? Can we afford those additional nurses? Why did labor expense jump? Is this program really losing money, or are we allocating costs badly? Should we buy the equipment? What happens if volume doesn't materialize? And perhaps the most dangerous question in management: Are we sure?

Those aren't accounting questions alone. They're management questions that use financial information. That's why this course doesn't stop at calculation. You may calculate working capital, current ratio and debt-to-assets, but then you'll be asked to interpret liquidity risk and leverage posture and decide what management should do about it.

> **The calculation is necessary. The interpretation is where things get interesting. The judgment is where you earn your paycheck.**

Then Along Came AI

AI can calculate, explain a balance sheet, analyze a variance, build a break-even model, compare alternatives, suggest staffing assumptions and draft a recommendation to your CFO. It can do many of those things astonishingly well. Great. We're going to use that.

This is not a course built around pretending AI doesn't exist. Nor are we going to spend eight weeks playing a game of “Did the Student Use ChatGPT?” There's a much more interesting opportunity.

For years, educators could make a reasonable assumption: if you did the work, you probably did much of the thinking required to produce the work. The assignment itself created the practice. A difficult case made you retrieve what you knew, transfer it to a new situation, calculate, interpret, question, compare, connect and decide.

AI weakens that old connection because it can now perform much of the work that once caused the thinking to happen. Our response isn't to abandon the traditional idea of learning through application. It is to preserve what is valuable about doing the work while adding AI in a way that can make both the work and the thinker better.

> **Keep the work. Keep the thinking. Add AI to make both better.**

The Path to Judgment

Good healthcare decisions rarely come from one calculation or one flash of insight. Experienced leaders work their way toward judgment. They try to understand the situation, examine what is really going on, weigh alternatives and consequences, and then decide what they believe should be done.

UNDERSTAND — What do I know about the situation?

Get the facts, numbers, context, constraints and relevant concepts straight. Identify what you know, what you think you know, and what you still need to know.

EXAMINE — What is really going on?

Work the problem. Calculate, compare, classify, interpret, trace causes and question assumptions. Ask whether the apparent problem is actually the problem.

WEIGH — What are the alternatives, tradeoffs and consequences?

Consider competing explanations, options, risks, uncertainty, organizational consequences and future implications. Decide what matters most and what you are willing to accept.

JUDGMENT — What do I conclude, recommend or do...and can I defend why?

Eventually somebody has to make the call. The answer may be Go, Don't go, Not yet, or I need more information before I'm willing to recommend anything.

The path is shown in a tidy order because diagrams have that luxury. Real decisions don't. New information can send you backward, a challenged assumption can make you examine the problem again, and a consequence you hadn't considered can change what you're willing to recommend.

The Path to Judgment™ framework developed by Mike Rockwood and Yvonne Rockwood for this booklet (2026).

The Thinking the Work Develops

Healthcare finance develops thinking partly through the work itself. You calculate a break-even point, diagnose a variance, question a staffing assumption, compare alternatives, consider what else a decision might affect and eventually recommend what management should do. We don't need to interrupt that process with a long taxonomy lesson. We do, however, want you to notice some of the capabilities the work is asking you to exercise.

QUANTITATIVE — What do the numbers tell me?

ANALYTICAL — What explains what I'm seeing?

SYSTEMS — What else does this affect?

STRATEGIC — Given where we're trying to go, what choices best position us for the future?

You don't need to memorize four definitions or treat them as separate processes. You'll develop these capabilities the way people have always developed much of their thinking: by using them. From time to time, we'll simply stop and notice which ones the problem is asking you to use.

That becomes especially important with AI. If AI can perform some of the work that once caused you to exercise these capabilities yourself, we need to be deliberate about how it enters the work. That is where the Thinking Loop comes in.

The Thinking Loop

AI is extraordinarily good at answering questions. Ask it for a break-even calculation, an explanation of a variance, possible causes of a staffing problem or alternatives for reducing costs, and it may give you something very useful.

