Charta | LESSON NOTES

AI and Financial Stability

FRM Part II

1 A story: one bank, one market, one supplier

- Tessaly Bank: AI once for fraud, now almost everywhere
- Loans, credit papers, customers, code
- Much of it rented from one provider, Nimbrel
- Most rivals rent from Nimbrel too
- Stability: can the whole system absorb a shock?

IN THE LESSON

- Picture Tessaly Bank, a mid sized lender.
- Five years ago it used A I in one place, a fraud filter on card payments.
- Today it uses A I to score loans, to draft credit papers, to answer customers and to write code.
- Much of that runs on models rented from one outside provider, which we will call Nimbrel, a supplier of cloud computing and large models.
- Here is the twist.
- Most of Tessaly's rivals rent from Nimbrel too.
- Each bank's choice looks sensible on its own.
- The question for this lesson is what happens to the whole system when many firms make the same sensible choice.
- That is the question of financial stability, which means the system's ability to keep providing credit and payments through a shock.

2 What we mean by AI

- AI: computers doing tasks that seem to need intelligence
- Machine learning: the model learns its rules from data
- Deep learning: large networks with many layers
- Generative AI: creates text, code or images
- Foundation model: trained once, adapted by many

IN THE LESSON

- First, some words.
- Artificial intelligence, or A I, means computer systems that do tasks we usually think need human intelligence, such as spotting patterns or writing text.
- Machine learning is the main way it is built today: instead of following rules written by people, the model learns the rules from data.
- Deep learning uses very large networks with many layers, which suit images, speech and text.
- Generative A I creates new material, such as text, code or images, rather than just scoring or classifying.
- Its best known form is the large language model, trained on huge amounts of text.
- And a foundation model is a large model trained once on broad data, which many users then adapt to their own tasks.

Charta Learning · charta-learning.com

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.

3 Two forces behind the advances since 2017

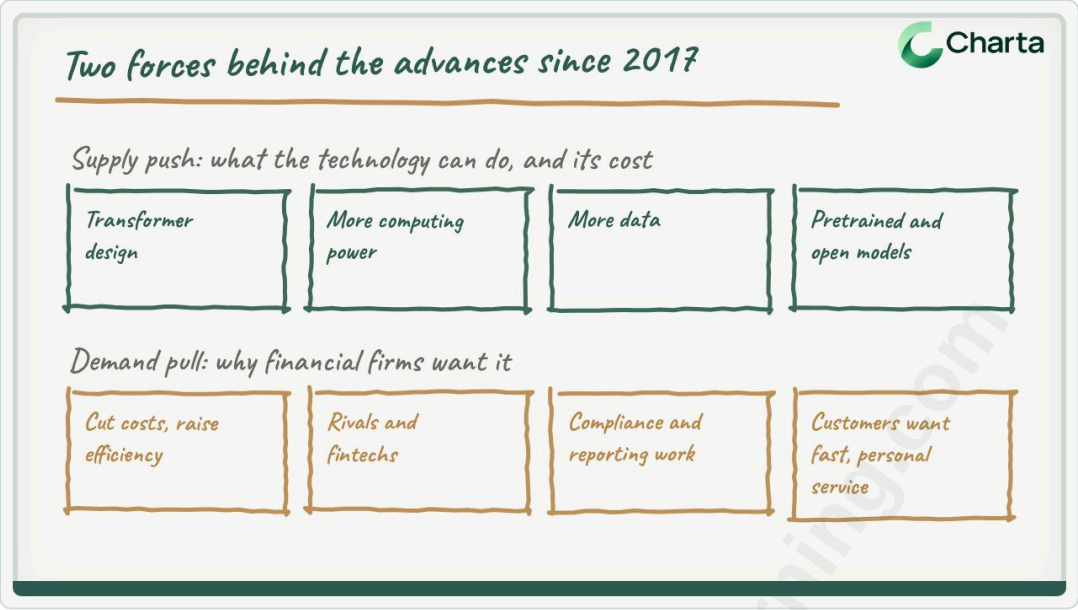

IN THE LESSON

- Why has A I advanced so fast since two thousand and seventeen?
- The curriculum sorts the reasons by where the push comes from.
- Look at the first row, supply push.
- These are developments in what the technology can do and what it costs.
- There are four: the transformer design, more computing power, more data, and pretrained and open source models.
- Now the second row, demand pull.
- These are the reasons financial institutions want the technology.
- There are four again: cutting costs and raising efficiency, competition from rivals and fintechs, heavier compliance and reporting work, and customers who expect fast, personal service.
- A simple test separates them.
- Supply is about the technology.
- Demand is about the firm's own reasons.

