



Clare Kitching

Transform your AI & data ambition into action | xQuantumBlack, xMcKinsey | Global top 100 Innovators in Data & Analytics | AI & data strategy, governance and capability building

54k followers • [View Profile](#)

7

POSTS / WEEK

9 AM, UTC

MAIN POSTING TIME

600

AVG REACTIONS
275 comments

100%

CALL-TO-ACTION
FREQUENCY

#1 SUCCESS FACTOR

She turns confusing, high-stakes AI topics into simple "truth + picture" lessons that busy leaders can instantly understand, save, and share. Her profile promises the same thing (turn ambition into action), and her posts consistently end with a low-friction next step (follow or newsletter), so the trust built by the frameworks naturally turns into an owned audience and inbound demand.

FORMAT BREAKDOWN



■ Image **100%** 600 avg reactions

7 Tactics You Can Steal

01 [Start with a bold "everyone wants this, few do that" line](#)

02 Turn the lesson into one simple picture people can reuse

03 Explain big topics with a ladder of levels, one step at a time

04 Use "what people think" vs "what it really is" to kill confusion

05 End with one easy question that makes people pick a side

06 Open with a scary risk, then calm people with a clear plan

07 Sell softly: give the full lesson, then one small next step

7 Tactics You Can Steal

1 Start with a bold “everyone wants this, few do that” line

She often opens with a blunt contrast that calls out the gap between what people want and what they avoid. It works because readers instantly pick a side and feel the tension, then they keep reading to see the explanation. This style shows up in several of her highest performers; the “foundations before AI” version averages about 2.42x typical engagement, and her best “under the hood” explainer reached about 3.70x. To copy it: write two short sentences that clash (want vs do), then spend the rest of the post proving the point with a simple list or diagram.

[Example 1 →](#)

[Example 2 →](#)

[Example 3 →](#)

2 Turn the lesson into one simple picture people can reuse

Every single post in this set is an image, usually a diagram or comparison chart, so the content is built to be saved and shared. It works because a picture can be dropped into a slide deck or team chat without needing extra context. With medians around 598 likes and 275 comments, this “shareable asset” approach clearly sustains daily posting volume. To copy it: take one idea, reduce it to a single visual (two columns, four boxes, a ladder, a Venn diagram), and make the labels readable on a phone.

[Example 1 →](#)

[Example 2 →](#)

[Example 3 →](#)

3 Explain big topics with a ladder of levels, one step at a time

She frequently teaches by moving up a clear ladder (level one, then level two, then level three), so readers always know where they are. This works because it turns a scary, messy subject into a calm path with “what changes at each step,” which makes leaders feel in control. Her “stack/layers” explainers average about 2.05x typical engagement across 3 posts, and the first-time version was far stronger than later repeats (about 3.70x vs about 1.81x and 0.64x). To copy it: pick 3–5 levels, name each level in plain words, and add one sentence per level that says what new work, risk, or cost shows up.

[Example 1 →](#)

[Example 2 →](#)

[Example 3 →](#)

4

Use “what people think” vs “what it really is” to kill confusion

She often uses a simple two-part structure: first list the common wrong idea, then replace it with the real checklist. It works because it respects what the reader already believes (“you’re not stupid”), but still gives them a cleaner way to explain it to others. Governance posts as a whole average about 1.29x typical engagement, and the clearest “split the buckets” governance framing averages about 1.81x. To copy it: write 3–5 bullets under “people think,” then 5–8 bullets under “it’s really about,” and make the second list practical enough to use in a meeting.

[Example 1 →](#)[Example 2 →](#)[Example 3 →](#)

5

End with one easy question that makes people pick a side

Most posts finish with a question that’s simple to answer without writing an essay. It works because it turns a lecture into a conversation, and it lets readers share their situation (which naturally creates long comment threads). In this set, engagement questions show up in 94 out of 100 posts, and the median comment count is 275—so comments are a core engine, not a bonus. To copy it: ask a forced-choice question (A/B/C), a time-bound test (“could you do this in 10 seconds?”), or a “what would you do first?” question.

