AI-Powered Market Intelligence:

Lessons from a Scan of Wealth-Building Initiatives

JANUARY 2026

Dalberg

AI-Powered Market Intelligence:

Lessons from a Scan of Wealth-Building Initiatives

Introduction

1

This brief at a glance

- In 2025, Dalberg Advisors tested a new approach to market intelligence using generative AI tools.
- The AI-enabled landscape identified and analyzed over 250 initiatives across six states in under five weeks reducing the timeline of a typical scan by an estimated 70%.
- The AI-enabled PEST analysis surfaced macro-level trends shaping the wealth-building ecosystem, enabling faster foresight, more strategic framing, and structured prioritization of external risks and opportunities.
- The methodology expanded visibility, introduced structured tagging, and reduced manual synthesis time – complementing (not replacing) human judgment and local expertise.
- The method also surfaced new limitations to monitor going forward: digital visibility bias, dependency on external platforms, and the need for rigorous human validation and contextual interpretation. AI technologies evolve rapidly, and users should view these workflows as starting points and update them to stay current and aligned with responsible AI practices.
- Funders and ecosystem builders can strengthen this work by actively sharing lessons from their own AI experiments, co-developing responsible use standards, and building open tools and frameworks that others can adapt and improve.

Funders and ecosystem leaders often want to better understand where their efforts can add the most value or avoid places where they may be redundant. Traditional approaches to market intelligence[1] that rely on known networks and static document reviews are labor- and time-intensive, risk becoming quickly outdated, and may bias results toward well-established or well-resourced actors.

Recent advances in generative AI (AI) offer a potential complement to traditional research methods. By synthesizing large volumes of publicly available data and applying consistent tagging, AI tools can help identify directional patterns, surface emerging initiatives, and support more timely and inclusive mapping efforts.

This piece shares insights from a 2025 initiative led by Dalberg Advisors. The effort applied an AI-enabled workflow to deliver market intelligence through two discrete analyses: (1) a scan of the landscape of wealth-building initiatives across New England and (2) a Political, Economic, Social, and Technological (PEST) analysis to surface macro-level trends shaping the broader operating environment. Both workstreams explored how this approach could reduce research timelines, expand visibility into community actors, and generate useful insights for strategy, investing, and collaboration, all while maintaining rigor and human oversight. The AI-enabled approach reduced the workload for a scan that previously took around five months to one that was completed in less than five weeks, representing time savings of roughly 70%. It identified and analyzed over 250 initiatives, across six states and 15+ sectors, offering a timely view of the regional landscape. In parallel, the accompanying PEST analysis was also accelerated through an AI-assisted process that took one week, enabling the team to rapidly identify and prioritize macro trends with strategic relevance faster than through a traditional approach.

This brief shares practical insights and strategic implications from applying generative AI to market intelligence research through a regional landscape scan and PEST, offering lessons for funders and ecosystem actors seeking faster, broader, and more structured field awareness.

1 In this context, market intelligence refers to both the landscape scan of wealth-building initiatives and the PEST analysis of broader trends shaping the ecosystem.

Methodology snapshot: A regional landscape scan and PEST analysis using AI-enabled methods

2

How this methodology was designed

The AI workflows used in this guide were developed through rapid iteration. Early prompts returned hundreds of results, many irrelevant. In the landscape scan, some outputs included inactive organizations or short-term relief programs outside the intended scope of the scan.

Iteration helped refine prompt structure, apply a multi-step search approach, and embed inclusion criteria across both the landscape scan and the PEST. The methodology leveraged AI to extract data, evaluate relevance, apply tagging logic, and group trends, cutting manual review time in half.

Each round of testing led to improvements in clarity, accuracy, and efficiency, ultimately yielding a replicable method.

This methodology combined generative AI with rigorous human validation to produce a timely view of wealth-building activity in New England. AI tools enabled the rapid organization of large volumes of data, while human oversight ensured nuance and contextual relevance in the findings. The objective was to explore how AI could help deliver more timely, consistent, and scalable insights while maintaining rigorous human validation throughout.
Regional landscape scan: The scan systematically extracted publicly available information, applied a standardized taxonomy, and generated structured tags across key variables such as strategy type, target population, and geography to compile a broad-based view of the landscape scan.