4 Supply push in more detail

- Transformer, 2017: trains in parallel on huge text sets
- Specialised chips and rented cloud capacity
- Large digital datasets to learn from
- Pretrained and open-source models: adapt, not build

IN THE LESSON

- Take the supply side one factor at a time.
- First, the transformer, an architecture for neural networks published in a research paper in two thousand and seventeen.
- It lets a model weigh every word against every other word, and, crucially, it trains in parallel, so models can learn from very large collections of text.
- Second, far greater computing power, from specialised chips built for this kind of maths and from cloud capacity that firms can rent by the hour.
- Third, the growth of large digital datasets to learn from.
- Fourth, the release of pretrained and open source models.
- A firm no longer has to build a model from nothing.
- It can take one that already exists and adapt it cheaply to its own task.

Charta Learning · charta-learning.com

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.

5 Demand pull, and how the two reinforce

- Cut costs and raise efficiency
- Competition from rivals and fintechs
- Heavier compliance and reporting workloads
- Customers expect fast, personal service
- They reinforce: cheaper supply makes more uses worthwhile

IN THE LESSON

- Now the demand side, the reasons inside financial firms.
- First, pressure to cut costs and raise efficiency, for example by automating slow manual work.
- Second, competition, from rival banks and from fintechs, meaning young technology led finance firms, which force incumbents to keep up.
- Third, heavier compliance and reporting workloads, from rules on money laundering, sanctions and disclosure.
- Fourth, customers who expect fast, personal service at any hour.
- The two sides reinforce each other.
- As supply gets cheaper, more of the uses that firms want become worth paying for.
- And as firms want more, providers invest more, which makes supply cheaper still.
- Our first example shows the first half of that loop in numbers.

6 Example 1: cheaper supply makes a use case pay

20,000 memos a year, 2 hours saved each, $50 an hour. Tool: $3.0m or $0.5m

Time saved a year	20,000 × 2 hours	**40,000 hours**
Value of time saved	40,000 × $50	**$2.0m**
Build its own model	2.0 − 3.0	**−$1.0m**
Adapt a pretrained model	2.0 − 0.5	**+$1.5m**

IN THE LESSON

- Tessaly writes twenty thousand credit memos a year, the papers that explain each lending decision.
- A I could save two hours on each, and staff time costs fifty dollars an hour.
- A tool to do it costs either three million or half a million dollars a year, depending on how it is built.
- The time saved is twenty thousand times two, forty thousand hours a year.
- At fifty dollars an hour, that is worth two million dollars a year.
- Some years ago, Tessaly would have had to build its own model, at about three million dollars a year.
- Two minus three is a loss of one million dollars, so the idea was shelved.
- Today it can adapt a pretrained model rented through the cloud for half a million dollars a year.
- Two minus zero point five is a gain of one point five million.
- The demand never changed.
- Cheaper supply made the same use case worthwhile.

Charta Learning · charta-learning.com

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.

7 Example 2: supply or demand? A classic mistake

Sort each factor: supply push (the technology) or demand pull (the firm's reasons)

Factor at Tessaly	Side
Board wants to cut back-office costs	Demand
Cloud computing price per hour falls	Supply
A fintech rival offers instant loans	Demand
An open-source model is released	Supply
New rules add to the reporting workload	Demand

IN THE LESSON

- Now sort five factors at Tessaly.
- Each one is either supply push, about the technology, or demand pull, about the firm.
- First, the board wants to cut back office costs.
- That is demand.
- This is the classic mistake: cost saving sounds like a supply matter because it is about money, but it is the firm's own motive, so it pulls from the demand side.
- Second, the price of cloud computing per hour falls.
- That is the technology getting cheaper, so supply.
- Third, a fintech rival offers instant loans.
- That is competition, demand.
- Fourth, an open source model is released.
- That is supply.
- Fifth, new rules add to the reporting workload.
- That is a compliance burden on the firm, so demand.

