


AI and the Future of Work

What It Means for Laredo – and for You

- Jack Shaw
- AI Transformation Advisor and Keynote Speaker
- Author of The AI Imperative Series
- Texas A&M International University
- September 16, 2026



Thank you, Dr. Clarke. And thank you to Amy Palacios and everyone at the Center for the Study of Western Hemispheric Trade.

It's a privilege to open this year's series during President Maynard's inauguration week. Dr. Maynard, since you're about to be inaugurated, I thought I'd offer a little advice.


I was lucky enough to attend Yale under one of its great presidents, Kingman Brewster. As the story goes, when one of his colleagues left to lead another university, that colleague asked Kingman for the one piece of advice that mattered most. Kingman told him: take care of your A students and your B students, because some of them will come back as your best professors. But take even better care of your C students. Because someday, one of them will come back and build you a fifty-million-dollar science laboratory.

Now, let me tell you what this talk is not. It's not a session on how to use AI tools. Most of you are already better at that than most of the executives I typically brief.

What I want to talk about instead is the organizations you're about to work for, and what's happening inside them right now. That's the part that will shape your career.

This morning at the bridge

- Trucks queued before dawn on the bridge
- Manifests, inspections, brokers, drivers, waiting
- Every crossing is a chain of decisions
- A wrong one holds a trailer, misses a window



2

This morning at the bridge there were already trucks waiting. Not a few. A line of them.


Stand there ten minutes and you see the shape of it. A driver checking paperwork against a phone. A broker three miles away deciding which lane to file under. An inspector deciding which trailer to open. A dispatcher deciding whether the load still makes its window.

And when one of those decisions goes wrong, it isn't abstract. A trailer sits. A load misses its window. A customer promised Thursday doesn't get Thursday, and takes their business elsewhere.

From a distance it looks like traffic. Up close it's a chain of decisions, thousands every day, made from incomplete information.

The busiest land border crossing

- Laredo is the nation's largest inland port
- 2,945,388 trucks entered from Mexico in 2025
- Nearly one in four entering the United States
- Two national economies meet at this point
- You live at the center, not the edge



*Bureau of Transportation Statistics, Border Crossing Data
Annual Release, February 2026. CBP inbound counts.*

3

Most people don't know that Laredo is the nation's largest inland port. The busiest land border crossing in the United States. Not one of the busiest. The busiest.

Last year, 2,945,388 trucks entered the United States through Laredo from Mexico. Nearly one in four of every truck that entered this country by land, from either border, came through here.

On value, about \$354 billion moved through this port in 2025. Third among all United States ports of entry, and first among every land port in the entire Western Hemisphere.

Which means you're not studying at the edge of the American economy, looking in. You're studying at one of its hinges.

All of it runs on decisions

- Every truck, container, invoice, inspection, and loan
- Behind each one, someone decided something
- AI is coming for the decisions, not the paperwork
- That is what changes everything else



4

Every truck, every container, every invoice, every inspection, every loan behind that scene at the bridge represents a series of decisions that somebody makes. And artificial intelligence is coming for those decisions.

Not just for the paperwork. The paperwork story is a productivity story. Real, and boring.

The decision story is the one that changes what an organization is and what it needs people for.

Everything I say for the next forty minutes is about that – decisions.

You already use AI

- You use these tools for coursework already
- Research, drafting, troubleshooting, revision, study
- No demonstration and no prompting tutorial today
- You would be correcting me inside five minutes



5

Most of you already use AI tools. Every day. For reading you didn't finish, for a first draft you couldn't start, for code that wouldn't compile, for checking your own reasoning at one in the morning.

So there's no demonstration today and no prompting tutorial. If I showed you how to write a prompt, several of you would likely be correcting me inside five minutes!

The real question

- You know what AI does for you
- The harder question:
What happens when employers use it
- We are early. Organizations move slower than tools
- That gap is where careers get made
- Five sectors, and I will get to all five



7

So if the question isn't what AI can do for you, what is it?

It's this. What happens when the organizations you're going to work for use AI?

As students, when you use these tools, you get faster. But when a whole organization uses them properly, the organization itself changes shape. And two things change with it. What it can do, and what it needs people for.