Figure 1: 10-step methodological workflow for landscape scan

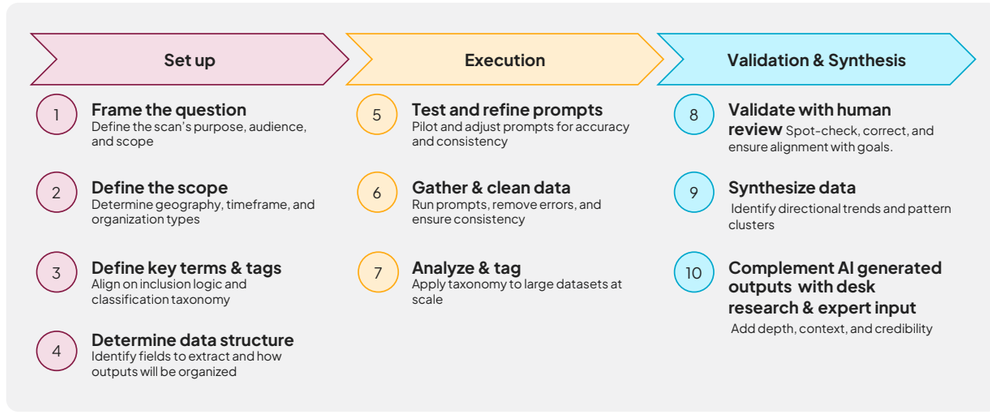

PEST analysis:

In parallel, a PEST analysis situated the landscape findings within broader regional trends. Traditionally a time-intensive exercise, the PEST was accelerated using AI for rapid signal gathering and synthesis. This enabled the analysis to consider how external forces, such as demographic shifts, market volatility, and technological change, may shape future gaps, risks, and opportunities in the wealth-building ecosystem.

Figure 2: 5-step methodological workflow for PEST analysis

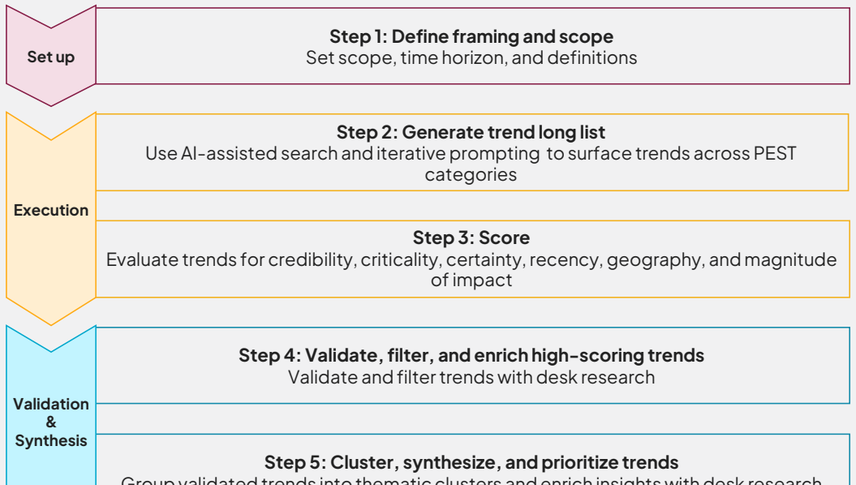

AI in market intelligence: Tradeoffs in speed, visibility, structure, and accuracy

3

Applying AI to the market intelligence analyses enabled several practical shifts in speed, coverage, structure, and accuracy.

An AI-first method is best used to accelerate and broaden market intelligence research, not to replace traditional research or local engagement. When paired with careful validation and field input, AI can reduce manual burden, increase visibility into emerging actors, and support faster, more strategic decision-making.

In the methodology outlined above, this translated into four specific shifts in how market intelligence research was conducted.

Table 1: Comparison of research approaches

	DIMENSION	TRADITIONAL RESEARCH APPROACH	AI-ENABLED RESEARCH APPROACH
1	**Time and efficiency**	Multi-month effort involving manual search, tagging, and synthesis.	Multi-week effort with structured prompting and partial automation.
2	**Coverage and visibility**	Draws from known networks, internal referrals, and well-documented actors.	Pulls from public sources, websites, and news coverage via structured prompts.
3	**Structure and repeatability**	Relies on human nuance and expertise making output high-quality but inconsistent if replicated across different teams.	Uses standardized prompts and AI logic for consistent tagging and extraction.
4	**Accuracy & Trust**	High accuracy based on vetted sources and expert validation, though susceptible to human error.	Moderate accuracy: depends on prompt quality, public data availability, and requires post-processing and validation to maintain credibility.

