Charta | LESSON NOTES

AI in Capital Markets

FRM Part II

1 A story: one morning at Thornbury

- A model scores the night's news and filings
- An algorithm slices a large sell order
- A model flags odd order patterns among millions
- A chatbot drafts a summary; an analyst checks it
- All called AI, but very different tools and risks

IN THE LESSON

- Picture a morning at Thornbury Capital.
- Before the market opens, a model has scanned the night's news and company filings and scored each one as good or bad news.
- When trading starts, an algorithm decides how to slice a large sell order so that it moves the price as little as possible.
- In the compliance team, another model flags a handful of odd order patterns out of millions of orders.
- And in research, an analyst asks a chatbot to draft a summary of a bond prospectus, then checks every line of it.
- All four are called A I.
- But they are very different tools, with very different risks, and the first task of this lesson is to tell them apart.

2 What AI and machine learning mean

- AI: computers doing tasks that need human judgement
- Machine learning: the system learns its rules from data
- Trained on past data, then applied to new cases
- Most AI in markets today: conventional ML, narrow tasks
- Generative AI: creates new text, code or data

IN THE LESSON

- Some definitions first.
- Artificial intelligence, or A I, is the broad name for computer systems that perform tasks we would normally think need human judgement, such as recognising patterns, making predictions or writing text.
- Machine learning, or M L, is the part of A I in which the system learns its rules from data, instead of a programmer writing them by hand.
- A model is trained on past data, then applied to new cases.
- Most A I in markets today is this conventional machine learning, used for narrow, well defined tasks.
- Generative A I is a newer branch.
- Built on large language models, it creates new text, code or data rather than a score or a forecast.

Charta Learning · charta-learning.com

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

3 Three ways a machine learns

Conventional machine learning comes in three kinds

Kind	Learns from	Typical use
Supervised	Labelled examples	Credit scoring
Unsupervised	No labels: clusters, outliers	Surveillance, AML
Reinforcement	Trial and reward: a policy	Execution schedule

IN THE LESSON

- Conventional machine learning comes in three kinds.
- Supervised learning learns from labelled examples, meaning past cases where the right answer is known, such as which borrowers defaulted.
- It then predicts that label for new cases, so it suits credit scoring and default prediction.
- Unsupervised learning has no labels.
- It finds structure on its own, grouping similar cases into clusters or picking out outliers, the cases that look unlike the rest.
- That suits market surveillance and anti money laundering screening.
- Reinforcement learning learns by trial and reward.
- It tries actions, sees the result, and learns a policy, meaning a rule for which action to take in each situation.
- A trading execution schedule is the classic case.

4 Where machine learning is used today

Established uses: mostly conventional machine learning

Area	Task
Trading	Algorithms and optimal execution
Surveillance	Flag abnormal order patterns
Risk and compliance	Credit and market risk; AML
Research	News and filings: sentiment
Clients and operations	Robo-advice; back office

IN THE LESSON

- Established uses are mostly conventional machine learning.
- Here is where these tools already work in capital markets.
- In trading, algorithms handle algorithmic trading and optimal execution, which means splitting a large order to keep its cost down.
- In surveillance, models flag abnormal order patterns that may signal market abuse.
- In risk, they support credit and market risk analytics, such as default prediction and scenario analysis.
- In compliance, they screen clients and payments for anti money laundering.
- Research uses natural language processing, the branch of A I that reads text, to turn news and company filings into sentiment signals, scores of how positive or negative the text is.
- And for clients and operations there is robo advice, meaning automated investment advice, and back office automation, such as matching trades and reading documents.

Charta Learning · charta-learning.com

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

5 Example 1: Thornbury's execution algorithm

Thornbury sells 500,000 shares; price on arrival $40.00. Implementation shortfall?

All at once: avg $39.70	(40.00 − 39.70) × 500,000	**$150,000**
20 slices: avg $39.88	(40.00 − 39.88) × 500,000	**$60,000**
Saving on one order	150,000 − 60,000	**$90,000**

IN THE LESSON

- Thornbury must sell five hundred thousand shares, and the price when the order arrives is forty dollars.
- The cost of trading is measured by the implementation shortfall, the gap between that arrival price and the average price actually achieved.
- Sold all at once, the order eats through the buyers' bids and averages thirty nine dollars seventy.
- The shortfall is forty dollars minus thirty nine seventy, thirty cents a share, times five hundred thousand, which is one hundred and fifty thousand dollars.
- Thornbury's reinforcement learning model instead splits the order into twenty slices, timed to when buyers are present, and averages thirty nine dollars eighty eight.
- That shortfall is twelve cents a share, sixty thousand dollars.
- The saving on one order is ninety thousand dollars.