Here's the encouraging part. We're still very early. The tools have gotten very good very quickly, but organizations haven't. That gap is why so many companies have spent real money for very little, and why this is a better moment to start a career than it might look like.

So I'm going to describe an AI framework, and then spend the last part of this talk showing how it could apply in five places you might actually work. Trade and logistics. Finance and accounting. Government and the border. Healthcare. And education.

What is Tactical AI

- Tools applied to specific tasks, right now
- Measurable wins inside twelve months
- Drafting, summarizing, classifying, answering, checking
- Useful, fundable, and easy to start
- It does not change the operating model



8

I work with two categories of organizational AI.

The first is Tactical AI. Tactical AI is a tool applied to a specific task. Drafting. Summarizing. Classifying. Answering a routine question. Checking work that somebody else did.

Think of it as a Swiss Army knife. Genuinely useful, always at hand, and it doesn't require you to reorganize your life around it.

It's easy to justify. It pays back inside a year, the finance function can see the number, and nobody has to be brave. That's why almost every organization starts here.

But here's the part to keep in mind. Tactical AI changes nothing about how the organization works. The same people make the same decisions in the same meetings. They just get their inputs faster.

What is Strategic AI

- Changing how the organization actually works
- How information moves and who decides
- Predictive maintenance, dynamic scheduling, integrated planning
- Three to five years, not three months
- This is where advantage compounds



9

The second category is Strategic AI, and it's a completely different animal.


Strategic AI changes how the organization works. Not what one person can produce in an afternoon, but how information moves through the organization and who has the authority to act on it.

Predictive maintenance is a good example. Predicting when a bearing might fail is a task. Redesigning maintenance so the prediction sources the part, schedules the technician, and moves the production plan is an operating model. The first is just a tool. The second creates a different company.

Strategic AI looks ahead three to five years. It's harder to fund, because the payback sits past most executives' planning horizons. But it's also harder to reverse. And it's the only place where advantage compounds.

Two paths, one destination

- Tactical funds the journey and builds capability
- Strategic changes the operating model itself
- Run both. Neither one works alone
- Tactical alone becomes pilot purgatory
- Strategic alone never gets funded



10

Now, the obvious question is which one you should do Tactical, or Strategic. And the answer is both, at the same time, deliberately.

I call it the Dual Path. Tactical funds the journey and builds capability, because visible wins keep the budget alive and teach people how these systems behave. Strategic changes the operating model the journey is for.

Run tactical alone and you get pilot purgatory. Dozens of small successes, none of them connected, an organization measurably busier but not measurably better. I've watched serious companies spend the past three years there.

Run strategic alone and it never gets funded, because you're asking a board to buy a three-year story with no evidence.

So, both. That's the process. Now let me tell you where the process is taking us.

The autonomic nervous system

- Your heartbeat runs without your attention
- So does breathing, and digestion
- That is not the absence of mind
- It is what frees the mind for judgment



11

To explain the destination, I'm borrowing an analogy from biology.


Your heart is beating right now and you're not managing that. Breathing, digestion, body temperature, none of it reaching your conscious attention. That's your autonomic nervous system.

And notice what that buys you. It isn't that your body decided it didn't need a mind. Quite the opposite. Because the routine runs itself, your mind is free for everything that does require you.

Organizations are now building the same capability.

The Autonomic Enterprise™

- AI runs the routine sensing, routing, and response of operations at machine speed.
- The way the autonomic nervous system runs the body, without conscious attention.
- Machines act inside limits the organization sets, and escalate what falls outside them.
- Human judgment is reserved for the decisions that cannot be delegated.
- Autonomous means removing people. Autonomic means protecting the decisions that need them.



12

So, here's the destination.

I call it the Autonomic Enterprise.

AI runs the routine sensing, routing and response of operations at machine speed. The way the autonomic nervous system runs the body, without conscious attention.

The AI acts inside limits the organization sets, and it escalates to people anything that falls outside those limits. And human judgment is also reserved for some decisions that can't be delegated to AI at all.

Now, I want to be very careful about one word here. We hear a great deal about autonomous systems. But autonomous means removing people. Autonomic means protecting the decisions that need people.