[Example 1 →](#)[Example 2 →](#)[Example 3 →](#)

6

Open with a scary risk, then calm people with a clear plan

Some of her strongest governance posts start by making the risk feel real (what could go wrong, who gets blamed, what breaks), then they immediately give a short plan to fix it. This works because the fear grabs attention, but the plan earns trust and prevents the post from feeling like doom. Governance content is already an above-average performer here (about 1.29x), and the best governance framing is the kind that gives leaders something specific to build (about 1.81x). To copy it: start with a one-sentence risk, then offer a numbered list or “layers” plan that tells the reader what to put in place next week.

[Example 1 →](#)[Example 2 →](#)[Example 3 →](#)

7

Sell softly: give the full lesson, then one small next step

The posts are built to stand on their own, then the “ask” shows up as one small line at the end. This works because people don’t feel tricked—value comes first—so the follow, repost, or newsletter link feels like a natural continuation. Newsletter links appear in 68 out of 100 posts, and the engagement difference is small (about 1.14 with a newsletter link vs 1.10 without), which suggests the real benefit is long-term audience building, not short-term reach. To copy it: teach first, then add one simple CTA (follow, subscribe, download) and keep it optional and friendly.

[Example 1 →](#)[Example 2 →](#)[Example 3 →](#)[Profile →](#)

Profile Breakdown

HEADLINE FORMULA

Transform your AI & data ambition into action | xQuantumBlack, xMcKinsey | Global top 100 Innovators in Data & Analytics | AI & data strategy, governance and capability building

[Clear outcome you deliver in plain language] | x[1-2 trusted brands/roles] | [one strong third-party proof point like an award/ranking] | [3-5 buyer search keywords]. Why it works: the outcome comes first (so executives self-select fast), the credibility stack reduces perceived risk, and the keyword cluster makes the profile easy to find for common needs (strategy, governance, capability building).

BANNER STRATEGY

CAMBIQ

Turn AI & data ambition into action.

Follow for insights on AI, data and making it matter.



Brand-style banner with one core promise (“turn AI & data ambition into action”) + a simple top-of-funnel CTA (“Follow for insights”) + credibility logos (recognisable employers/partners). Why it works: it repeats the same promise as the headline, builds instant trust with visual proof, and asks for the easiest next step on a profile visit (follow) instead of pushing a hard sell.

Content Strategy

PILLAR	%	WHAT CRUSHES	EXAMPLES
Execution-first AI (foundations, strategy, problem framing, adoption)	45%	<p>The strongest posts in this bucket are “foundations before shiny AI” reality checks: a small set of these posts averages about 2.42x typical engagement. Problem-first leadership checklists also work well because they give leaders meeting language they can reuse immediately. What underperforms here is motivational behaviour-change content (habits/fluency/fears): 5 posts in that style average ~0.57x, suggesting this audience responds more to practical operating guidance than encouragement.</p>	Post 1 Post 2 Post 3
AI concepts and mental models (definitions, stacks, agents, taxonomy)	26%	<p>The standout is “layers/stack” explaining (turning a messy topic into a step-by-step ladder): 3 posts of this type average ~2.05x, driven by a first-time version that hit ~3.70x, while later near-identical reposts dropped to ~1.81x and ~0.64x. That pattern says novelty matters: the first clear diagram wins, repeats decay fast. More generic overviews (“AI is like a toolbox/vehicle spectrum”) are weaker: 3 posts average ~0.90x.</p>	Post 1 Post 2 Post 3
Governance, risk, and decision accountability (AI + data)	18%	<p>Governance is a reliable outperformer overall: 18 posts average ~1.29x typical engagement. The clearest winning frame is splitting “inputs” (data) from “outcomes/decisions” (AI): 4 posts using that idea average ~1.81x, because it helps leaders quickly understand what’s missing. Softer ‘governance encouragement/truths’ content is weaker: 4 posts average ~0.86x, so the audience seems to reward concrete ownership, monitoring, and decision-focused checklists over general reminders.</p>	Post 1 Post 2 Post 3
Curated learning resources (courses,	11%	<p>Practical course roundups win: 3 posts average ~1.60x, especially when the list is easy to act on immediately. Report reading lists are much weaker (5 posts average ~0.86x), and meta-</p>	Post 1 Post 2 Post 3

PILLAR

%

WHAT CRUSHES

EXAMPLES

roundups, reading lists)

learning posts underperform too (6 posts average ~0.78x). The difference is "do this next" utility: courses feel like an instant win, while long reading paths feel like homework.