6 Example 2: surveillance alerts

2,000,000 orders a day. Old rules against an unsupervised model

Old rules: 1,500 alerts	12 genuine: 12 ÷ 1,500	**0.8%**
Model: 200 alerts	16 genuine: 16 ÷ 200	**8%**
Who decides?	analysts review every alert	**people**

IN THE LESSON

- Thornbury's compliance team watches two million orders a day for signs of abuse, such as spoofing, which means placing orders you never intend to fill to move the price.
- Its old fixed rules raise one thousand five hundred alerts a day, and on review only twelve are genuine, a hit rate of zero point eight per cent.
- An unsupervised model, which learns what normal trading looks like and flags what does not fit, raises two hundred alerts, of which sixteen are genuine.
- That is a hit rate of eight per cent, ten times better, and it found more genuine cases from far fewer alerts.
- But notice the last row.
- Analysts still review every alert.
- The model ranks the cases; people decide what is abuse.

7 Generative AI: what is new

- Large language models: trained to predict the next word
- Uses: drafting research, summaries, code, client queries
- Mostly pilots and internal tasks, a human reviews
- Hallucination: fluent output that is false
- Opacity: no one can fully explain the answer

IN THE LESSON

- Now look closer at generative A I.
- Its large language models are trained on vast amounts of text to predict the next word, which lets them write fluent new text, code or data.
- In capital markets its likely uses are drafting research, summarising long documents such as prospectuses, writing and checking code, and answering client queries.
- For now these are mostly pilots and internal tasks, with a human reviewing the output.
- Two limits explain why.
- The first is hallucination, meaning fluent output that is false, such as an invented figure, stated just as confidently as a true one.
- The second is opacity: nobody can fully explain why the model gave a particular answer.
- Both limit fully autonomous use.

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

8 Established ML against emerging GenAI

Two families of AI in capital markets

	Established ML	Emerging GenAI
Output	Prediction, score, action	New text, code, data
Tasks	Execution, surveillance, AML	Drafts, summaries, code
Maturity	Embedded in production	Pilots; human reviews
Key limit	Data quality, overfitting	Hallucination, opacity

IN THE LESSON

- Two families of A I work in capital markets.
- Put them side by side.
- First, the output.
- Established machine learning produces a prediction, a score or an action, such as a default probability or a trade.
- Generative A I produces new text, code or data.
- Second, typical tasks.
- Machine learning does execution, surveillance, credit and anti money laundering.
- Generative A I drafts, summarises, writes code and answers queries.
- Third, maturity.
- Machine learning is embedded in production, in daily use for years.
- Generative A I is mostly in pilots, with a human reviewing what it produces.
- Fourth, the key limit.
- For machine learning it is data quality and overfitting, meaning a model that fits the past too closely and fails on new cases.
- For generative A I it is hallucination and opacity.

9 Example 3: the classic mistake, all AI is GenAI

A board paper says Thornbury runs eight 'GenAI tools'. Sort them

Tools called 'GenAI'	in the board paper	**8**
Conventional ML	execution, surveillance, credit…	**6**
Generative AI	research drafts, code helper	**2**
Share really GenAI	2 ÷ 8	**25%**

Hallucination checks on the 2; ordinary validation on the 6

IN THE LESSON

- Here is the classic mistake.
- A paper for Thornbury's board says the firm runs eight generative A I tools, and proposes one set of controls for them all.
- Sort them.
- Six are conventional machine learning: the execution algorithm, the surveillance model, the credit score, anti money laundering screening, the news sentiment model and the document reader in the back office.
- Only two are generative: the research drafting assistant and the coding assistant.
- So the share that is really generative A I is two out of eight, twenty five per cent.
- The footnote gives the consequence.
- Hallucination checks belong on those two, while the other six need ordinary model validation.
- Calling everything generative A I points the controls at the wrong risks.

Charta Learning · charta-learning.com

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

10 The exam point on AI uses

- Established: narrow ML tasks across the firm
- Supervised, unsupervised or reinforcement learning
- Emerging: GenAI drafts, summarises, codes, answers
- Keep the two apart: different risks and controls
- Hallucination and opacity limit autonomous use

IN THE LESSON

- So the exam point for this first outcome.
- A I and machine learning are already embedded in capital markets, but in narrow, well defined tasks: trading and execution, surveillance, risk analytics, compliance and anti money laundering, sentiment from text, robo advice and the back office.
- These are conventional machine learning, supervised, unsupervised or reinforcement.
- Generative A I is the emerging part: drafting, summarising, coding and answering clients, mostly with a human checking.
- Keep the two apart, because their risks and controls differ.
- And remember why generative A I is not yet trusted to act on its own: hallucination and opacity limit autonomous use.