Those aren't variations on a theme. They're opposite intentions, and they produce opposite organizations. One of them is just about cutting staff. The other is about transforming for the future.

When the machine calls for help

- A rush order for a product never built
- The system schedules, sources, and prices it
- Then it stops and calls an engineer
- The machine knows what it does not know



13


Now imagine: A rush order arrives for a product variation a plant has never built before. The system reads it, checks capacity, finds materials, prices the job, and schedules the line, in the time it takes to read an email.

Then it stops. And it calls a senior engineer to approve the plan.

Notice the process. Not the speed. The stop. The engineer isn't there because the AI system failed. The engineer is there because the people who built that system decided what it should not decide.

Think, decide, act

- Every operating decision follows one loop
- Take in signals. Judge them. Act
- More than one plan, and the option to switch
- Machines can run that loop continuously
- The question is how far it runs



14

I told you this talk was about decisions, so let me show you how they work when properly designed.


Strip away the vocabulary and every operating decision follows the same loop. Take in signals. Judge what they mean. Act. Find out whether you were right, and go round again. That's what a purchasing manager does, a nurse, a credit officer, an inspector at the bridge.

And you do a version of it every morning yourself. You have more than one route to campus. The expressway and the surface streets. You keep both in mind, watching for anything that might make you switch, and sometimes you switch. That's the difference between a plan and a decision.

A machine can run that loop continuously. So the question was never whether it can. It's how far round the loop it should go without human intervention.

Five stages, and two thresholds

- Sense, assess, reason, respond, escalate
- Two thresholds, not one line
- A purchasing agent, limit ten thousand dollars
- Samantha raises it to fifty thousand
- A person decides what the machine may decide



15

When I was first building AI systems, over twenty years ago, we described that loop in five stages. Sense: collect signals. Assess: is something happening we ought to know about. Reason: work out what's probably true. Respond: act, within the authority granted. And escalate: decide when a person must be told. That fifth stage is a decision about decisions.

Here's what it looks like when a real person sets one.

We have a purchasing agent. Not a person, an AI-enabled system. It watches supply and prices, and when it sees a shortage coming, it buys. It has a limit. Let's say, ten thousand dollars. Above that it stops and asks a human being if it's all right to make the purchase.

Then a supplier fails, or prices move, and that ten-thousand-dollar limit, perfectly sensible when it was set, is now the thing standing between the company and a stockout. So Samantha, who runs vendor relations, decides to raise the limit going forward to fifty thousand.

Now what just happened? A human being decided how much authority a machine gets. She didn't decide what to buy or when to buy it. She decided what the system may decide without her.

And that's the shape of it. There's not one line between people and systems. There are three states with two thresholds. Some decisions the system completes alone. Some it prepares, and a person approves. And some never leave a human desk.

We built that in 20 years ago because it was good engineering. No law required it. That came later.

The decision that cannot be delegated

- A licensing case, flagged and fully prepared
- The system lays out the options
- A person makes the adverse determination
- The affected person gets an explanation
- That is not courtesy. It is design
- Somebody's whole job is guarding that line



16

Now the other end. A case where the answer has to be no.

A licensing board. A complaint against a practitioner. The system has read the file, cross-referenced the history, checked it against the standards, and laid out the options with the reasoning attached. Weeks of preparation for people, done in minutes.

And then a person makes the determination. A human being decides that this practitioner loses their license. Afterwards, the practitioner get a hearing where the decision is explained in terms they can argue with if they choose to.

That's not politeness. That's design. Somebody decided, in advance, that this decision would never be delegated to AI, no matter how good the system got.

And at that regulator, there's a person whose entire job is ethics, privacy, bias and security in these systems. Not a committee. A person. That job didn't exist a few years ago. There'll be many more in the near future, and some will be yours.

Where the law now sits

- Regulators are working from different problems
- They are converging on one requirement
- Machine authority must be made explicit
- Human responsibility must be made explicit
- Say what is reserved, and why



17

Here's what's changed in the last eighteen months, and it's why I'm more confident about this than I've ever been.

Regulators have started writing it down. Not one regulator. Several, working on completely different problems.