Conversion Strategy

Their Offer(s)

Cambiq's B2B advisory and delivery support to turn AI & data ambition into action (strategy, governance/responsible AI, capability building/adoption, and program setup/turnaround).

Newsletter: "RETURN ON INTELLIGENCE" (fortnightly) via featured signup link

Executive education sessions for leaders/teams

Speaking, advisory, and interim executive support

How They Promote in Posts

CTAs are effectively universal: all posts include at least one CTA, and about 99% include explicit end-of-post CTA lines. The most common pattern is value first (a complete diagram/lesson), then an end block with an easy engagement question (94/100 posts), plus "follow" and "repost" prompts, plus a soft newsletter link (68/100 posts). Conversion asks are usually low-friction and placed at the end; newsletter-link posts average ~1.14 relative engagement vs ~1.10 without the newsletter link (a small difference), which fits a strategy focused on compounding trust and owned audience over time.

[Example 1 →](#)

[Example 2 →](#)

[Example 3 →](#)

How They Promote in Profile

The profile acts like a clean funnel: headline and banner repeat the same outcome promise and stack credibility, the banner asks for a simple follow, the Featured section makes the newsletter the main click (owned audience), and the About section reads like a short services page with an offer ladder (education → advisory/delivery → interim) and a direct contact path (hello@cambiq.com).

[View Profile →](#)

Top 10 Posts

1



Everyone talks about AI models.

A clear ladder explains how real-world AI often moves from basic models to retrieval systems, then to tools that can take actions, then to bigger multi-part systems. It ties each step to what leaders actually need to worry about: ownership, trust, monitoring, and governance as autonomy increases.

♡ 2725 💬 341 ↻ 381

2




Everyone wants the AI lift.

A strong metaphor shows that many organisations queue for shiny AI solutions while avoiding the less glamorous foundation work. It lays out a simple set of foundations—data basics, workflow redesign, and clear accountability—and asks readers to choose what they'd invest in first.

♡ 2270 💬 449 ↻ 168

3



The fastest way to understand AI governance?

A quick reframe separates governing data inputs from governing AI outcomes and decisions. It explains why AI governance rarely has one “natural owner” and highlights the need to coordinate across functions that share responsibility for impact.

♡ 1823 💬 404 ↻ 254

4



Everyone wants AI magic.

A hype-vs-reality message calls out the messy foundations that block AI value, like inconsistent definitions, scattered data, and unclear ownership. It ends with a short, practical list of “boring” basics leaders need before expecting AI outcomes.

♡ 1889 💬 404 ↻ 157

5



McKinsey taught me that brilliant people fail when they rush to answer the wrong question.

A leadership lesson warns that teams often jump to “add AI” before agreeing on the real problem they're solving. It offers a short checklist to test whether the problem is clear, outcome-based, and worth solving even without AI.

♡ 1542 💬 540 ↻ 131

6



Microsoft just opened the door to AI mastery.

A curated list of free AI courses is organised by what different readers want to do: learn basics, apply tools, lead teams, or build solutions. It uses a clear “save this” angle and asks readers to pick which course they'll start.

♡ 1454 💬 275 ↻ 274

7




46% of employees are using AI tools their company never approved.

A set of risk and adoption statistics highlights how common unapproved AI use has become. It then lays out a four-layer way to build guardrails, moving from policy to real behaviour and technical controls.

♡ 1331 💬 467 ↻ 204

8



Governance breaks down the moment no one owns the decision.

This post argues that governance fails when “shared responsibility” means no named decision owner. It uses a high-stakes moment—being asked to explain harm to a regulator, customer, or executive—to show why clarity matters before something goes wrong.

♡ 1225 💬 453 ↻ 131

9




Behind every confident decision

A myth-vs-reality structure reframes governance as a way to make faster, safer “yes” decisions—not bureaucracy. It lists practical benefits like shared definitions, trustworthy quality, access, and security that enable confident action.

♡ 1212 💬 400 ↻ 156

10



Think your data governance covers AI risks?

A challenge hook explains why governing data inputs is necessary but not enough when AI affects real decisions and behaviour. It lists what AI governance needs to include, like ownership of outcomes, monitoring after launch, and the ability for humans to intervene.

♡ 1253 💬 328 ↻ 175