11 AI and market behaviour: two directions

- Normal times: better price discovery
- Deeper liquidity, better execution quality
- Concern: correlation
- Same foundation models, vendors and datasets
- Alike trades become herding in stress

IN THE LESSON

- Now the second outcome: what wider use of A I could do to the way markets behave.
- It cuts both ways.
- In normal conditions, it may help.
- Faster processing of information can improve price discovery, the process by which new information finds its way into prices.
- It can deepen liquidity, meaning how easily you can trade without moving the price, and it can improve execution quality, as Example one showed.
- The concern is correlation.
- If many firms rely on the same foundation models, meaning large general models that others adapt, and on the same vendors and the same datasets, their signals and their trades become alike.
- In stress, alike trades become herding, everyone moving in the same direction at the same time.

Charta Learning · charta-learning.com

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

12 From common inputs to a crash

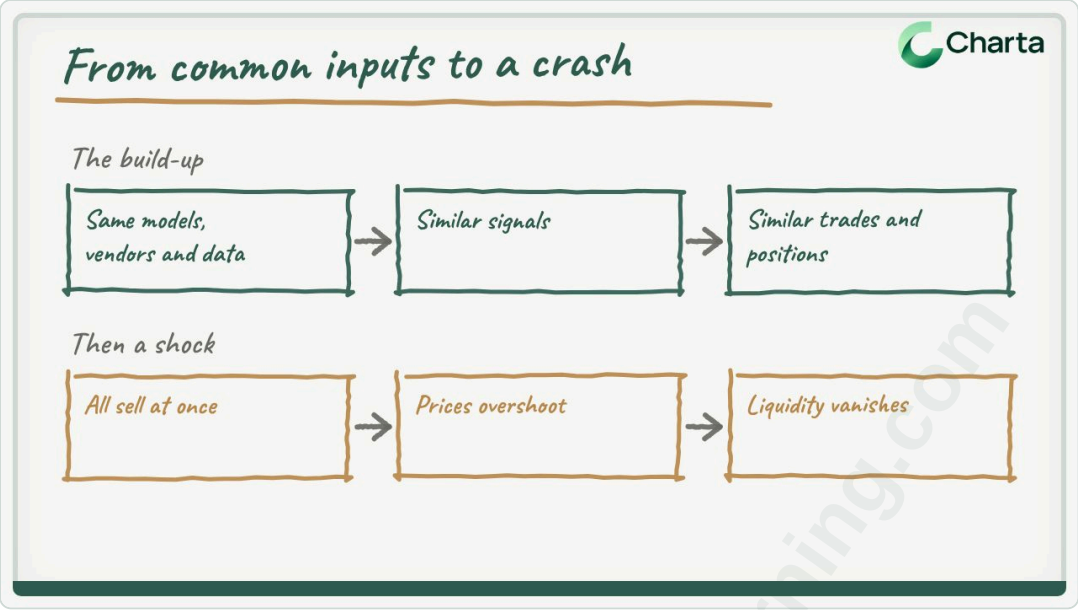

IN THE LESSON

- Here is the mechanism in two rows.
- The first row is the build up.
- Many firms use the same models, vendors and data.
- So their models produce similar signals: they interpret the same news the same way and see the same risks.
- Similar signals give similar trades and similar positions.
- The second row is what happens when a shock arrives.
- Because their models agree, the firms all sell at once.
- That selling is procyclical, meaning it adds to the move that is already under way instead of leaning against it.
- Prices overshoot, falling further than the news justifies.
- And liquidity vanishes, because the buyers who would normally step in are running the same models and are selling too.

13 Example 4: herding in the Calderra market

10 funds hold $100m each. Price falls 1% per $50m sold. A seller sells 20%

Each seller sells	20% × $100m	**$20m**
Diverse models: 2 sell	2 × 20 = 40; 40 ÷ 50	**0.8% fall**
Common model: 10 sell	10 × 20 = 200; 200 ÷ 50	**4% fall**
Same headline, bigger move	4 ÷ 0.8	**5 times**

IN THE LESSON

- Let us put numbers on it.
- Ten funds each hold one hundred million dollars of a Calderra bond, and in this market the price falls one per cent for every fifty million dollars sold.
- A bad headline arrives, and any fund whose model says sell will sell twenty per cent of its holding, twenty million dollars.
- First, diverse models.
- Each fund built its own, and only two of the ten say sell.
- They sell forty million dollars, and the price falls forty over fifty, zero point eight per cent.
- Now a common model.
- All ten funds use versions of the same foundation model, and all ten say sell.
- They sell two hundred million, and the price falls two hundred over fifty, four per cent.
- The same headline moves the price five times as much.