8 Established uses at financial firms

- Fraud detection and anti-money-laundering monitoring
- Credit scoring and insurance underwriting
- Algorithmic trading and portfolio signals
- Risk management and stress testing
- Customer service chatbots

IN THE LESSON

- Now the uses of A I, starting with financial firms.
- Several are well established.
- First, fraud detection and anti money laundering monitoring, where models flag unusual payments and transactions for staff to check.
- Second, credit scoring and insurance underwriting, where models estimate how likely a borrower is to default or a policyholder is to claim.
- Third, algorithmic trading and portfolio signals, where models help decide what to buy or sell and when.
- That is the subject of a separate lesson on capital markets, so here we note it and move on.
- Fourth, risk management and stress testing, where models help measure exposures and explore bad scenarios.
- Fifth, customer service chatbots, which answer routine questions.

9 Newer uses: generative AI, mostly internal

- Summarising documents: contracts, reports, notes
- Drafting and checking code
- Supporting compliance and research
- A human reviews the output before it reaches a client

IN THE LESSON

- The newer uses come from generative A I, and they are mostly internal, meaning they help staff rather than face clients.
- First, summarising documents, such as long contracts, reports or meeting notes.
- Second, drafting and checking computer code, which speeds up the work of developers.
- Third, supporting compliance and research, for instance by finding the relevant rule or pulling facts out of filings.
- The common thread is a human in the loop.
- A person reviews the output before anything reaches a client or drives a decision.
- That matters because generative models can be wrong in a fluent, confident way, as our next example shows.

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.

10 Example 3: why a human reviews the drafts

AI drafts 500 client letters a week; 4% contain errors; reviewers catch 95% of them

Letters with errors	4% × 500	**20**
Caught by reviewers	95% × 20	**19**
Reach clients with review	20 − 19	**1 a week**
Reach clients, no review	all of them	**20 a week**

IN THE LESSON

- Tessaly's generative tool drafts five hundred client letters a week.
- Suppose four per cent contain an error, such as a wrong fee or an invented fact.
- Four per cent of five hundred is twenty letters with errors each week.
- Staff review every draft, and they catch ninety five per cent of the errors.
- Ninety five per cent of twenty is nineteen caught.
- So twenty minus nineteen, just one flawed letter a week, reaches a client.
- Without the review, all twenty would go out.
- That is twenty times as many.
- This is why firms start with internal uses and keep a person checking the output.

11 The official side: suptech and regtech

- Suptech: supervisors and central banks using AI
- Collect and analyse reported data
- Flag anomalies; scan disclosures; monitor markets
- Regtech: firms automating their own compliance

IN THE LESSON

- Now the official side.
- Supervisory technology, or suptech, is the use of technology, including A I, by supervisors and central banks to do their own job.
- They apply machine learning and natural language processing, meaning software that works with human language, in four broad ways.
- First, to collect and analyse the data that firms report.
- Second, to flag anomalies, such as a bank whose figures suddenly look odd.
- Third, to scan disclosures and other documents at a scale no team of people could manage.
- Fourth, to monitor markets for signs of manipulation or stress.
- Regtech, or regulatory technology, is different.
- It is the firm side automation of compliance.
- Do not confuse the two labels.

12 Example 4: regtech or suptech?

Who uses the tool decides the label: a firm (regtech) or an authority (suptech)

Tool	Label
Tessaly screens payments against sanctions lists	Regtech
Supervisor scans bank filings for anomalies	Suptech
Supervisor's AI watches trading for manipulation	Suptech
Tessaly fills in its regulatory returns	Regtech

IN THE LESSON

- Label four tools.
- The test is simple: who is using the tool?
- First, Tessaly screens its payments against sanctions lists.
- A firm is meeting its own compliance duty, so that is regtech.
- Second, the supervisor scans bank filings for anomalies.
- The authority is doing its job, so that is suptech.
- Third, the supervisor's A I watches trading for manipulation.
- Again the authority, so suptech.
- Fourth, Tessaly automatically fills in its regulatory returns.
- The firm again, so regtech.
- A tempting error is to call the third one regtech because it is about regulation.
- It is not.
- Regulation is the subject of both; the user decides the label.

Charta Learning · charta-learning.com

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.