But notice what has happened. AI has given you a piece of the answer, perhaps a very good piece. You still have to decide what that information means, how it connects with everything else you know, whether you trust it and what you should do because of it. In that sequence, AI is primarily supplying pieces and you are connecting the dots afterward.

There's nothing wrong with that. Sometimes that's exactly what you need. But when the problem requires judgment, there is another way to work: begin connecting the dots yourself and then bring AI into the reasoning process while it is happening.

We use a simple process: FRAME → ANCHOR → TEST → REFINE. And throughout the Loop, we follow one practice:

> **THINK FIRST. THEN THINK WITH AI.**

Why Think First?

Not because there's virtue in making the work harder, and certainly not because using AI is somehow cheating. We want you to use AI - a lot. The reason to think first is that the first attempt at thinking is itself valuable practice.

Before AI begins suggesting how to frame the problem, what facts matter, what conclusions are reasonable or how your recommendation might be improved, your brain has to retrieve what it already knows, transfer that knowledge to a new situation, generate possibilities, make connections, encounter uncertainty and form an initial position. That little bit of cognitive friction is not wasted effort. It is part of how capability develops.

Then bring in AI. Now the experience changes. Instead of simply receiving AI's way of seeing the problem, you can compare it with your own. AI may challenge an assumption, show you something you missed, offer another explanation, test the numbers or expose a consequence you hadn't considered. You have to decide what to accept, reject, investigate or change.

> **Your first attempt gives you one opportunity to exercise the capability. Your interaction with AI gives you another, richer one because somebody is pushing back.**

Why This Can Produce Better Work Too

Think First. Then Think With AI. is not a tradeoff in which we sacrifice some of AI's power for the sake of exercising your brain. Done well, you get the benefit of both.

You begin the reasoning. AI challenges it. You compare. You reconsider. You revise. AI challenges again. That back-and-forth can uncover weak assumptions, alternative explanations, missing information and consequences that neither your first attempt nor AI's first answer surfaced on its own.

The result should be a stronger basis for judgment. And something else has happened along the way: you've practiced the thinking required to reach that judgment.

> **You improve the answer while developing the person who has to judge the answer.**

AI as a Capability Accelerator

This is why we describe AI as a potential capability accelerator. Used primarily as an answer machine, AI can save cognitive effort and make you more productive. Used as a thinking partner, it can create a rapid cycle of attempt, challenge, comparison, reconsideration and refinement.

In some ways, it begins to resemble having a remarkably patient personal mentor sitting beside you. You try something. Your mentor questions it. You defend it. The mentor points out something you missed. You reconsider. The mentor pushes again. Eventually your reasoning is stronger, and so are you for having gone through the process.

No, AI is not Aristotle. It can be wrong, overconfident and occasionally ridiculous. But something once scarce - an interactive thinking partner available whenever you want to test an idea - is becoming widely available. Eat your heart out, Aristotle.

> **That gives us a win-win worth pursuing: a better-supported basis for judgment today and greater human capability tomorrow.**

How the Loop Works

The four steps are simple. What matters is the rhythm inside them: think first, then think with AI. You do not have to complete the problem alone before AI is allowed into the room. You simply give your own thinking a chance to begin before AI starts shaping it.

FRAME — What is the real question or decision?

Think first: generate possible ways of framing the situation and choose your initial frame. Then Think With AI: ask AI to challenge that frame, identify what may be missing and suggest another way to see the problem.

ANCHOR — What do I know and trust?

Think first: identify the facts, assumptions, evidence and missing information as you understand them. Then Think With AI: ask AI to test your grounding, identify gaps and question assumptions that deserve more scrutiny.

TEST — How might I be wrong?

Think first: develop an initial interpretation or position worth testing. Then Think With AI: ask AI to attack it, test assumptions, offer competing explanations and follow consequences you may have missed.

REFINE — What should change?