1. Time and efficiency: From manual synthesis to structured starting points

Traditionally, market intelligence is collected and interpreted manually with human researchers going through website by website to pull out key pieces of information across multiple documents. In this effort, by contrast, AI was prompted to extract relevant details from a range of sources and assign consistent tags (e.g., sector, strategy, population focus) based on detailed descriptions allowing faster turnaround. While this tagging required human validation to ensure reliability, the initial outputs offered a structured base from which analysis could begin.

AI did not eliminate the need for review, but it made it possible to structure hundreds of entries faster and with greater internal consistency than manual methods allow.

Case Study

Learnings from the 2025 scan of wealth building initiatives: AI pulled significant detail for each data point, requiring human cleaning to properly tag and sort the data. However, this extra information ultimately proved valuable, making it easier for reviewers to validate and cross-check complex cases that didn't fit neatly into a single category.

2. Coverage and visibility: From known networks to digitally visible actors

Traditional market intelligence often relies on information from known networks, such as grantees, institutional partners, and convening participants, which can unintentionally reinforce the visibility of well-connected actors. By contrast, the AI approach used web-based search and scraping of defined keywords to identify organizations beyond these "known actors," surfacing smaller initiatives. In the PEST analysis, it also revealed emerging signals, across a broad range of sources, sectors, and experts.

At the same time, the AI approach introduces some new limitations. Since AI's reach is limited to publicly available online content, initiatives and trends without strong digital footprints, such as newer grassroots efforts or informal networks, may be underrepresented. Additionally, since results rely heavily on keyword-driven web scraping, the choice and scope of these keywords significantly shape what initiatives are surfaced. Even so, by moving beyond institutional knowledge and tapping into a broader digital sphere, the AI approach expands overall visibility and helps democratize the coverage of the landscape. To ensure a more inclusive and representative analysis, this method is best complemented with direct field outreach, iterative keyword refinement, and stakeholder engagement.

Case Study

The AI's reliance on predefined keywords meant some types of actors were overlooked in early research drafts. Early searches prioritized organizations providing direct services, which unintentionally excluded key enablers like CDFIs and philanthropic funders whose roles, such as grantmaking and capital deployment. didn't fit the original keyword framing. To fill this gap, additional targeted searches were run with refined keywords, and human reviewers manually validated entries to ensure these contributions were included.

3. Structure and repeatability: From one-off scans to reusable components

Each step in the process, from prompt design to validation protocols, was carefully documented. This ensures that the approach can be replicated in future efforts. The result is a modular workflow that can be adapted to new regions or issue areas, such as small business ecosystems or early childhood programs. AI does not create an off-the-shelf solution, but it does enable teams to codify and reuse elements (such as prompts, tagging logic, and inclusion criteria) that would otherwise need to be rebuilt each time. It is important to note, however, that these components are not universally transferable. Without thoughtful adaptation, reused frameworks can reinforce outdated assumptions or miss context-specific nuances. Regular refinement based on feedback and field conditions is essential to ensuring ongoing relevance.

Case Study

In completing the PEST analysis, a systemized approach to identifying, selecting, and prioritizing trends was essential to enable a repeatable process. This required multiple iterations to develop, refine, and define assessment criteria that could evaluate trends objectively, consistently, and with an eye toward relevance and impact. Through trial and adjustment, the process yielded a well-calibrated set of criteria that supported more structured decision-making and laid the foundation for future, repeatable analyses.

4. Accuracy and trust: From surface signals to informed insights

Traditionally, market intelligence research is human-led, grounded in high-accuracy data drawn from vetted sources and expert validation. By contrast, while AI-generated scans are faster and broader, they introduce risks to data quality. Even with well-structured prompts, the model may produce false positives, fabricated details, or misclassified edge cases. Human oversight was required at every stage to ensure accuracy and contextual relevance.

While AI can efficiently classify what initiatives claim to do, it cannot exercise value judgements on how well they do it, how community-aligned they are, or how they operate within broader systems. In the context of the PEST, it can also surface macro-level forces, such as policy shifts or economic trends, but lacks the contextual judgment to interpret relevance, nuance, and scope of impact. As a result, practitioner insight, field interviews, and community engagement remain essential complements. Without them, there is a risk of overstating the reliability of AI outputs and misrepresenting the ecosystem the scan and analysis is meant to illuminate.

Case Study

In the 2025 wealth-building scan, AI tagged several programs as "asset-building" based on generic terms like "financial goals." On review, many turned out to be short-term relief efforts and not aligned with long-term wealth-building criteria–underscoring the need for human validation.

