

Strategy Mapping: using AI to do the Legwork

A whitepaper on the high-level steps to build a strategy map faster and sharper, using AI as the research and pressure-testing engine, not the author of the strategy itself.

The job I actually burn AI's time on

Last time I attended a strategic planning session, people were discussing next year's budget and which initiatives to fund. This session wasn't strategic planning; it was about establishing financial controls. According to Dr Roger Martin, strategy is the set of choices about where to compete and how to win there, deliberately made at the cost of other options you could have chosen instead. Sadly, this confusion between strategic objective setting and financial planning is common across many parts of UK Plc. In these circumstances, the organisational strategy is thought to be implicit. Everyone is assumed to know the enterprise's objective and purpose. Yet if this assumption were true, why would we have a strategy execution gap? And why would portfolio planners, IT and Finance Directors talk about aligning investments with strategy? Furthermore, why do we have AI initiatives which do not increase organisational performance?

As an AI consultant, I have been looking at applications which go beyond a Chatbot or Robotic Process automation. I have been exploring a Digital Twin to build enterprise agility, but I've hit a wall of unclear strategic objectives. As an ex-Boy Scout and Mountain Leader, I'm used to using a map to create direction. I therefore set about creating a map using AI to give senior leaders, planners, and employees a definition of organisational strategy, using a combination of internal and external documentation and established strategic tools. This map can be used to sense the need for directional change and to identify where the organisation needs to adapt.

This paper sets out, at a high level, how I use AI across the process of building a strategy map: the Strategy Choice Cascade, the Business Model Canvas, the Wardley Map, the Value Proposition Canvas, the Reason-to-Buy Map, the Operating Model Canvas, the Value Creation Model and the Brand Equity Pyramid. It isn't a paper about handing strategy to AI. It's a paper about where AI earns its place in helping leaders build the map and, based on organisational strategy, where it has to stay firmly on the sidelines.

The job AI is actually good at here.

Strategy mapping has two very different kinds of work in it, and it's worth being honest about which is which. There's the judgement work: deciding where to play, arguing about how to win, naming the one thing AI itself must never be allowed to touch. That's an AI search of documents validated by a room full of people with skin in the outcome, and no model does it for them. Then there's the legwork underneath that judgement: pulling together what's already known about a market, a competitor, a customer segment, drafting a first pass at a canvas so the room isn't staring at a blank

page, checking whether a claim someone just made in the room actually holds up against the evidence sitting in a shared drive somewhere.

That second category is where AI belongs. It's fast at synthesis, tireless at cross-referencing, and it doesn't get precious about a first draft being wrong, which matters more than it sounds. A blank canvas intimidates a room into vague answers. A flawed first draft, put up on the screen and torn apart, gets you a sharper answer faster, because arguing with something specific is always easier than inventing something from nothing. Furthermore, AI can scan the marketplace, identify and analyse competitors' actions, evaluate customers' actions, and outline the need for adaptation.

The high-level steps

1. Feed it the baseline before the workshop, not during it

Before I open a Business Model Canvas or a Wardley Map with a client, I point AI at whatever the organisation already has: previous strategy decks, board papers, customer research, competitor filings, analyst notes, transcripts of earlier conversations. I ask for a first-cut draft of the canvas, built strictly from that material, with every claim flagged back to its source. Nothing in that draft is trusted. All of it is useful, because it means the workshop opens with something to react to instead of something to invent.

2. Use it to widen the room's view of the landscape, then narrow it yourself

On the Wardley Map in particular, AI is genuinely useful for surfacing where a component actually sits on the evolution axis: what's still genuinely novel, what's already been commoditised by three vendors the room hasn't heard of yet. I ask it to research the landscape broadly, then I bring that research into the room as an input to argue about, never as the map itself. The room still decides what's worth building in-house and what's worth buying off the shelf. AI makes sure that decision isn't being made on a partial view of what's actually out there.

3. Let it draft the Reason-to-Buy chain, then make a real buyer test it

Based upon the definition of the corporate brand, the Reason-to-Buy Map is the one place in the whole sequence where a claim has to survive contact with a specific, named alternative a specific, named buyer would otherwise choose. AI is a decent first pass at drafting that chain, especially where a Value Proposition Canvas already exists to draw the pains and gains from. But I never let a drafted Reason-to-Buy claim stand unchallenged. I put it back to the room with one test: would a real buyer in this segment actually say this, in these words, about why they chose us? If nobody in the room believes a real customer would say it, it's not a reason to buy. It's a hope, dressed up in AI's confident prose, and that's exactly where this goes wrong if you let it stand.

4. Use it to check the map for gaps, not to fill them silently

Once you have a first pass at the map, I ask AI to check it for internal consistency: does every Revenue Stream on the Business Model Canvas trace back to a named Value Proposition? Does every Reason-to-Buy claim have a Proof that actually points back to something real in the Value Proposition Canvas? Does the Value Creation Model's delivery stream look, on paper, capable of

honouring what marketing is promising? These are exactly the questions a tired room stops asking by hour three of a workshop, and exactly the questions AI never gets tired of asking. What it flags, a person decides. It doesn't get to patch a gap itself quietly.

5. Put it to work as the Sensing Loop once the map is live

A strategy map is a snapshot, and it starts decaying the moment the workshop ends. This is where AI does its most valuable ongoing work, not in the room building the map but afterwards, watching it. Behavioural data, competitive movement, unstructured market signals, macro shifts: AI is well suited to scanning all of that continuously and flagging when a Reason-to-Buy claim looks like it's drifting out of date, or when a Wardley component has moved further toward commodity than the map currently shows. It escalates a signal to a named person. It doesn't get to rewrite the Segment Grid on its own authority, no matter how confident the signal looks.

6. Name, once, what AI is never allowed to decide on this map

This is the step that has to happen before any of the above, not after. Somewhere in How to Win, before AI touches a single canvas, I ask the client to name the one thing about the way they win that AI must never be allowed to quietly undermine the organisation, whether that's a fairness commitment, a brand promise, or the visible presence of a human expert at the point a customer actually experiences the business. That line goes into the cascade in the client's own words. Every use of AI in every step above sits underneath it, not the other way round.

Why this order matters

Get this backwards, and the workshop suffers in a specific, predictable way. Hand AI the whole first draft of a strategy map with no boundary named first, and the room spends its limited time either rubber-stamping something that sounds plausible but is never actually tested, or throwing it out and starting again because nobody trusts where it came from. Neither outcome is worth the time saved.

Do it the other way round: boundary named first, AI doing the legwork under it, a person testing every claim that reaches the page, and the room gets back exactly what a workshop is for: time spent arguing about the choices that decide whether the strategy actually wins, instead of time spent hunting for a market figure that should have taken two minutes to find. The map that comes out the other end is still entirely the room's, built faster, against a wider set of evidence than any one person in the room could have gathered alone, and never once handed over to something that can't be held accountable for the answer.

If you're building a strategy map and want a hand getting AI doing the useful part of that work without letting it anywhere near the choices that are actually yours to make, I'd be glad to talk it through.

For further information regarding Strategy Mapping, please contact [Beneficial Consulting](http://www.beneficialconsulting.co.uk).

Jonathan Ward



Beneficial Consulting Ltd
36-38 Wigmore Street,
London
W1U 2BP

Web: <http://www.beneficialconsulting.co.uk>
Mobile Tel: +44 (0) 7802 884598