Charta Learning · charta-learning.com

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

14 The same headline, two price paths

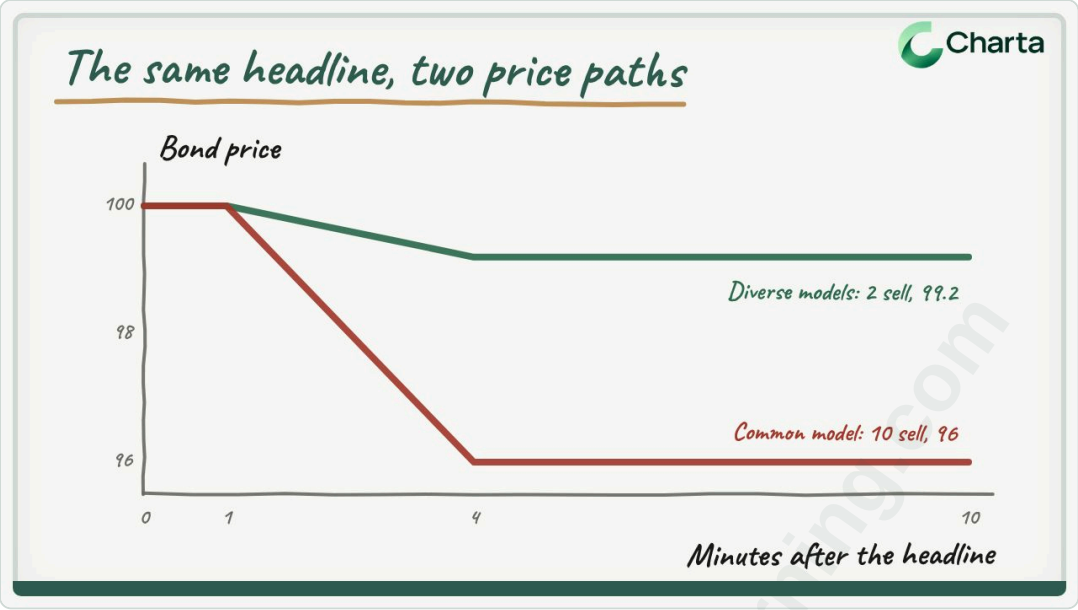

IN THE LESSON

- Here is the same example as a price path.
- Along the bottom is time, in minutes after the headline.
- Up the side is the bond's price, starting from one hundred.
- Selling starts after one minute and lasts three minutes.
- The green line is the market with diverse models.
- Two funds sell, and the price slips from one hundred to ninety nine point two, then holds.
- The red line is the market with a common model.
- All ten funds sell, and the price falls to ninety six.
- Both paths start at the same price and react to the same news.
- The only difference is how similar the models are.
- That is the exam point in one picture: common models and data create correlated selling and bigger price moves.

15 Speed, and liquidity that disappears

- Speed: hours become milliseconds
- A move can run before any human looks
- AI market makers trained on similar data
- They widen quotes or withdraw together
- Liquidity vanishes exactly when it is needed

IN THE LESSON

- Two features make this worse.
- The first is speed.
- Reactions that once took hours, as people went through the news and phoned their brokers, now happen in milliseconds, so a move can run its course before any human looks at it.
- The second is liquidity.
- Much market making, which means standing ready to quote prices at which you will buy and sell, is now done by algorithms, and increasingly by A I models.
- If those models are trained on similar data, they all see the same danger at the same moment.
- They widen their quotes, the gap between the price at which they buy and the price at which they sell, or they withdraw altogether.
- So liquidity can vanish exactly when it is needed.