Think first: decide which challenges you accept, which you reject and how you would revise your position. Then Think With AI: ask AI to test the revised reasoning and identify what remains weak, uncertain or difficult to defend.

Then comes judgment. The Loop can continue as long as it is improving your thinking, but eventually somebody has to decide what should be done and why. AI can still challenge the final recommendation or play the skeptical board member. The judgment remains yours to make, explain and defend.

Now Let's Try It: The Speer Street Scandal

Mile High Behavioral Health announces a shiny new “robotic pharmacy” and 24/7 crisis intake. Unfortunately, the robot turns out to be essentially a repurposed cabinet with a leased tablet. Someone has booked it as a $480,000 capital asset and started depreciating it over seven years.

The supposed 24/7 crisis operation drops back to 12 hours a day almost immediately, but the break-even presentation is still using the 24/7 volume assumptions. Pharmacy inventory is disappearing. Staffing reports don't match actual schedules. Stakeholders are asking questions. Other than that, things are going beautifully.

This is exactly the kind of case in which the spreadsheet alone won't save you. Let's work it through the Loop and notice the thinking capabilities as they appear.

FRAME

Think First. Before you open AI, decide what you think the real decision is. The crisis program isn't meeting its projections, but is that really the problem? Perhaps the original business case was flawed. Perhaps 24/7 operation never made sense. Perhaps the operating model is wrong. Or perhaps the program is strategically important enough that Mile High should tolerate losses while it develops. Generate a few possibilities, then choose the framing that seems most useful.

> *Should Mile High continue investing in the crisis program, and if so, under what operating model and conditions?*

Notice the thinking. You're using analytical thinking to distinguish the apparent problem from the underlying problem and strategic thinking to ask what the organization is ultimately trying to accomplish. The case is teaching the capability because you are doing the work.

Then Think With AI.

> *I've framed the decision this way: Should Mile High continue investing in the crisis program, and if so, under what operating model and conditions? What important dimension might my framing be missing? Don't answer the question yet. Challenge how I've framed it.*

Now AI is inside the reasoning. It didn't replace your first attempt; it is testing it.

ANCHOR

Think First. Go back through the case and separate what you actually know from what you're assuming and what you still need to know. You know the original model assumed 24/7 operation, actual operations fell to 12 hours, the break-even analysis continued to use the original assumptions, inventory is disappearing and staffing reports don't match actual schedules.

Be careful about what you infer. Missing inventory doesn't automatically tell you why it's missing. A bad forecast doesn't tell you what the correct forecast should have been. The leased pharmacy equipment raises an accounting question, but you still need to understand the arrangement before deciding how it should be treated.

Here quantitative and analytical thinking become especially visible. You may need to distinguish fixed, variable and step-fixed costs, recast the break-even model, test volume assumptions, examine staffing against operating hours and decide which numbers are trustworthy enough to use.

Then Think With AI.

> *Here are the facts I believe are established, the assumptions I think we're making, and the information I think is missing. Check my classification against the case. What have I misclassified or overlooked? Don't fill missing information with assumptions.*

You're not asking AI to establish reality for you. You're asking it to test the ground you've already established.

TEST

Think First. By now you should be leaning somewhere. You don't need a polished recommendation, but you do need a position worth testing. Suppose your initial thinking is:

> *I think Mile High should continue the crisis program, but operate it 12 hours per day until volume justifies expansion.*

Write down why you think that. What evidence supports you? What assumptions are you making? What bothers you about your own conclusion? Forming a tentative position requires you to bring the pieces together before anyone tells you what to think.

Then Think With AI.

> *Assume my recommendation is wrong. Build the strongest case against it. Challenge it quantitatively, analytically, systemically and strategically. Which assumption, if wrong, would most seriously undermine my recommendation?*

This is where the four capabilities become useful as instructions rather than definitions. Test the numbers quantitatively. Challenge the explanation analytically. Follow the consequences systemically. Ask whether the recommendation moves the organization in the right direction strategically. Let AI disagree with you. A thinking partner who only tells you you're brilliant isn't much of a thinking partner.