Start here in Texas. Since September of last year, if your health insurer is going to deny a treatment because it says the treatment isn't medically necessary, a machine is not allowed to make that call. Not wholly, and not partly. A qualified human being has to make it, and the denial notice has to tell you what criteria were used. The same law explicitly permits AI for the administrative work and for fraud detection. So it isn't a ban. It's a constraint.

Similarly, Medicare reserves nonpayment recommendations to licensed clinicians. Lending regulations requires the specific reasons for an adverse decision be stated to the applicant. Aviation guidance requires designers to say which responsibilities belong to people.

These aren't all the same rule. Different problems, different results.

But they converge on one demand. Make the system's authority explicit. Make the human's responsibility explicit. Be able to say which decisions are reserved, and why.

That's a design requirement now, not a philosophy.

The Autonomic Economy™

- Autonomic Enterprises transact directly with one another, in real time.
- Goods, services, information, and funds move on the same terms.
- Machines negotiate, schedule, and settle.
- People set the terms they work within.



18

So far, everything I've described happens inside one organization. Now put a few thousand of these organizations next to each other, and what happens?

I call it the Autonomic Economy.

Autonomic Enterprises transact directly with one another, in real time.

Goods, services, information and funds move on the same terms. And I want to stress that, because people hear this and picture container ships, or eighteen-wheelers. But it's not only physical goods. A service being delivered, a record being exchanged, money moving between accounts. These are all autonomic transactions in exactly the same sense.

Systems negotiate, schedule and settle. People set the terms they work within.

And here's what matters most to anyone who runs any organization. An enterprise can defer its own autonomic transformation. But it can't defer its counterparties' transformations.

Which means the timetable stops being yours. The slow link is the one every counterparty routes around, and nobody sends you a letter first.

A transaction with no one in it

- A shipment is delayed. Nobody is called
- The schedule adjusts. The parties are notified
- The record writes itself
- Rebates are calculated and paid
- No person appears anywhere in that



19

Let me show you one of those transactions.

A shipment is delayed. Weather, at a port a long way from here.

Nobody makes a phone call. The delay is simply detected. The schedule adjusts. Every party with an interest is notified, in the form each of them uses. The record writes itself. Contract terms are recalculated. Any rebates owed are worked out and paid.

So where was the person in that?

There wasn't one. Not because somebody was fired. Because nothing in that sequence required judgment. It required accuracy, speed and memory, and machines are better at all of these.

The slow link

- An enterprise can defer its own transformation.
- It cannot defer its counterparties' – its customers' or suppliers'.
- Nobody schedules a meeting to answer instantly
- The slow link is the one every counterparty routes around.
- That is a deadline you do not set



20

So here's why that matters to an organization that would prefer to take its time.

You can defer your own transformation indefinitely. What you can't defer is your counterparties'. Customers, suppliers, , carriers, banks, regulators.

And nobody schedules a Thursday planning meeting to answer a bid that expires in ninety seconds. If your organization is the slow link, you don't get an argument about it. You get routed around. Not out of malice. Out of arithmetic.

Where you will actually work

- Trade and logistics
- Finance and accounting
- Government and the border
- Healthcare
- Education




21

So let's get specific about where you're going to work.

The five places I mentioned, because that's where most of you are most likely to start.

When machines do the buying

- Documentation and routing: largely automated
- Machines buy continuously, not on a cycle
- More decisions and documents per dollar of trade
- Fewer trucks is possible. Fewer decisions is not
- Exceptions rise fastest of all
- Work that needs judgment, done here



22

Trade and logistics first, because that's Laredo. What's already true, briefly. Documentation and classification, largely automated at the bigger brokerages. Routing optimized continuously rather than set once a day by somebody with a spreadsheet. Exceptions flagged before a person looks. Invoice disputes settled against a shared record.

Now the part that's not yet well recognized. A human purchasing manager typically buys on a cycle. A monthly review, a quarterly contract, a forecast built in September for the following year. A system doesn't need a cycle. It re-optimizes continuously, and can move between suppliers as conditions move. And continuous re-planning means more decisions for the same trade. Every re-plan is another document, another classification, another compliance check, another filing.