13 How adoption tends to move

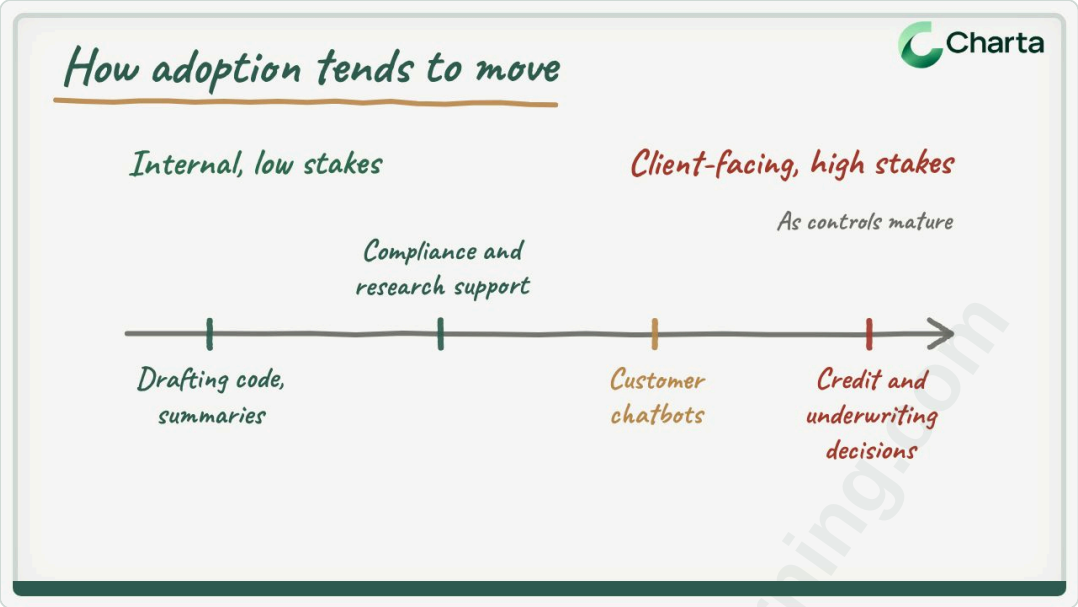

IN THE LESSON

- Adoption tends to follow a path.
- Here is a line from internal, low stakes tasks on the left to client facing, high stakes decisions on the right.
- The firm moves along it as its controls mature.
- At the far left is drafting code and summarising documents, where an error is caught inside the firm.
- Next comes support for compliance and research, still internal, but feeding real decisions.
- Further right are customer chatbots, which speak to clients directly.
- At the far right are credit and underwriting decisions made mainly by A I, where a model's output decides who gets a loan or a policy, and at what price.
- Older scoring models are already used here; the point is how much weight we hand to newer ones.
- Each step right raises the cost of an error, so each needs stronger testing, monitoring and governance before the firm takes it.

14 The exam point for drivers and uses

- Supply: technology. Demand: the firm's reasons
- A firm's motive, such as cost saving, is demand
- Established uses, plus generative uses kept internal
- Suptech is the authorities; regtech is the firms
- Adoption moves toward high stakes as controls mature

IN THE LESSON

- So here is the exam point for the first two outcomes.
- The advances since two thousand and seventeen come from two forces.
- Supply push is about the technology: the transformer, computing power, data, and pretrained models.
- Demand pull is about the firm's own reasons: cost, competition, compliance and customers.
- A firm's motive, such as cost saving, is always demand.
- The two reinforce each other.
- On uses, firms have established uses in fraud, credit, trading, risk and chatbots, with newer generative uses that are mostly internal and reviewed by a person.
- Authorities use suptech, while firms use regtech.
- And adoption moves from internal, low stakes tasks toward client facing, high stakes decisions as controls mature.

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.

15 How AI could affect financial stability

- Third-party dependence and provider concentration
- Market correlations: similar models, similar actions
- Cyber risk: AI-armed attackers, attacks on models
- Model risk, with data and governance
- Others: fraud, disinformation, misalignment

IN THE LESSON

- Now the heart of the topic: how could A I affect financial stability?
- The curriculum names several channels, and we will take them in turn.
- First, third party dependence and provider concentration: many firms relying on a few suppliers.
- Second, market correlations: firms using similar models and acting alike.
- Third, cyber risk: attackers armed with A I, and attacks on the models themselves.
- Fourth, model risk, together with data and governance: models that are opaque, biased or simply wrong.
- Fifth, a group of others: A I enabled fraud, disinformation, and misalignment.
- Remember that A I also brings benefits, such as better fraud detection.
- The question is which risks could spread beyond one firm.