Best practices in building an AI-enabled market intelligence

4

Given the shifts in speed, coverage, structure, and accuracy, applying AI to market intelligence requires thoughtful design choices and guardrails.

Each improvement comes with trade-offs that must be actively managed. As the AI and data landscape evolves rapidly, best practices will continue to shift; what works today may need to be recalibrated tomorrow. Staying current with new tools, techniques, and risks is essential to maintaining analytic rigor and equity. The market intelligence analyses of wealth-building solutions revealed key lessons on effectively harnessing AI for ecosystem scans and trend analyses. Beyond basic principles, these best practices enhance the reliability, relevance, and replicability of AI-driven analyses:

- **Always validate before using outputs:**
  Human review is essential, especially for low-confidence entries or ambiguous classifications. No AI output should be taken at face value without vetting.

- **Benchmark against network-based scans:**
  In contexts where large-scale landscape analysis is new, comparing AI-enabled methods to traditional, relationship-driven approaches can provide a valuable reference point. This helps identify gaps, validate directional findings, and anchor AI insights in on-the-ground knowledge.

- **Break the scan into smaller parts to improve accuracy:**
  Dividing the scan by topic, geography, or organization type helps improve accuracy and relevance. Asking the AI to execute specific and targeted steps reduces confusion, lowers error rates, and produces cleaner outputs.

- **Treat findings as directional, not exhaustive:**
  The scan will not capture every organization or initiative, but it reliably surfaces directional insights across sectors, useful for ecosystem-level pattern detection, not exhaustive counts.

- **Document assumptions and decisions:**
  Clear inclusion criteria, tagging rules, and prompt logic help ensure replicability and transparency.

- **Ensure data privacy and model containment:**
  When working with sensitive, proprietary, or unpublished data, it is critical to use enterprise-grade AI platforms that guarantee data security and prevent model training on uploaded content.

- **Combine AI with traditional inputs:**
  Desk research, expert interviews, and community conversations remain essential for interpreting what the AI can't see.

For an additional list of good practices please visit Prompt engineering best practices for ChatGPT.

Strategic implications for funders and ecosystem leaders

5

AI-enabled market intelligence offers practical value for funders and ecosystem leaders who seek to make more informed decisions in dynamic, fragmented fields.

- **When to use this method:**
  This approach is particularly helpful at moments of orientation, strategic refresh, and exploration. It is well suited to:
  - Rapid ecosystem mapping to support field convenings, partnership development, or collaborative investment design
  - Periodic refreshes of prior field analyses, particularly when resources are limited for fully manual scans
  - Surfacing directional trends or external forces that may shape or disrupt current strategies, especially in long-term or systems-oriented investments

- **How this fits alongside traditional approaches:**
  Rather than replacing field interviews, participatory design, or relationship-building, AI-enabled market intelligence can make those engagements more targeted and strategic. For example, AI-generated insights can help prioritize deeper research or guide which actors to engage in co-creation or consultation efforts.

- **Where collaboration could strengthen the field:**
  Funders, intermediaries, and public institutions each have a role to play in building more efficient and equitable ecosystem intelligence. Key opportunities include:
  - **Building institutional capacity and public goods** by sharing lessons learned and tips on various methodological approaches (e.g., supporting shared prompt libraries and modular workflows to reduce duplication and improve quality across landscape analyses)
  - **Embedding rapid scans into multi-stakeholder learning cycles**, such as funder collaboratives or backbone organizations seeking to align strategies
  - **Creating accountability for ethical and responsible AI use** through shared playbooks, templates, and case studies, and by advancing sector-wide standards or a dedicated body to guide responsible, inclusive, and context-aware use in impact-focused work
  - **Supporting education for equity in AI data design and use** by training teams for bias mitigation and inclusion criteria ensuring that AI-driven insights do not inadvertently reinforce power imbalances

In all cases, responsible use of AI should prioritize transparency, reproducibility, and partnership. When used in service of clarity, this approach can help advance more equitable and coordinated strategies across the field.

Looking ahead: Building toward better field awareness

6

As funders and ecosystem leaders seek more adaptive strategies in a rapidly evolving landscape, the value of timely, structured, and scalable insight generation will only grow. To meet this need, AI-enabled market intelligence provides a valuable approach. Rather than substituting for human judgment or trusted relationships, it works alongside them to strengthen insight.