Now, does that mean more trucks? Not necessarily. A system that re-optimizes will also consolidate loads and cut empty miles. This corridor could well move more trade on fewer crossings. In fact, last year there was roughly a 4% increase in the value of the goods coming in at Laredo with a 3% decrease in the number of trucks.

But notice what that doesn't touch. Fewer trucks is possible. But fewer decisions is not possible. And exceptions rise fastest of all, because the more often a plan changes, the more often something doesn't fit the rules.

And that creates more work. Brokerage, compliance, exception handling, all done by people. The routine end already automated. What grows is the end where a person has to decide.

So the question isn't whether AI reduces the work here. It's whether the work that remains gets done here, by you, or by somebody else.

Trade and logistics next

- Clearance predicted before the truck arrives
- Yard and warehouse operations largely self-managing
- Freight that schedules itself across carriers
- Compliance checked continuously, not at the gate
- Physical constraints remain. Decision delay falls



23

Now let's project three to five years out.

Clearance predicted before the truck reaches the bridge, so the decision about which trailer to open is made on the road rather than at the booth. Yards that largely manage themselves. Freight that schedules itself across carriers.

And compliance checked continuously rather than at the gate. That's where the advantage sits. A system that catches an unapproved substitution before it ships, proposes an approved alternative, and writes a log a regulator can read. That turns compliance from a cost into a reason to choose your company.

So go back to that line of trucks at dawn. Moving decisions upstream could shorten that queue. It won't eliminate it, because lanes, booths and staffing are physical constraints. What it changes is the part of the wait caused by deciding.

Robots and AI are one story

- TAMU's own robotics team, filmed here
- Assembly-line robots: one spot, one repeated task
- These sense, decide, and adapt as they go
- A robot is AI with a body
- Most people still treat them as separate



24

I want to show you about a minute of something, and it was filmed right here.

Now, most people still keep robots and AI in two separate compartments. For twenty-five years, robots have been arms bolted to assembly lines, doing one thing, in one place, over and over. And for the past few years, people have thought of AI as the thing on your laptop that writes.

But they're the same story. A robot is what AI looks like when you give it a body and put it someplace where physical objects actually move. And the interesting part is that the machine is deciding, in a place where conditions change.

Which is the same five stages I showed you earlier, running in a warehouse, or outdoors, instead of just on a spreadsheet.

Finance today

- Reconciliation and transaction review, increasingly automated
- Fraud detection running continuously
- Document processing faster than any team
- First-pass credit analysis already assembled
- Entry-level tasks are genuinely exposed



25

Finance and accounting.

Reconciliation. Transaction review. Fraud detection. Document processing. First-pass credit analysis. All of that is already being automated.

How far along any institution is varies enormously. Some large banks are well down this road. Plenty of community and regional institutions are not.

But that work is what the first two years of an accounting career used to consist of. So when somebody tells you the entry-level tasks in this field are exposed, they're right.

But let's look at what's actually happening to the work.

What the finance work becomes

- Exception judgment, where the rules run out
- Controls design, deciding what systems may do
- Client relationships nobody can automate
- Interpreting output no one else understands
- The first rungs move. They do not vanish



Here's a Tuesday morning, just a few years from now.

A customs broker applies for a working capital line. Thirty employees, eleven years of crossings behind her.

The bank's systems have read what they're permitted to read. Deposit history, filed returns, the shape of the receivables. And they notice something a credit committee wouldn't have seen for another month. Her receivables have lengthened. But the same lengthening is showing up this quarter across four other brokerage clients of the same bank. So, it isn't her business softening. It's something happening to the corridor, and she is carrying it on her books.

By eleven the file is complete, tested against credit policy, reasoning attached.

And if the terms need to change, somebody has to say why. Specifically. To the woman who asked.

Now notice what the machine actually did. It saw across the whole portfolio, which no individual credit officer does, because nobody reads five files side by side. What it couldn't do was have the call. A person carries that conversation, and carries the account afterwards.

That's the job that grows. Judgment where the rules run out. Deciding what the system may decide. Being the person who can explain an answer nobody else in the building can. Those are the new first rungs, and they're harder but way more interesting.

Government and the border

- Government here is a major employer
- Screening, case processing, benefits administration
- Licensing decided in hours, not months
- A state assistant answering thousands weekly
- Oversight built in: dashboards, reviews, logs



27

Government, and I include the border, because here in Laredo they're not separable. It's a major employer here precisely because of the crossing.