16 The two most systemic channels

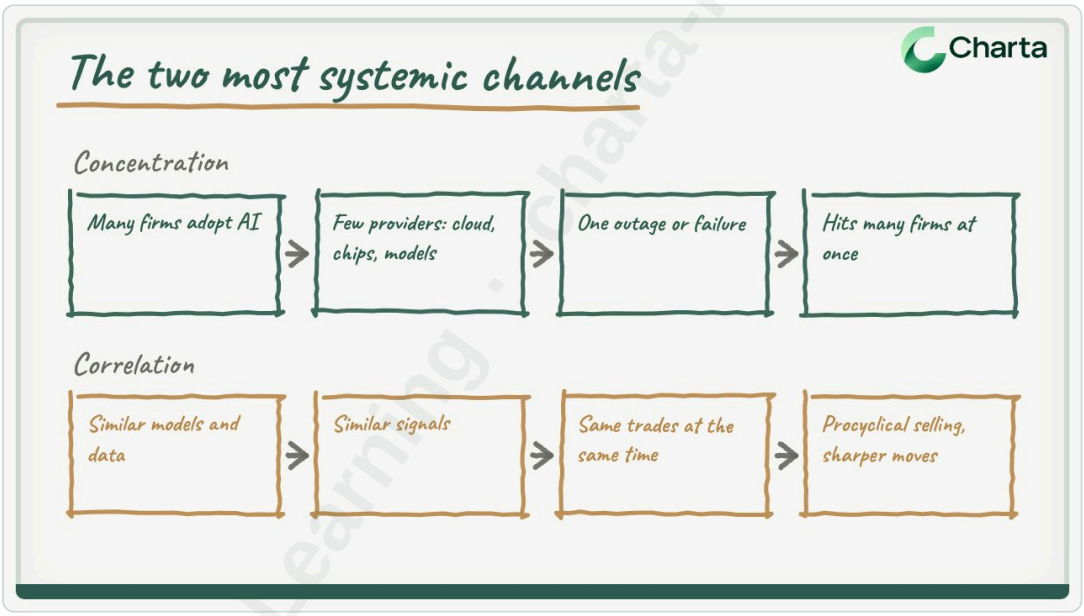

IN THE LESSON

- Two of these channels matter most for the system, so let us draw them first.
- The first row is concentration.
- Many firms adopt A I.
- They rely on a few providers of cloud computing, specialised chips and foundation models.
- So one outage or failure at a provider hits many firms at once.
- The second row is correlation.
- Firms use similar models trained on similar data.
- So they get similar signals.
- So they make the same trades at the same time, which is called herding.
- The result is procyclical selling, meaning selling that deepens a fall instead of cushioning it, and sharper price moves.
- Both rows end the same way.
- A problem that would hit one firm hits the whole system.

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.

17 Third-party dependence and concentration

- Few providers: scarce chips, compute and data
- Dependence: a service the firm cannot easily run itself
- Concentration: many firms rely on the same few
- Hidden links: fourth parties under the vendors
- Hard to substitute: switching takes time and testing

IN THE LESSON

- Look closer at the first channel.
- Building large models takes huge computing power, scarce chips and vast data, so only a few providers can supply them.
- That creates third party dependence: a firm relies on an outside supplier for a service it cannot easily run itself.
- Concentration adds a system view.
- When many firms rely on the same few suppliers, one outage, cyber attack or failure spreads widely.
- The links can be hidden.
- A bank's software vendor may itself run on a cloud provider and a foundation model, a fourth party the bank never contracted with.
- And these services are hard to substitute, because switching a model or a cloud takes time, data and testing.