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

16 Example 5: market makers pull back

Five AI market makers quote a Calderra bond, $10m each at the best price

Normal depth at best quotes	5 × $10m	**$50m**
Stress: 4 pull back	1 × $10m	**$10m**
Bid-ask spread (points)	0.10 → 0.60	**6 times wider**
A $30m sale at the quotes	10 ÷ 30	**only a third fills**

IN THE LESSON

- Five A I market makers quote the Calderra bond, each ready to trade ten million dollars at its best price.
- In normal times, the depth at the best quotes, the amount you can trade without moving the price, is five times ten, fifty million dollars.
- Now volatility jumps.
- Four of the five models were trained on similar data, and all four pull back.
- Only one maker still quotes, so the depth is ten million.
- Meanwhile the bid ask spread, the gap between the buying and selling price, widens from zero point one zero to zero point six zero price points, six times wider.
- A fund that must sell thirty million dollars finds only ten million at the quotes.
- Only a third fills; the rest must go at worse prices.

17 Two more behavioural risks

Beyond herding and vanishing liquidity

Risk	What happens	Why it is hard
Tacit collusion	Algorithms learn to keep spreads wide	No communication to find
False information	AI-made fake news or deepfakes	News-reading models act at once

IN THE LESSON

- Two more behavioural risks complete the picture.
- The first is tacit collusion.
- In research simulations, pricing algorithms that learn by trial and reward have discovered, on their own, that keeping spreads wide pays better than competing.
- They can end up behaving like a cartel without any communication at all, so there is no agreement for a competition authority to find.
- The second is false information.
- A I can now produce convincing fake news stories, documents or even video, known as deepfakes.
- Automated news reading systems, like Thornbury's sentiment model, may act on such a story before any human checks it, and the trades of one model become the news that the next model reads.
- A fake report can move a price in seconds.

Charta Learning · charta-learning.com

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

18 The exam point on market behaviour

- Error: sophisticated models mean diverse behaviour
- Common models and data make firms more alike
- Normal times: price discovery, liquidity, execution
- Stress: correlated behaviour and procyclicality
- Plus tacit collusion and false information

IN THE LESSON

- Here is the usual error.
- It is tempting to assume that more sophisticated models mean more diverse behaviour, because each model is cleverer and more finely tuned.
- Example four showed the opposite.
- When ten funds build on the same foundation model, the same vendor and the same data, their behaviour becomes more alike, not less.
- So remember three things.
- In normal times, A I can improve price discovery, liquidity and execution.
- In stress, common models and data create correlated behaviour and procyclicality, made faster by speed and worse by liquidity that disappears.
- And new behavioural risks, tacit collusion and false information, can arise with no human decision behind them.

19 Stability is a property of the system

- Stability: the system keeps lending, paying, pricing
- AI can create shared vulnerabilities
- Each firm's model can look sound alone
- Yet sound models can add up to correlated behaviour
- Judge risk across firms, not only firm by firm

IN THE LESSON

- Now the third outcome, financial stability.
- Financial stability means the financial system as a whole keeps doing its job, lending, paying and pricing risk, even when something goes wrong.
- Wider A I use matters for it because it can create shared vulnerabilities, weaknesses that many firms have at the same time.
- The important idea is that stability is a system property.
- Each firm's model can look sound when you examine it alone, validated, accurate and well governed.
- Yet if many firms run sound models on the same inputs, together they can still add up to correlated behaviour and a system wide problem.
- So risk must also be judged across firms, not only firm by firm.

20 Four channels to financial stability

Four main channels from AI to the whole system

Channel	How it spreads
Third-party concentration	Few providers serve many firms
Market correlation	Firms trade on similar signals
Cyber risk	AI-assisted attacks; data poisoning
Model risk, governance, data	Opacity; hallucination

IN THE LESSON

- There are four main channels.
- First, third party dependencies and concentration.
- A few providers of cloud services, data and foundation models serve many institutions, so a failure at one of them hits all of them at once.
- Second, market correlation, from many firms trading on similar signals, which we met in Example four.
- Third, cyber risk.
- Attackers can use A I to write convincing phishing messages or to fake a voice, and models themselves can be attacked.
- Data poisoning means corrupting the data a model is trained on, so that it learns the wrong thing.
- Fourth, model risk, governance and data quality, made harder by opacity and by hallucination in generative models.

Charta Learning · charta-learning.com

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

21 Example 6: one provider goes down

10 firms trade $20bn a day in Calderra. 7 run their models on one provider, Nimbra Cloud

Firms on Nimbra Cloud	Thornbury included	**7 of 10**
Their daily trading	of $20bn in all	**$13bn**
Their share of the market	13 ÷ 20	**65%**
Lost in a half-day outage	13 × 0.5	**$6.5bn**

Each firm saw one vendor; the system had one point of failure

IN THE LESSON

- Ten firms trade in the Calderra market.
- Seven of them, Thornbury included, run their trading models on the same A I cloud provider, a made up firm called Nimbra Cloud.
- Those seven firms do thirteen billion dollars of the market's twenty billion dollars of daily trading.
- So their share of the market is thirteen over twenty, sixty five per cent.
- One day Nimbra Cloud suffers an outage for half the trading day.
- All seven firms stop at once, and thirteen billion times one half, six point five billion dollars of trading, does not happen.
- Each firm had checked its own vendor and saw one supplier among many.
- Only from above could anyone see one single point of failure for most of the market.