Screening, case processing, benefits administration. Licensing decisions that took a graduate weeks or months, now returned in hours. State agencies running assistants handling thousands of inquiries a week without adding staff.

And the agencies doing this well built the oversight in. Dashboards showing what the system is asked. Humans reading transcripts, deliberately, looking for the answers the system got wrong.

That isn't bureaucratic caution. That's what makes the system trustworthy enough to keep using.

Healthcare

- Documentation, scheduling, imaging support, revenue cycle
- A medication error caught and escalated immediately
- The nurse and the supervisor told together
- Correction, not surveillance
- Presence and judgment stay with people



28

Healthcare, briefly.


The routine layer is documentation, scheduling, imaging support, revenue cycle. Volumes of necessary work nobody went to nursing school to do.

But here's the part worth keeping in mind. A medication error, caught as it happens, with the nurse and the supervisor told at the same moment.

Notice the design. Both told together, immediately, while it can still be fixed. That's correction rather than surveillance, and the difference is an escalation decision – a choice a person made.

Education

- Students already use these tools daily
- They cannot be blocked. They can be guided
- One framework beats scattered departmental tools
- Faculty are being asked to integrate AI
- Integration is not the same as decoration



29

Education. And now I'm talking to the faculty for a moment.

Your students already use these tools, every day. That's not a discipline problem, because it can't be blocked, and the institutions that have tried only made it invisible.

What works is guidance. And one framework across an institution beats each department buying its own tool and finding its own problems separately.

You may be asked to integrate this into your course loads. But integration isn't decoration. Adding a session or two about AI to a syllabus is just decoration. Changing what you ask a student to produce, so the thinking is the assignment and the drafting isn't, is integration. Only the second helps them.

What AI does for teaching

- Research, drafting, scheduling, first-pass evaluation
- Professors should not outsource the thinking either
- The ideal class is one professor, one student
- Nobody has that luxury with hundreds enrolled
- AI could get us closer within five years



29

Most of what these tools do inside a university is make the work of faculty and staff more efficient. Research. Organizing material. A first pass at evaluating papers, though I'd say the same to the faculty that I just said to the students. Don't outsource the thinking.

But here's where it goes, and this is a three-to-five-year idea rather than a next semester idea.


The ideal education has always been one experienced professor and one student. Somebody who knows the subject and the learning process, who can see what this student wants out of the class, where they're stuck, and can reposition the material so it lands.

But that's never been available to anyone except the very wealthy, because professors have hundreds of students, not one.

These systems can get us closer to that than we've ever been. Not by replacing the professor. By giving the professor the time and the visibility to teach each student as though that student were the only one.

The pattern

- AI absorbs the routine and the repetitive
- It amplifies the analytical
- It escalates what falls outside its limits
- Judgment at the boundaries stays with people
- Accountability stays with people, always



30

So, five sectors, and I hope you noticed I told the same story five times.

Here's the pattern. AI absorbs the routine and the repetitive. It amplifies the analytical. And it escalates what falls outside its limits.

Now let me be precise about the human part. Plenty of judgment is delegated to these systems, deliberately, inside limits. Samantha delegated fifty thousand dollars of it. What stays with people is judgment at the boundaries, where the rules run out.

Accountability is different. That doesn't get delegated at all, because there's no person on the other end to hold responsible.

The fear, and the answer

- The fear is not losing the job you hold
- It is whether the first rung still exists
- AI is replacing tasks, and some roles
- Most jobs get redesigned, not deleted
- People who use it will replace those who do not



Stanford Digital Economy Lab, Canaries in the Coal Mine, August 2026 update. ADP payroll data

31

Now. Something everybody here is still thinking about. When I talk to people who already have jobs, their fear is that the work they do now gets done by a machine. That's not the concern among the students in this room.

Yours is sharper. It's whether the first rung will even be there when you reach for it. Because you've heard, correctly, that what the first two years of a career used to consist of is exactly what automates first. So let me answer honestly.