18 Example 5: concentration in numbers

Ten banks, $50bn of assets each. Seven run key AI services on Nimbrel

System assets	10 × $50bn	**$500bn**
One bank's own model fails	50 ÷ 500	**10% of system**
Banks on Nimbrel	7 × $50bn	**$350bn**
Nimbrel outage hits	350 ÷ 500	**70% of system**

IN THE LESSON

- Tessaly's market has ten banks, each with fifty billion dollars of assets, and seven of them run their key A I services on Nimbrel.
- First, the system's assets: ten times fifty, five hundred billion dollars.
- Suppose one bank's own model fails.
- That hits one bank, fifty billion dollars, which is fifty over five hundred, ten per cent of the system.
- That is a firm level problem.
- Now the seven banks on Nimbrel.
- Seven times fifty is three hundred and fifty billion dollars.
- If Nimbrel has an outage, all seven are hit at once: three hundred and fifty over five hundred, seventy per cent of the system.
- The same kind of failure, but at a shared provider, turns a ten per cent problem into a seventy per cent one.

19 Market correlations and herding

- Similar models, similar data, similar decisions
- Herding: firms buy and sell together
- In a shock: one sell signal, many sellers
- Procyclical selling and sharper price moves
- Risk: not one bad model, but many agreeing

IN THE LESSON

- Now the second channel, market correlations.
- When many firms use similar models, trained on similar data, often from the same few providers, their decisions line up.
- That is herding: firms buy and sell together, not because they copied each other, but because their models think alike.
- In calm markets this may go unnoticed.
- In a shock, many models may give the same sell signal at once, and selling becomes procyclical, deepening the fall.
- Prices then move more sharply than they would if firms disagreed, and buyers on the other side become scarce.
- Speed adds to it, because automated systems can act within seconds.
- The risk is not one bad model, but many models agreeing.

Charta Learning · charta-learning.com

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.

20 Example 6: correlation in numbers

10 funds hold $1bn each. A sell signal: sell half. Price falls 1% per $1bn sold

Diverse models: 2 of 10 fire	2 × $1bn × ½	**$1bn sold**
Price fall	1% per $1bn	**1%**
Similar models: 8 of 10 fire	8 × $1bn × ½	**$4bn sold**
Price fall	1% per $1bn	**4%**

Same shock, same funds: only the number of models that agree changes

IN THE LESSON

- Ten funds in Tessaly's market each hold one billion dollars of the same corporate bonds.
- A shock hits, and any fund whose model gives a sell signal sells half its holding.
- Assume the price falls one per cent for each billion dollars sold in a day.
- First, diverse models.
- Only two of the ten fire.
- Two times one billion times a half is one billion dollars sold.
- At one per cent for each billion sold, the price falls one per cent.
- Now similar models.
- Eight of the ten fire together.
- Eight times one billion times a half is four billion dollars sold.
- At one per cent per billion, the price falls four per cent.
- Same shock, same funds.
- Four times the fall, purely from models thinking alike, and that fall can trigger yet more selling.
- The note at the bottom says it all: same shock, same funds, and only the number of models that agree has changed.

21 Cyber risk: attackers and attacked models

- Attackers use AI: convincing, tailored phishing
- Deepfakes: fake voice or video of real people
- Faster discovery of weak points
- Models attacked: data poisoning corrupts what they learn

IN THE LESSON

- The third channel is cyber risk, and A I raises it in two ways.
- First, attackers use A I.
- It writes convincing phishing emails, tailored to each target and free of the spelling errors that used to give them away.
- It makes deepfakes, meaning fake but realistic voice or video, so a caller can sound exactly like a bank's chief executive.
- It can also help find weak points in systems faster.
- Second, the models themselves can be attacked.
- In data poisoning, an attacker slips corrupted data into what a model learns from, so it behaves wrongly later.
- Defenders use A I too, but a successful attack on a widely shared provider links straight back to concentration.

Charta Learning · charta-learning.com

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.

22 Example 7: AI-written phishing

Illustration: 10,000 phishing emails; click rate 0.5% if generic, 2% if AI-tailored

Generic emails	10,000 × 0.5%	**50 clicks**
AI-tailored emails	10,000 × 2%	**200 clicks**
Footholds multiplied by	200 ÷ 50	**4 times**

IN THE LESSON

- Here is an illustration with made up rates.
- An attacker sends ten thousand phishing emails to staff across Tessaly's market.
- Suppose half of one per cent click on a generic email, and two per cent on one tailored by A I, with the right names, projects and tone.
- With generic emails, ten thousand times half of one per cent is fifty clicks.
- With A I tailored emails, ten thousand times two per cent is two hundred clicks.
- Two hundred over fifty is four.
- The attacker gets four times as many footholds inside the firms, for almost no extra cost.
- That is how A I scales up an old attack.