Moving forward, this approach has the potential to support a wider range of applications, including:

- Monitoring ecosystem shifts in real time to help decision-makers anticipate emerging trends and align strategies with future conditions rather than lagging indicators
- Optimizing the deployment of limited resources by identifying where risk, momentum, or policy windows create outsized potential for catalytic investment
- Informing place-based funding strategies with more regular updates

◆ Realizing this potential will require more than just technical tools.

It will require shared standards, responsible experimentation, and cross-sector learning. Funders and public institutions alike can play a role by investing in open-source taxonomies, piloting use cases across different domains, and embedding AI-enabled scans into their broader learning and strategy cycles. Done well, this method can help the field become more responsive, more inclusive, and better informed – unlocking faster visibility into what is emerging and where resources might be most catalytic.

Annex 1: A Detailed Replicable AI-enabled Method for Landscape Analysis

7

This section outlines a step-by-step approach for conducting a landscape analysis using generative AI. It is designed to support teams that want to apply AI in a structured, responsible way to identify actors, activities, and trends within a defined issue area or geography.

The method is modular and flexible. It can be adapted for different purposes, such as refreshing existing landscape data, identifying emerging interventions, or informing funding and partnership strategies. This 10-step framework guides the entire workflow, beginning with scoping and ending in synthesis. At each stage, it highlights the value of integrating AI with human insight and sector knowledge.

Flag: Because AI technologies and platforms continue to evolve rapidly, users of this guide should treat workflows and prompts as version-specific starting points. Future teams are encouraged to review and revise the methods outlined here to reflect technological advancements, maintain methodological integrity, and ensure continued alignment with best practices in responsible AI use.

This methodology was used to produce the Landscape Analysis of Wealth Solutions in New England (inclusive of initiatives founded between 2020-2025). This methodology follows a ten-step process designed to balance the efficiency of AI tools with the rigor of human validation:

Figure 3: 10-step methodological workflow for landscape scan

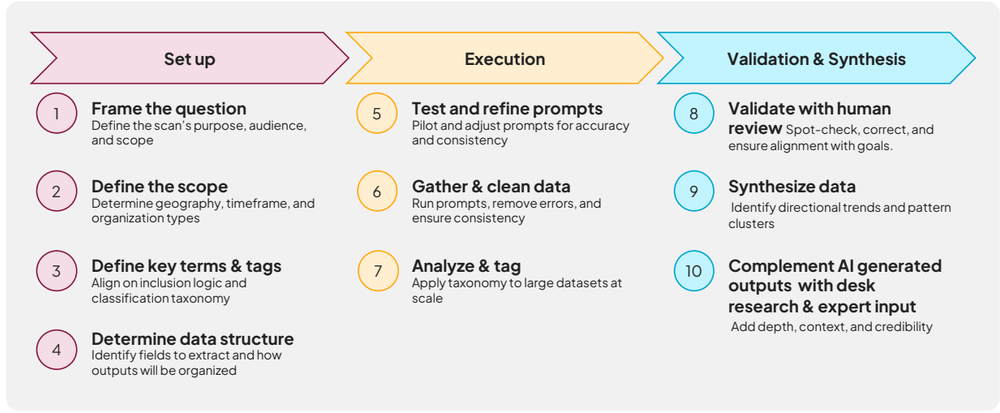

Step 1: Frame the question

Before using any AI tool, it is essential to begin with clear objectives, a defined audience, and a focused scope. A well-formulated research question is not only a starting point, but it anchors the entire landscape process. Without clear framing, AI tools are more likely to generate inconsistent, irrelevant, or overly broad outputs. Defining the purpose and scope early helps ensure that the prompts are tailored and contextualized, the inclusion criteria are specific, and the findings are actionable and relevant.

Before developing AI prompts or gathering data, it is important to clarify:

- **Purpose and objectives:** What the landscape analysis is intended to support (e.g., strategy setting, partnership development, funding alignment)
- **Audience:** Who are the primary users of the findings will be (e.g., funders, researchers, community stakeholders)
- **Temporal focus:** What focus is most appropriate (e.g., point-in-time snapshot, trend overview, sector-specific profile)
- **Primary research question:** Synthesize the audience and purpose into one central guiding question that can be translated into clear, targeted AI prompts

Example from the 2025 wealth-building landscape analysis:

The team began by clarifying the purpose, audience, format, and research question for this scan to ensure the AI had the full context to generate strategically useful outputs:

- **Purpose:** Equip funders and ecosystem leaders with a timely, data-grounded view of the regional landscape of financial resilience and wealth-building initiatives
- **Audience:** Funders, policymakers, practitioners, and ecosystem leaders interested in impactful or scalable wealth-building solutions
- **Temporal focus:** A point-in-time "pulse check" snapshot of initiatives currently active across the region, with basic segmentation by geography, strategy, and target population
- **Primary research question:** What does the current landscape of wealth-building and financial resilience solutions in New England look like–and what external trends are most likely to influence where funders can fill the most pressing gaps today?