REFINE

Think First. When AI comes back with its challenge, resist the temptation to ask, “Okay, what's the better answer?” Decide which objections are legitimate, which aren't supported by the case, what AI noticed that you didn't, and whether anything actually changed your mind. Then revise the recommendation yourself.

> *Continue the crisis program at 12 hours per day for now. Rebuild the financial model using realistic volume and staffing assumptions, resolve the inventory and accounting problems, and establish specific performance thresholds before considering expansion to 24/7 operations.*

That's not quite where you started. Good. Your thinking moved. You may have used all four capabilities again without stopping to label every move.

Then Think With AI.

> *Here is my revised recommendation. Here are the challenges I accepted, the ones I rejected, and why. What weakness remains in my reasoning? What should I be prepared to qualify or defend before presenting this recommendation to the board?*

AI gets another turn. So do you. That's why we call it a loop.

JUDGMENT

Eventually somebody has to stop prompting and make the call. What do you recommend? Why? Which assumptions are you willing to defend? What tradeoffs are you accepting? What are you still uncertain about? What evidence could cause you to change your mind?

AI doesn't have to disappear. Ask it to play the role of a skeptical board member and give you the three hardest questions you're likely to face. That's excellent preparation. But you answer the questions. Eventually the board isn't asking AI. It's asking you.

> **That's judgment.**

You'll Practice This Throughout the Course

This won't be something Professor Rockwood explains in a booklet and then expects you to remember on your own. Think First. Then Think With AI. will be built into lessons, assignments, discussions and assessments throughout the course.

Sometimes you'll be asked to frame a problem before consulting AI. Sometimes you'll make an initial calculation or interpretation and then use AI to challenge it. You may capture what you thought before using AI and explain what changed afterward. In a discussion, you may need to defend why you accepted one AI suggestion and rejected another. And sometimes you'll demonstrate the capability without AI at all.

The purpose isn't to catch you using AI. Quite the opposite. It is to give you repeated practice using AI in a way that develops your own capability while improving the quality of the work you produce.

The professor still designs the problem, decides what capabilities the work should exercise, structures how AI enters, asks you to make your thinking visible and evaluates the judgment you ultimately defend. AI becomes part of the learning environment, not a substitute for it.

Make the Thinking Visible

Your professor can't see inside your head, so the course will sometimes ask you to make the process visible. You may calculate, compare, trace, connect, challenge, generate alternatives, defend a conclusion and reconsider when the evidence changes. These are not another list to memorize. They are simply observable signs that thinking is happening.

Your discussion boards already do this. You take a position, engage another person's reasoning and then return to your own position. Along the way you may ASK, ADD, CHALLENGE, CONNECT, ALTER or EXTEND an idea. The point isn't to perform a discussion ritual. It is to make the movement of your thinking visible.

> **Don't show me what AI knows. Show me what you did with what AI knows.**

Make Your AI Use Visible Too

When AI is permitted in submitted work, include the disclosure your instructor or assignment requires. A useful disclosure identifies the tool, explains how it participated in your work, and makes clear what you verified, changed, rejected or decided yourself. A few sentences are usually enough unless the assignment requires something different.

> *AI Use Disclosure - I used ChatGPT to challenge my initial interpretation of the staffing variance and identify assumptions I may have overlooked. I checked its suggestions against the case information, rejected one suggestion that was not supported by the data, and revised part of my recommendation. I verified the calculations and sources used in my final submission. The final analysis and judgment are my own.*

The more useful question isn't simply, “Did you use AI?” It is: What did you do with it? Did you question it? Check it? Improve on it? Reject something it suggested? Did it change your thinking? Can you explain and defend the final recommendation?

Your Voice Still Matters

Professional life is not one long take-home paper. Someone asks, “Why are you recommending this?” and they're looking at you. A physician challenges your numbers. A board member asks what happens if reimbursement drops. Your CEO says, “Give me the bottom line.”

