

FROM “FIND US SOME AI” TO AN “EXTREMELY VALUABLE” TOOL THEY COULD KEEP

How a subject-matter leader at a regulated, national consumer mortgage lender turned an open-ended mandate into a compliant, resilient, value-producing system in about four months.

Mandates to find AI opportunities arrive without necessary objectives, metrics, funding, and ownership. Organizations don't realize these words don't mean what they used to. Most AI that gets built anyway delivers little or no value.

MODELED FROM THE LENDER'S OWN ECONOMICS, ON THE CONSERVATIVE END OF THE RANGE

\$150,000

DEVELOPMENT SPEND AVOIDED

Twelve months at the \$12,500 per month the lender was on track to spend, using the shortest waste window rather than the longest.

\$243,000

**TOTAL BASE-CASE EXPOSURE
AVOIDED**

The development spend above, plus the loan production the leader was pulling off their own book while researching, valued on industry averages.

1,390%

TWO-YEAR ROI, OR 13.9X

From the lender's own two-year model, corrected, and net of the \$160,000 of cost the model carries over that period.

AT A GLANCE — CHALLENGES

- An open-ended mandate to find AI opportunities, with no metrics, funding, or named owner
- A subject-matter leader, not a technologist, asked to design and build
- Roughly twenty hours of self-directed research already spent, with no clearer path
- Twelve to eighteen months of development ahead that would look productive and return nothing

AT A GLANCE — GOALS

- Establish whether the mandate was worth pursuing at all
- Find the narrowest, highest-value use case
- Clarify where the sponsoring VP actually stood on ownership and expectations
- Build something maintainable in-house, without outside help
- Iterate into ROI by managing risk

CHALLENGE

A VP at a national consumer mortgage lender handed one of their industry subject-matter leaders an opportunity: Find ways to leverage AI for the business. The leader, tech savvy but not a technologist, put twenty hours into articles and videos working out how to design and build the impressive capabilities they envisioned. They started to see how broad the task really was. Their questions grew faster than their answers.

A mutual acquaintance referred the leader to Ben Hamilton. The leader brought questions about models and architecture, but Ben read the situation as a risk

pattern he sees often: the directive, while high-visibility and cutting-edge, was too ambitious and too vague to be successful. It included no metrics, no targets, no major funding, and only nominal executive sponsorship. AI systems that get built under these constraints look functional and productive but deliver no or negative value. The challenge, to Ben, was to help the subject matter leader know what their operational realities were, how to build AI efficiently to show them and their executives what AI could do under their constraints, in a way that looks successful, *is* successful, and justifies off ramps or additional investment.

SOLUTION

Ben started instead with a ninety-minute conversation about what the field can and cannot currently do, from an engineer's perspective. He "was able to quickly point out the potholes that are overlooked with newbie enthusiasm to AI" (their words), so they could tell a real opportunity from an expensive diversion. The Client and Ben partnered to shape the boundaries of the AI work, with Ben threading the Leader's deep domain expertise and organizational understanding with Ben's engineering and product expertise. The leader narrowed their big vision to a single, high-value use case on a single workflow, with quantitative and qualitative measures of success developed by the pair.

Hamilton also prepped the client to "vibe check" (their words) the VP's actual level of engagement, expectations and ownership. This confirmed what

Hamilton expected, so Ben advised the client on moving forward with confidence, finally prepping them on the project narrative and questions they would encounter while they were executing.

Over the following months, Ben advised the client on numerous AI failure modes, on realistic expectations for output quality and improvements, on communicating imperfect technology development to executives and users, and on using the lender's regulatory constraints to simplify how much to automate and where to interface with people. Ben also required that whatever got built had to be simple enough for the leader to maintain with their existing level of resources – even if that meant just the leader – to encourage them to make difficult trade-offs to align with their reality.