22 Example 7: match each scenario to its channel

Thornbury's risk team sorts five scenarios

Scenario	Channel
Cloud outage halts 7 firms	Third-party
Funds sell on the same signal	Correlation
Fake entries in a vendor's dataset	Cyber: poisoning
Summary invents a covenant	Model risk
Deepfake voice orders a payment	Cyber: attack

IN THE LESSON

- Thornbury's risk team sorts five scenarios, matching each to its channel.
- First, a cloud outage halts seven firms at once.
- That is third party concentration.
- Second, funds all sell on the same model signal.
- That is market correlation.
- Third, fake entries are slipped into a dataset that a vendor sells to many firms, so every model trained on it learns the wrong thing.
- That is cyber risk, through data poisoning.
- Fourth, a generative summary invents a covenant, a promise in a loan agreement, that is not in the document.
- That is model risk, through hallucination.
- Fifth, a deepfake voice of Thornbury's finance chief orders an urgent payment.
- That is cyber risk again, an A I assisted attack.

Charta Learning · charta-learning.com

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

23 Example 8: the mistake of judging firm by firm

Same market as Example 4. Each fund tests selling $20m alone

Alone: price falls	20 ÷ 50	**0.4%**
Loss on the $80m kept	0.4% × 80	**$0.32m**
All ten sell: price falls	200 ÷ 50	**4%**
Loss on the $80m kept	4% × 80	**$3.2m**
Across ten funds	10 × 3.2	**$32m**

IN THE LESSON

- Same market as Example four.
- Each fund's own stress test asks what happens if it sells twenty million dollars of the bond.
- Alone, that moves the price twenty over fifty, zero point four per cent.
- On the eighty million dollars it keeps, the loss is zero point four per cent of eighty, zero point three two million dollars, or three hundred and twenty thousand.
- The test passes.
- But every fund runs the same common model, so all ten sell together, two hundred million dollars.
- The price falls two hundred over fifty, four per cent, and the loss on the eighty million each keeps is three point two million, ten times the firm level estimate.
- Across ten funds, that is thirty two million dollars.
- Every firm level test was right on its own assumption; the system was not.

24 The exam point on financial stability

- Shared vulnerabilities, even if each model looks sound
- Four channels: third-party, correlation, cyber, model
- Match each scenario to its channel
- Error: judging risk only firm by firm
- Provider concentration: operational and systemic risk

IN THE LESSON

- So the exam point for this outcome.
- Wider A I use can create shared vulnerabilities across the system, even when each firm's model looks sound.
- Know the four channels: third party dependencies and concentration, market correlation, cyber risk, and model risk with governance and data quality.
- In a question, match each scenario to its channel, as in Example seven.
- Remember the usual mistake, judging risk only firm by firm, which Example eight showed can miss a loss ten times larger.
- And note one classification carefully.
- Concentration in a few providers of cloud, data and models is an operational and systemic risk, not a market risk.
- It comes from failures in services, not from price moves.

Charta Learning · charta-learning.com

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

25 Policy for firms: start from the rules you have

- Two sides: regulated firms, and supervisors themselves
- Existing rules are mostly technology-neutral
- Review them for gaps first; AI-specific rules after
- Governance with clear senior accountability
- Risk management proportionate to materiality

IN THE LESSON

- The last outcome is recommended regulatory practice, and it has two sides: advice on how regulated firms should use A I, and advice on how supervisors should use it themselves.
- We start with the firms.
- The first recommendation is about the rule book.
- Most existing rules are technology neutral, meaning they apply whatever tool a firm uses: a trading rule binds a human trader and an algorithm alike.
- So supervisors should first review existing rules for gaps, and only then write A I specific rules where gaps are found.
- Banning A I is not the advice.
- Next, firms need governance with clear senior accountability, a named senior person answerable for each use.
- And risk management should be proportionate to the materiality of each use.