Use AI to prepare. Ask it what questions you might get. Let it play the skeptical executive. Then close the laptop and explain what you think in your own words. If AI disappeared five minutes before the meeting, you should still understand what you're presenting and be able to defend it.

> **That's not anti-AI. It's pro-capability.**

One Last Thought

AI can make you faster. It can make information easier to obtain. It can calculate, explain, analyze, compare, model, generate, question and challenge. And it's going to get better. Good.

The opportunity is bigger than productivity. Used thoughtfully, AI can help you reach a better-supported judgment while giving you repeated opportunities to practice the capabilities that judgment requires. You get better work today and, if you use the process deliberately, a more capable thinker tomorrow.

So think first. Retrieve what you know. Generate an idea. Make a connection. Take an initial position. Then Think With AI. Let it challenge you, surprise you, expose an assumption, offer another possibility and make you reconsider something you were sure about five minutes ago. Then decide what you think.

> **Because the real product isn't the spreadsheet, the analysis, the presentation or even the answer. It's the capability you develop while producing it.**

Notes & References

This booklet grew from Yvonne Rockwood's work designing and teaching healthcare finance and Mike Rockwood's work on human thinking, judgment and AI developed in Think With AI. Together, those perspectives shaped the approach presented here. The Path to Judgment™ and the four thinking capabilities as applied to healthcare finance are practical teaching frameworks developed for this booklet; they are not presented as a new theory of cognition.

1. Foundational framework.

Rockwood, M. (2026). Think With AI. The Thinking Loop and the broader argument for using AI to strengthen rather than replace human thinking originate in this work.

2. Learner agency and cognitive offloading.

Wang, G., Wang, W., Yang, D., & Ren, J. (2026). Generative AI, cognitive offloading, and learner agency in higher education: A scoping review. Behavioral Sciences, 16(7), 1150. https://doi.org/10.3390/bs16071150

3. Critical autonomy and uncritical AI use.

Suazo Galdames, I. C., Santiesteban Velázquez, M., Saracostti, M., & Chaple Gil, A. M. (2026). Between cognitive offloading and critical autonomy: A systematic review of the epistemic implications of generative AI in higher education. Frontiers in Education, 11, 1878665. https://doi.org/10.3389/feduc.2026.1878665

4. Scaffolding critical thinking with AI.

Vendrell, M., & Johnston, S.-K. (2026). Scaffolding critical thinking with generative AI: Design principles for integrating large language models in higher education. Computers and Education: Artificial Intelligence, 10, 100572. https://doi.org/10.1016/j.caeai.2026.100572

5. Learning-to-learn in an AI-rich environment.

Schorr, I., Bardach, L., Bühler, B., & Kasneci, E. (2026). Learning-to-learn in the age of generative AI: A scoping review and conceptual framework. Computers and Education: Artificial Intelligence, 10, 100575. https://doi.org/10.1016/j.caeai.2026.100575

6. AI and critical-thinking skills.

Guo, Y., Huang, L., Zhang, C., Li, Q., & Chen, M. (2026). ChatGPT in education: A systematic review of its impact on critical thinking skills and dispositions. Thinking Skills and Creativity, 60, 102106. https://doi.org/10.1016/j.tsc.2025.102106

MSU Denver AI Guidance

The approach to generative AI presented in this booklet is aligned with MSU Denver's current guidance for responsible student use of AI and AI-aware assessment. Course and assignment directions govern each task. Students remain responsible for submitted work, should verify AI-generated facts, calculations and sources, and should not enter confidential, proprietary or patient-identifiable information into AI tools.

MSU Denver. Student Resources: AI Use Guidance. https://www.msudenver.edu/ai/student-resources/

MSU Denver. AI and Assessment: Guidance for MSU Denver Faculty. https://www.msudenver.edu/ai/ai-and-assessment/