23 Model risk, data and governance

- Opacity: hard to explain, so errors hide
- Poor data quality: gaps and stale data
- Bias against groups of customers
- Hallucination: fluent, confident and false
- Weak governance and validation

IN THE LESSON

- The fourth channel is model risk, together with data and governance.
- Model risk is the risk of loss from a model that is wrong or used wrongly.
- A I sharpens it in five ways.
- Opacity: complex models are hard to explain, so errors are hard to spot.
- Poor data quality: a model learns only what its data show, so gaps and stale data become wrong answers.
- Bias: a model can treat groups of customers unfairly if its data were unfair.
- Hallucination: a generative model can produce fluent, confident output that is simply false.
- And weak governance and validation: no clear owner, no independent testing, and no monitoring once the model is in use.

24 Example 8: a model trained on good years

$2bn loan book. Model predicts 1% defaults; downturn brings 4%. Half of each default is lost

Model's expected loss	2,000 × 1% × 50%	**$10m**
Loss in the downturn	2,000 × 4% × 50%	**$40m**
Shortfall	40 − 10	**$30m**

IN THE LESSON

- Tessaly's A I credit model was trained only on years of strong growth, a data quality problem.
- On a two billion dollar loan book it predicts a default rate of one per cent.
- Assume half of any defaulted loan is lost, which is the loss given default.
- The model's expected loss is two thousand million times one per cent times fifty per cent, ten million dollars.
- Then a downturn comes, and the true default rate is four per cent.
- The loss in the downturn is two thousand million times four per cent times fifty per cent, forty million dollars.
- The shortfall is forty minus ten, thirty million dollars, which Tessaly had not provisioned for.
- If rivals trained similar models on similar data, they all miss it together.

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.

25 Other channels: fraud, disinformation, misalignment

- AI-enabled fraud: deepfake voices, fake identities at scale
- Disinformation: false videos or reports spread fast
- It can trigger a run or a sharp market move
- Misalignment: AI pursues objectives nobody intended

IN THE LESSON

- The fifth group covers other channels.
- A I enabled fraud uses the same tools as the attackers: deepfake voices that authorise payments, and fake identities built at scale to open accounts and take loans.
- Disinformation is false information spread on purpose.
- A fake video of a bank's chief executive, or a false report of losses, can spread across social media in hours, and that can trigger a run, meaning depositors rushing to withdraw, or a sharp market move.
- Misalignment is where an A I system pursues objectives its designers did not intend.
- A system told to maximise returns, for example, might find a way to do so by taking risks or breaking rules its designers never spelled out.

26 Example 9: disinformation and a run

Deposits $40bn; liquid assets $6bn. A deepfake triggers withdrawals of 20% in a day

Outflows	20% × $40bn	**$8bn**
Liquid assets on hand		**$6bn**
Gap it cannot meet	8 − 6	**$2bn**

IN THE LESSON

- Tessaly has forty billion dollars of deposits and six billion dollars of liquid assets it can turn into cash at once.
- Then a deepfake video shows its chief executive admitting to huge losses.
- It is false, but it spreads fast.
- Suppose twenty per cent of deposits are withdrawn in a day.
- The outflows are twenty per cent of forty billion, eight billion dollars.
- Liquid assets on hand are six billion dollars.
- Eight minus six leaves a gap of two billion dollars that Tessaly cannot meet from its own liquid assets.
- It must sell other assets at a loss, borrow in a hurry, or turn to the central bank.
- A healthy bank can fail for lack of cash, and the panic can spread to banks that look similar.

27 Example 10: three vendors, one provider

Tessaly buys AI services from three vendors. Is it diversified?

Vendors under contract	looks diversified	**3**
Cloud underneath all three	Nimbrel	**1**
Foundation model underneath	Nimbrel	**1**
Real points of failure	not 3	**1**

IN THE LESSON

- Here is a classic mistake.
- Tessaly's risk team reports that it is diversified, because it buys A I services from three different vendors.
- But look underneath.
- All three vendors rent their computing from one cloud underneath, Nimbrel's.
- And all three build on one foundation model underneath, also Nimbrel's.
- So the real number of providers that could stop all three services at once is one, not three.
- Counting the contracts hides the fourth parties.
- The fix is to map the supply chain down to the cloud and the model, and to plan for that shared point failing.
- Supervisors face the same task across the whole system, since they need to see which providers sit under many firms at once.