AI is replacing tasks. And it's replacing some roles outright. Call centers. Data entry. Transcription. First-pass document review. If a job is nothing but tasks a machine does faster and cheaper, that isn't a hypothetical problem. It's a timeline problem.

And I can put a number on the part that concerns you directly. Stanford's Digital Economy Lab tracks payroll records across millions of workers. For people aged twenty-two to twenty-five, in the occupations most exposed to AI, employment is now running about nineteen percent below where it would be if it had kept pace with the same age group in less exposed work. A year ago that gap was fifteen percent. And it's happening through reduced hiring rather than layoffs, which is exactly what a narrowing first rung looks like.

But most jobs are not being deleted. They're being redesigned around what's left.

And people who understand how to use these tools will replace people who don't. That's not a slogan, it's a mechanism. The technology doesn't come for your job. A colleague who's fluent in it produces work you can't match, and the organization notices.

The first rung is moving, and moving is not nothing. Anybody who tells you it will simply be fine is asking you to trust them rather than showing you the reasoning. So here's the reasoning.

The machine raises people

- A few experts perform best in any organization
- Their judgment can be captured and shared
- Everyone else can move toward that standard
- Then the tools assist the experts too
- Designed well, the curve lifts rather than flattens



32

First. In any organization, a small number of people are dramatically better at the work than everybody else. They're the ones people walk down the corridor to ask.

Second. These AI-enabled systems are now good at capturing how those people decide. Not just what they know, which you can put in a manual. How they judge. Which exceptions matter, what they check first. That now becomes available to everybody else, so the person in their second year works from the instincts of the person in their twenty-second.

Third. The one people miss. The tools also assist the experts. So, the top doesn't stand still while the bottom catches up.

Designed well, the curve lifts rather than flattens. But this isn't automatic. Built carelessly, the same technology levels everybody down to the average of its training data.

Now, in July of 1995, I stood in front of a room of manufacturing executives and showed them something new called a website, where I'd just ordered my own book. And when I asked how it might apply to their businesses, one senior man told me, politely, that it wouldn't, because manufacturers do not sell books.

The easy lesson is that he lacked vision. That's the wrong lesson, and it isn't fair to him. He was reasoning from thirty years of experience in a world where that had been true every day. His experience is what trapped him.

There's a version of this running in professional services right now. Firms are finding that a first-year associate who uses these tools well can outproduce a partner of twenty-five years who won't touch them. The hierarchy inverts. That isn't a technology problem. It's a skills problem, and it's the most fixable one I've described.

You students don't have thirty years of experience. Traditionally that would be a great disadvantage. For the next few years, that's not only NOT a disadvantage. It may even be an advantage.

**Show your work.
And your prompts**

- Third grade arithmetic:
half the grade was the answer
- The other half was showing how you got there
- Same rule now, including the prompts you used
- Writing is how a thought gets processed
- The degree is the thinking, not the paper



32

Now let me answer another question you're probably being asked, which is whether using these tools means you're not learning anything.

Think back to third grade arithmetic. Your teacher handed out a quiz and said: show your work. Half the grade is the right answer. The other half is the steps. So one addition error didn't destroy you if you understood the process, and guessing right earned you very little if you didn't.

That rule still works. It just needs one word added. Show your work, and show your prompts. Not a two-thousand-word paper on Lewis and Clark. Show me how you searched, what you asked, what came back, what you kept, what you threw away, and why.

That does two things. It proves you thought your way through it. And it teaches you to use the tool properly, which is a skill nobody is grading today but every employer will value.

Professor Clarke said something to me recently that I haven't stopped thinking about. Nobody assigns a paper on Shakespeare because they expect a nineteen-year-old to advance Shakespearean scholarship. They assign it because it forces discipline in reading and in writing, and both of those are thinking. Writing is how a thought gets processed. If you never write, you never find out what you actually believe.

The machine can produce the paper. It can't do the thinking for you. That part is the degree.

One more thing, then I'll get practical. When I was working on my MBA, I took a part-time job keeping the books for a small business. Excellent, I thought, real accounting at last. Ninety percent of it turned out to be sorting piles of receipts and writing down numbers.

That work is going away. Good. What's left is the part you came here for.