This framing allowed us to distinguish short-term stabilization efforts from longer-term wealth-building strategies such as asset accumulation, credit-building, and intergenerational transfers.

◆ Step 2: Define the scope

AI performs best when guided by clearly bounded criteria. Scoping decisions shape the relevance, accuracy, and replicability of results. They also guide how AI prompts are constructed and how the resulting dataset is interpreted:

- **Geography:** What is the spatial focus of the analysis? This could include specific cities, states, regions, or multi-state groupings
- **Timeframe:** What period should the data reflect? This may include only current initiatives or a defined launch window (e.g., 2020–2025)
- **Organization types:** Which types of actors should be included? Options may include nonprofits, community-based organizations, CDFIs, government programs, philanthropic funders, private enterprises, etc.
- **Thematic inclusion criteria:** What characteristics make an organization or initiative relevant to the topic of interest? This could include specific intervention types, target populations, or operational scale

Example from the 2025 wealth-building landscape analysis:

In the pilot scan, the team established clear parameters to focus the AI on relevant entries and ensure consistency in how results were generated and interpreted:

- **Geography:** The analysis focused on six New England states–Massachusetts, Connecticut, Rhode Island, Vermont, New Hampshire, and Maine
- **Timeframe:** The analysis included initiatives launched between 2020 and 2025 that were still active at the time of data collection
- **Organization types:** The team intentionally did not limit the scan by actor type. Given the complexity of the financial resilience landscape, nonprofits, CDFIs, government programs, philanthropic funders, fintech startups, and private enterprises were included to capture the full range of contributors
- **Thematic inclusion criteria:** Only initiatives that supported long-term wealth-building or financial resilience were included. This was defined as interventions contributing to decreasing liabilities, increasing assets, sustaining wealth, and empowering wealth-building

◆ Step 3: Define key terms & tags

Clear, consistent definitions are critical when using AI to support landscape analysis. Generative AI models operate based on language patterns, which means ambiguous or undefined terms can lead to inconsistent or inaccurate outputs. Defining key terms up front helps ensure the AI interprets inclusion criteria correctly and that analysts can apply consistent logic when reviewing and tagging results.

This step includes two core activities:

- **Create a mini glossary of sector terms and definitions:** Clarify any word or phrase that signals inclusion or exclusion. This may include broad concepts (e.g., "wealth-building," "financial resilience") as well as narrower terms (e.g., "initiative," "currently operating")
- **Specify the keywords used for guiding the search:** Define a focused set of terms and phrases that signal inclusion based on the scan's thematic scope (e.g., legacy planning, debt counseling, budgeting workshop, etc.)

Example from the 2025 wealth-building landscape analysis:

To ensure clarity across hundreds of entries, a core set of definitions and tags was developed and used throughout the analysis:

Non-exhaustive excerpt of defined terms:

- Wealth-building solution- any program, service, or policy intervention that helps individuals or households improve their long-term financial position. Categorizing wealth-building solutions into four distinct but interrelated strategies: decreasing liabilities, increasing assets, sustaining wealth, and empowering wealth-building. These solutions must have a clear, tangible effect on a person's ability to move from financial precarity toward financial security and autonomy. For example, a wealth-building solution might take the form of a down payment assistance program that helps first-time homebuyers build equity, a community loan fund that supports local entrepreneurship, or a financial coaching initiative that enables participants to repair their credit and access affordable financial products

- New England- Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, and Vermont. An initiative qualified as "New England–serving" if it was either headquartered in one of the six New England states or explicitly served residents of one or more New England communities
- Low- and moderate-income (LMI) households- LMI households are those earning no more than 80 percent of the local area median income (AMI). Within this group, "low-income" typically refers to households at or below 50 percent of AMI, while "lower-middle-income" (or moderate-income) denotes those earning between 51 and 80 percent of AMI

Step 4: Determine data structure

Before running AI prompts at scale, it is important to define the data structure that will guide how results are organized, tagged, and analyzed. This step ensures that the output collected through AI is usable and aligned with the analytical goals of the landscape analysis.

Similar to column headers in a spreadsheet, a strong data structure includes a clearly defined set of fields that match the types of insights the analysis is intended to generate. These fields also shape how prompts are written, what information is extracted, and how human validation is conducted.