RESULT

“The scope and the approach drastically changed”, the leader said. What went live about four months later was a retrieval-augmented generation (RAG) chatbot for a single, mid-size team of employees. It was deliberately focused, targeting real needs, and maintainable by the person who owned it, matched to their operational reality.

Efficiency and relief before the build. The SME reported that working with Ben saved *them* ten hours a week for six months doing prep work, with the frustration they avoided during that time worth at least as much as the hours. They put the time back into revenue-generating work.

“ He likely saved me 10+ hours a week for 6 months, and even more frustration. The scope and the approach drastically changed.

SUBJECT-MATTER LEADER, NATIONAL CONSUMER MORTGAGE LENDER

BENEFITS OF AN INDEPENDENT ENGINEERING READ

01 SCOPE BEFORE BUILD, (OR AFTER, IF YOU'VE ALREADY DONE IT)

The most valuable month of AI development is the one that never happens, especially when the waste comes 12, 18 or more at a time. Narrowing a mandate to a defensible use case, before major capital is committed, is where the path to the return gets discovered, and organizational intelligence grows into broader and more complex cases. This may mean that flashy AI becomes boring AI, or even non-AI. Engineering can defend this.

02 REGULATORY CONSTRAINTS AND HUMAN CONSTRAINTS AS DESIGN INPUTS

In regulated industries the technology compliance surface gets treated *by technologists* (yes, my people), as an obstacle that slows down innovation. Adoption is treated the same way. Teams pick what the AI should be, then add human adoption and compliance after the fact. These become huge time and resource sinks. Adult engineers know complex constraints are not always easy to meet, but *they must be considered from the beginning*. Missing this is a skill issue, not an operational or technical issue. Doing compliance and adoption well can block ideas that will never meet the standard, which included about half of the SME leader's own ideas. It can also simplify ideas that are too complex until they become genuinely *innovative*, like their RAG system.

03 NO DEPENDENCIES

AI systems require support from the day they are online, to the day they are sunset. If you can't tell whether your system is degrading, it is. Not always in intuitive ways. Building to the client's own maintenance capacity is a scoping decision made at the start, not a handover problem solved at the end. This forces trade-offs as well, which is important. And as their operational capabilities grew, so too did the kinds of systems they could build. Successful use cases demonstrate their value before scale-up happens.

04 YOUR PEOPLE CAN'T SAY IT. YOUR VENDORS WON'T SAY IT. I WILL.

Misaligned incentives. The client had opportunities to partner with big names, or use big label technologies. Ben helped them choose what was right for them. Of course vendors want to help, and want to sell, and want you to be successful, but if big named vendors unlocked some secret path to AI value, they would say so. That doesn't mean they are not good at what they do. Ben will be the first person to point to a specific brand if that is the right one.



BEN HAMILTON

Principal Consultant
BDHamilton Consulting ·
Seattle

Ben Hamilton is the Principal Consultant of BDHamilton Consulting, an independent technology consulting practice in Seattle spanning strategy, technology, organization, and process. He is a chemical engineer with a PhD from the University of Minnesota who spent more than a decade in pharma, nanomaterials, and process engineering, including four years in Houston's energy sector during the fracking boom, where failure meant a fire, a spill, or a human life. He moved into data science and AI in 2017 and has consulted since 2018, including his years at PwC on the highest-risk AI workflows in financial services. He works with specialty and industrial chemical manufacturing, natural gas utilities and distribution, oil and gas, and pharmaceutical and biotech manufacturing, and others on referral.

HE HAS BEEN TALKING PEOPLE OUT OF AI SINCE 2018: ANTI-GARBAGE, NOT ANTI-TECH.

THE MOST VALUABLE MONTH OF AI DEVELOPMENT IS THE ONE THAT NEVER HAPPENS.

READY TO TALK?

Bring the decision or risk you are facing. No pitch, and an honest read on whether Ben is the right call.

BOOK 30 MINUTES

calendly.com/ben-bdhamilton-consulting/30min