Your unfair advantages

- Two languages, used properly, every day
- Working knowledge of how crossings actually function
- Relationships on both sides of the river
- Judgment about a place, formed by living in it
- Tacit and current. No AI model holds it



33

Now let me tell you something about this room specifically. Most of you are from here, and a lot of you are going to stay. And that's sometimes framed as a limitation. I'd like to reframe it. You operate in two languages. Not just conversationally. Properly, in a working setting, with the register right. You know how a border crossing actually functions. You have relationships on both sides of the river, and judgment about a place formed by living in it.

Now let me be accurate about what that's worth. These AI models know a great deal about Laredo. Spanish, customs procedure, trade law. What they don't have is what's true here this month, what a particular broker will actually accept, and the hundred things people in this business and in this city know and have never written down.

That knowledge is tacit, current and local. So when you walk into an employer trying to work out how this applies to a corridor economy, you are not necessarily the junior person in the room on that question.

The next twelve months

- Learn what your field's exceptions look like
- Practice explaining a decision, in writing
- Tell the tool to argue with you
- Check AI output in both your languages
- Ask employers what their systems may decide
- And employers: one key change



34

So, practically, what should you do in the next 12 months? Five things, all startable this semester. And none of them is learn AI, because that isn't advice, that's just common sense.

First, learn what the exceptions look like in your field. In accounting, the transactions that don't fit the rule. In logistics, the shipments that break the pattern. The routine work is going away. The exceptions are the job that stays.

Second, practice explaining a decision in writing. Not summarizing. Explaining. Why this and not that, in language the affected person can understand and argue with.

Third, and do this tonight. Tell your AI tools to argue with you. An executive told me recently that AI flatters him, and something about that felt wrong. He was right. So. write standing instructions telling your AI to be skeptical, to check your reasoning, and to say what you haven't considered. Then read what comes back instead of simply skimming for agreement.

Fourth, and this one is very specific to you. Become the person who can check AI output in both languages. Not translate. Check. Whether the Spanish term the system used is the right term in a customs filing, or a loan document. That's narrow, verifiable and employable, and very few people have it.

Fifth, in every interview, ask whoever would actually supervise you: what are your systems allowed to decide without a person? Then the follow-up, which matters even more. When your systems gets one wrong, who finds out, and how fast?

An organization that has done the work answers both in about a minute. One that hasn't will tell you it's just a tool and humans review everything, which is what people say when nobody has checked. Either way, you have just learned what your first year is going to consist of if you go to work for this organization.

And employers, one thing. Entry-level work is the easiest thing to automate and the most expensive thing to lose, because it's where judgment gets built. Redesign the first rung consciously. Don't try to just delete it.

Back to the bridge

- The same trucks. The same crossing
- Now you see the decisions inside it
- Some of them belong to the machines
- The consequential ones belong to you
- That work is being redesigned, not deleted



35

Let's go back to the bridge.

Same trucks. Same queue. Nothing about that picture has changed in the last forty minutes.

What's changed, I hope, is that you can see what's inside it. Thousands of decisions made every day. And in five years a great many will be made by machines, correctly, faster than any of us could, and nobody will miss making them.

But the ones that carry consequence for a person. The ones where somebody is owed an explanation. The ones where the rules run out and somebody has to choose what the rule should be. Those are going to belong to people.

They're going to belong to you.

That work isn't disappearing. It's being redesigned. And it's getting harder, and more interesting, and considerably more important.

Where to go next

- Slides for today's talk
- The Executive AI Briefing, free to download
- Both at [HTTPS://JackShaw.com/tamiu](https://JackShaw.com/tamiu)
- Jack Shaw
Jack@JackShaw.com
770-910-5969
[HTTPS://JackShaw.com](https://JackShaw.com)



Now, two more things and then I'll take your questions.

First, there's one page, [jackshaw.com slash tamiu](https://jackshaw.com/tamiu), and it carries copies of these slides and my comments. My Executive AI Briefing, which goes deeper into this framework, is on the same page, and it's free.

Second, my contact details are on screen. Please reach out to me if you have questions – whether tomorrow or years from now.

Thank you for listening.

Jack Shaw

AI Transformation Advisor – Keynote Speaker – Author of The AI Imperative Series

JackShaw.com