When designing the structure, it is useful to work backward from the intended synthesis questions. For example, if the goal is to compare the types of organizations in particular field across states, include fields for "organization type" and "state served."

Example from the 2025 wealth-building landscape analysis:

To support analysis of patterns across organizations, strategies, and populations, various data fields were used to structure AI prompts and organize outputs

Non-exhaustive excerpt of data collected:

- **Organization information**
  - Organization name
  - Type (nonprofit, CDFI, public agency, etc.)
  - Organization launch date
- **Initiative details**
  - Initiative name
  - Description of services or programs
  - Year initiative launched
  - States served
  - Wealth-building strategy (e.g., increase assets, decrease liabilities)
- **Population Served**
  - Population income level
  - Demographic focus, if stated (e.g., LMI households, immigrant communities, otherwise return "Unclear")

Additional notes

- Rationale for inclusion
- Missing data flags
- Source quality or certainty rating

Defining these fields in advance allowed the AI outputs to be automatically sorted into a structured dataset, significantly reducing the time required for downstream cleaning and synthesis. It also enabled the team to surface patterns by geography, population, and strategy type with minimal additional formatting.

Step 5: Test and refine your prompts

Before deploying AI across a full dataset, test the prompt to ensure outputs are accurate, consistent, and aligned with your inclusion criteria and classification schema. Prompt quality is a major determinant of AI output quality where specificity, structure, and tone all matter.

Key techniques for prompt development include:

- **Use clear, direct language:** Write, affirmative instructions avoid repetition or ambiguity; and embed illustrative examples whenever possible
- **Embed definitions and classification logic:** Reinforce key terms, eligibility rules, and data field requirements directly in the prompt
- **Segment the search space:** Break the analysis into smaller segments by geography, organization type, or time period to improve accuracy and focus
- **Test with multiple models (optional):** For complex classification tasks, consider running prompts through multiple models (e.g., GPT-4 and Claude) to compare consistency and minimize model-specific biases

In all cases, it is essential to:

- Conduct error analysis by reviewing initial outputs for common issues such as hallucinated entries, false inclusions, or missing tags. These findings should be used to revise the prompt language, clarify inclusion logic, and improve the overall quality and reliability of the results
- Pilot test the prompts on a diverse set of 10 to 20 example organizations that vary by size, sector, geography, and organizational type. This step helps identify where the AI performs accurately, where it misses key edge cases, and where human review will be most necessary

Example from the 2025 wealth-building landscape analysis:

During testing, prompts were iterated multiple times to improve alignment with the inclusion criteria. Early versions produced a high volume of results, but included many programs focused solely on short-term financial relief rather than long-term wealth-building. To improve precision, the team made the following adjustments:

- **Embedded positive indicators:** Keywords like "matched savings," "down payment assistance," "business capital," and "financial coaching" were added to signal relevance
- **Added exclusion logic:** Terms associated with short-term emergency relief were used to deprioritize programs that did not meet the wealth-building definition

- **Segmented the analysis:** Prompts were run separately for each New England state and filtered by year of launch (e.g., "Connecticut-based wealth-building initiatives launched in 2021")
- **Requested citations:** Each AI result was required to include a public source link (e.g., organization website or news article) to support downstream validation

Step 6: Gather and clean your data

Once prompts have been refined and tested, the next step is to run the full set of AI queries and compile the resulting data. At this stage, the focus shifts from prompt development to data collection and quality control. Even when prompts are well-designed, the output may contain inconsistencies, incomplete fields, or formatting issues. Cleaning is therefore an essential step before tagging, analysis, or synthesis begins.

The goal of this step is to transform raw AI-generated outputs into a structured dataset that is accurate, complete, and aligned with the fields defined earlier in the process.

Recommended practices include:

- **Standardize formats:** Ensure consistency in how information is presented across fields such as dates, geographic names, organization types, or initiative categories
- **Identify and flag gaps:** Note entries with missing information or unclear classification for later human review
- **Resolve duplications:** Combine or remove duplicate entries, particularly if prompts were run in overlapping segments
- **Check internal consistency:** Review whether entries that appear similar are tagged similarly, and whether any fields conflict with others (e.g., an initiative marked as "inactive" but described in present tense)

Step 7: Analyze and tag

After the dataset has been cleaned and standardized, the next step is to classify each entry based on the predefined taxonomy. This step transforms a collection of raw descriptions into a structured, analyzable landscape. Generative AI can assist with tagging at scale, allowing teams to apply consistent classification logic across hundreds of records more efficiently than manual methods alone.