26 Proportionality: controls follow materiality

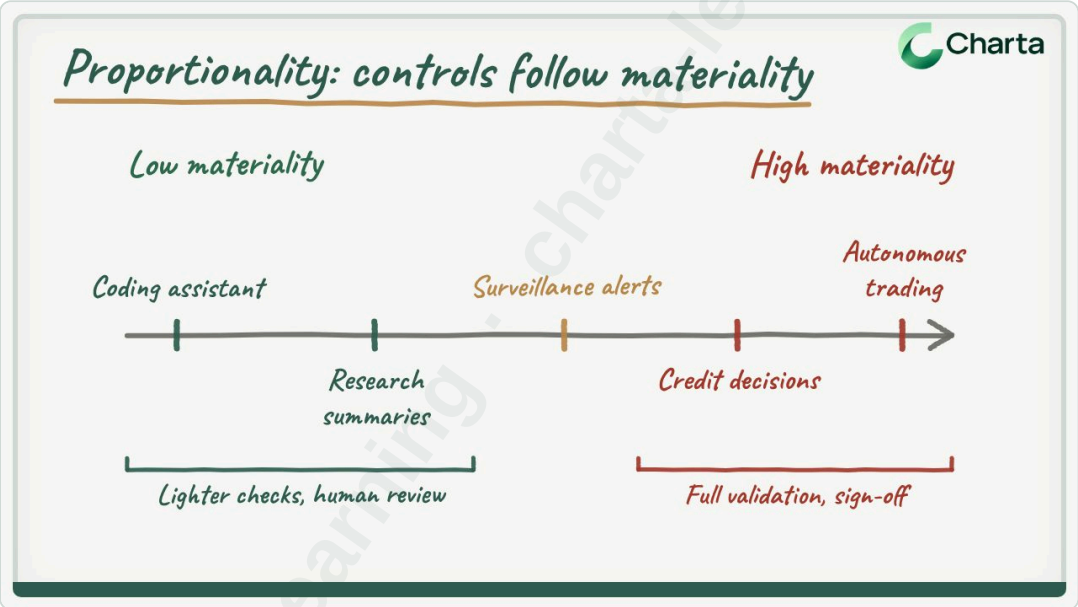

IN THE LESSON

- Proportionality means the strength of the controls matches the materiality of the use, how much harm it could do if it went wrong.
- Here are Thornbury's uses on a line from low to high materiality.
- At the low end is the coding assistant, whose output a developer tests before anything runs.
- Next come research summaries, checked by an analyst before use.
- In the middle are surveillance alerts, which feed decisions about possible abuse.
- Higher up are credit decisions about real clients.
- And at the top is autonomous trading, where a model acts in the market with no human in between.
- Below the line, the low end gets lighter checks and human review, and the high end gets full validation and senior sign off.

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

27 Controls every material use needs

- Validation before use; ongoing monitoring after
- Sound data quality
- Explainability and human oversight
- Third parties: due diligence and exit plans
- Outsourcing never transfers responsibility

IN THE LESSON

- What should the controls contain?
- First, model validation before use and ongoing monitoring after, because a model that worked last year can drift as markets change.
- Second, sound data quality, since a model is only as good as the data it learns from.
- Third, explainability, meaning the firm can say why the model produced a given output, and human oversight, a person who can question and overrule it.
- Fourth, third party risk management: due diligence before choosing a provider, and an exit plan, a tested way to move to another provider or bring the work in house.
- And fifth, a principle that runs through all of it.
- Outsourcing never transfers responsibility for outcomes.
- The firm remains answerable.

28 Example 9: Thornbury tiers its controls

Thornbury's 8 AI tools, by materiality. Validation effort in hours a year

High: 2 tools	2 × 200	**400**
Medium: 3 tools	3 × 60	**180**
Low: 3 tools	3 × 10	**30**
Proportionate total	400 + 180 + 30	**610**
Full checks on all 8	8 × 200	**1,600**

IN THE LESSON

- Take the eight A I tools from Example three, sort them by materiality, and set validation effort in hours a year.
- Two are high materiality, the execution algorithm and the credit score, at two hundred hours each: four hundred hours.
- Three are medium, such as surveillance, at sixty hours each: one hundred and eighty hours.
- Three are low, such as the coding assistant, at ten hours each: thirty hours.
- The proportionate total is six hundred and ten hours.
- The alternative, full checks on all eight at two hundred hours each, would take one thousand six hundred hours, much of it spent on tools that could do little harm.
- Proportionality puts the effort where the risk is, without leaving any tool unchecked.