Charta Learning · charta-learning.com

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.

28 The channels compared

How each channel spreads, and how far it tends to reach

Channel	How it spreads	Reach
Concentration	One failure hits many	System-wide
Correlation	Same trades at once	System-wide
Cyber	Attacks scaled by AI	Can spread
Model risk	Flawed outputs	Mostly one firm
Fraud, disinformation	Losses, runs	Can spread

IN THE LESSON

- Here are the channels side by side, with how each spreads and how far it reaches.
- Concentration spreads because one failure hits many firms, so its reach is system wide.
- Correlation spreads because firms make the same trades at once, so it too is system wide.
- Cyber risk spreads through attacks scaled up by A I.
- It often starts at one firm, but it can spread, especially through a shared provider.
- Model risk works through flawed outputs, and mostly hits one firm, unless rivals use similar models.
- Fraud and disinformation cause losses and can spark runs, so they can spread.
- The pattern is clear: the channels that link firms together are the ones that reach the whole system.

29 The exam point for financial stability

- Five channels: concentration, correlation, cyber,
- model risk with data and governance, and others
- Concentration and correlation: the most systemic
- They turn firm-level failures into system-wide ones
- Ask: who depends on whom, and who thinks alike?

IN THE LESSON

- So here is the exam point for the third outcome.
- A I can affect stability through third party dependence and provider concentration, market correlations, cyber risk, model risk with data and governance, and others such as fraud, disinformation and misalignment.
- Concentration and correlation are the two most systemic channels, because they turn firm level failures into system wide ones: one provider failing hits many firms, and many similar models act as one.
- The other channels mostly start inside a firm, but they become systemic when they run through a shared provider or a shared model.
- So the questions to ask are always the same: who depends on whom, and who thinks alike?

30 Mistakes to avoid

- Cost saving labelled as supply: it is demand
- Confusing suptech (authorities) with regtech (firms)
- Assuming generative AI already makes client decisions
- Counting vendors instead of the shared cloud and model
- Treating all channels as equally systemic

IN THE LESSON

- Before the quiz, here are the mistakes to avoid across the whole topic.
- Do not label a firm's own motive, such as cost saving, as a supply factor.
- Supply is about the technology; motives are demand.
- Do not confuse suptech, used by the authorities, with regtech, used by firms.
- Do not assume generative A I is already making client decisions; most uses are internal, with a human reviewing the output.
- Do not measure diversification by the number of vendors; look through to the shared cloud and model underneath.
- And do not treat all the channels as equal.
- Concentration and correlation are the most systemic, because they turn failures at one firm into failures across the system.

Charta Learning · charta-learning.com

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.

31 Try three quick questions

1. A bank adopts AI because a fintech rival approves loans almost instantly. Supply or demand?
   Answer: Demand pull: competition is the firm's own reason.
2. A central bank uses language models to scan banks' disclosures. Regtech or suptech?
   Answer: Suptech: an authority is the user.
3. Six of ten equal-sized banks use one AI provider. Share hit by an outage? Which channel?
   Answer: 6 ÷ 10 = 60%: provider concentration, highly systemic.

IN THE LESSON

- Try three quick questions, then listen to the answers.
- Question one.
- A bank adopts A I because a fintech rival now approves loans almost instantly.
- Is that supply push or demand pull?
- It is demand pull, because competition is the firm's own reason.
- Question two.
- A central bank uses language models to scan the disclosures of the banks it supervises.
- Is that regtech or suptech?
- It is suptech, because an authority is the user.
- Question three.
- Six of ten equal sized banks run their key A I services on one provider.
- What share of the system does an outage at that provider hit, and which channel is this?
- Six over ten is sixty per cent, through provider concentration, one of the two most systemic channels.

Summary notes for the Charta whiteboard lesson. Figures are from the lesson's made-up examples.

Charta Learning · charta-learning.com

© Charta Learning · charta-learning.com · For your own study. Please do not share or redistribute.