29 Example 10: the vendor is not responsible

Thornbury buys a credit model from a vendor. A flaw causes an $8m loss

Loss from the vendor's model	bad lending decisions	**$8m**
Vendor's liability cap	in the contract	**$1m**
Thornbury bears	8 − 1	**$7m**
Answerable to regulator	and to clients	**Thornbury**

The work was outsourced; the responsibility stayed

IN THE LESSON

- Now a classic mistake.
- Thornbury buys a credit scoring model from a vendor and assumes the vendor carries the risk.
- A flaw in the vendor's model leads to bad lending decisions and a loss of eight million dollars.
- The contract caps the vendor's liability at one million dollars.
- So Thornbury bears eight minus one, seven million dollars itself.
- And when the regulator asks who is answerable to it and to the clients, the answer is not the vendor.
- It is Thornbury, and the senior manager accountable for that use.
- Due diligence and an exit plan were Thornbury's job, not the vendor's.
- The footnote is the lesson: the work was outsourced, but the responsibility stayed.

Charta Learning · charta-learning.com

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

30 Supervisors: skills, data, and their own AI

- Build skills and data
- Close gaps on AI use and provider dependencies
- Use AI to analyse data and detect misconduct
- Human review, validation and explainability
- So every supervisory decision can be justified

IN THE LESSON

- The second side is the supervisors themselves.
- First, they should build skills, hiring and training people who understand A I, and build their data.
- Second, they should close the gaps in their information: which firms use A I, for what, and which providers of cloud, data and models they depend on.
- Without that, nobody can see a concentration like Nimbra Cloud.
- Third, supervisors should use A I themselves, to analyse large data sets and to detect misconduct such as market abuse, which is often called suptech, short for supervisory technology.
- Fourth, they must hold themselves to the standard they set for firms: human review, validation and explainability, so that every supervisory decision can be justified.

31 Example 11: the supervisor's own model

A made-up regulator scans Calderra trade reports for market abuse

Trade reports scanned	each month	**3,000,000**
Flagged by the model	unusual patterns	**600**
Reviewed by staff	all of them	**600**
Inquiries opened	30 ÷ 600	**30 (5%)**

No action on the model's word alone; every decision explained

IN THE LESSON

- The Calderra securities authority, a made up regulator, applies a machine learning model to three million trade reports a month, looking for signs of market abuse.
- The model flags six hundred unusual patterns.
- Staff review all six hundred, because the model ranks cases but does not decide them.
- After review, thirty inquiries are opened, thirty out of six hundred, five per cent.
- The footnote is the rule.
- No action is taken on the model's word alone.
- For every inquiry the authority records the evidence and the reasons, so that if a firm challenges it, the decision can be explained and justified.
- That is human review and explainability, applied by the supervisor to itself.

32 Mistakes to avoid

- Treating all AI as generative AI
- Assuming sophisticated models mean diverse behaviour
- Judging risk only firm by firm
- Calling provider concentration a market risk
- Banning AI, or thinking a vendor carries the blame

IN THE LESSON

- Before the quiz, the mistakes to avoid across the topic.
- Do not treat all A I as generative A I.
- Most of it is conventional machine learning, with different risks.
- Do not assume more sophisticated models mean more diverse behaviour.
- Common models and data make firms more alike.
- Do not judge risk only firm by firm, because sound models can add up to a system wide problem.
- Do not call provider concentration a market risk: it is operational and systemic.
- And do not think regulators should ban A I, or that a vendor carries the firm's responsibility.
- The advice is proportionality and human accountability.

Charta Learning · charta-learning.com

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

33 Try three quick questions

1. AI groups trades and flags those unlike the rest. Which kind? Is it GenAI?
   Answer: Unsupervised learning: conventional ML, not GenAI.
2. 8 of 10 firms use one cloud and model provider, and it fails. Channel?
   Answer: Third-party concentration: operational, systemic.
3. Each of 10 funds sells $15m; price falls 1% per $50m. One alone? All ten?
   Answer: Alone 0.3%. All ten sell $150m: 3%.

IN THE LESSON

- Try three quick questions, then listen to the answers.
- Question one.
- A risk team uses A I that groups trades and flags the ones unlike the rest.
- Which kind of learning is it, and is it generative A I?
- It is unsupervised learning, conventional machine learning, not generative A I.
- Question two.
- Eight of ten firms run on one cloud and model provider, and it fails.
- Which channel is that, and what kind of risk?
- Third party concentration, an operational and systemic risk, not a market risk.
- Question three.
- Ten funds each sell fifteen million dollars, and the price falls one per cent for every fifty million sold.
- What is the fall if one fund sells alone, and if all ten sell?
- Alone, zero point three per cent.
- All ten sell one hundred and fifty million, so three per cent.

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.