AI can be prompted to assign values for specific fields such as sector, initiative type, population focus, or strategy used. These tags enable pattern recognition, allow for filtering and cross-comparison, and help surface insights in the landscape.

To ensure accuracy, it is important to use well-structured prompts that:

- Reference the exact taxonomy terms defined earlier in the process
- Provide examples or instructions for how to distinguish between similar categories
- Ask AI for a confidence score to ensure detailed human verification
- Include logic for confirming eligibility (e.g., whether the initiative meets inclusion criteria)

When used responsibly, this step can reduce the burden of manual tagging while preserving analytic consistency.

Example from the 2025 wealth-building landscape analysis:

In this application, AI was used not only to classify organizations by taxonomy but also to assess whether each initiative met the criteria for inclusion in a wealth-building landscape. A five-point scoring rubric was developed, grounded in core principles of wealth building. AI was prompted to review initiative descriptions and assign a score from 1 to 5 based on how well each entry aligned with these criteria.

In addition, AI was used to extract and summarize the financial solution each initiative offered (e.g. mortgage assistance, small business budget support, etc.) and to group similar solutions into broader categories. This enabled more structured comparisons across initiative types and helped surface patterns in how different strategies were being applied across sectors and geographies.

Step 8: Validate and source check with human review

Even when AI is used to support classification and synthesis, human validation remains essential. Generative models can misinterpret context, overlook nuance, or inconsistently apply classification logic, particularly for edge cases or initiatives with limited online presence. A structured review process ensures that outputs are accurate, relevant, and aligned with the goals of the landscape analysis.

Human review should focus on three main areas:

- **Accuracy:** Confirm that each entry has been correctly classified and that the AI did not hallucinate, mislabel, or fabricate information
- **Consistency:** Ensure similar initiatives are tagged the same way and that definitions are applied uniformly across the dataset
- **Completeness:** Identify entries with missing fields, low-confidence outputs, or ambiguous eligibility that require follow-up research

A tiered review approach can improve efficiency, allowing analysts to focus attention on entries most likely to contain errors.

Recommended validation practices include:

- **Spot-check a representative sample:** Review a subset of entries from each segment (e.g., by state, sector, or population) to verify the quality of AI-generated classifications
- **Review ambiguous or low-certainty cases:** Prioritize entries flagged by the AI or the analyst team as borderline, unclear, or incomplete
- **Cross-check sources:** Compare AI outputs with official websites, program descriptions, and public documents (such as Form 990s) to confirm accuracy
- **Apply consistent logic:** Confirm that tagging logic has been applied uniformly, especially for common patterns or repeated solution types

Example from the 2025 wealth-building landscape analysis:

In the landscape scan of wealth-building initiatives, the team asked AI to assign a certainty score (e.g., high, medium, low) based on its confidence in meeting the inclusion criteria. A human analyst then manually validated all entries marked as medium or low certainty – cross-referencing program descriptions, websites, and public records. This approach helped us efficiently catch false positives and ensure edge cases were correctly classified without having to review every record individually.

We also implemented a structured quality check by manually reviewing every 10th entry in full. This helped identify recurring errors, test the consistency of AI-applied tags, and ensure that inclusion criteria were being interpreted and applied uniformly across the dataset.

Step 9: Synthesize data

With a structured and tagged dataset, AI can be used to support early synthesis by surfacing potential patterns for discussion, describing clusters, and generating draft narrative summaries. While AI cannot replace strategic analysis, it can act as an accelerant providing first-pass insights that a human team can refine.

Use AI to:

- **Describe segment trends** (e.g., "Write a summary of programs serving moderate-income populations in rural areas")
- **Group entries by shared traits** (e.g., geography, strategy type, maturity) and provide plain-language comparisons
- **Identify early signals or shifts** by comparing time-based cohorts (e.g., initiatives launched pre-2022 vs. post-2022)

AI summaries can support team-level synthesis and help produce audience-specific versions of key insights. However, they are best used as a draft. Final outputs should be shaped through human validation and deeper analysis.

Step 10: Complement AI generated outputs with desk research and expert input

AI-generated synthesis can provide a useful starting point, but it should always be grounded in traditional research methods and practitioner insight. Desk research and expert input help validate findings, fill in context that AI cannot detect, and ensure that conclusions reflect real-world dynamics, not just digitally visible data.

This step is especially important when landscape findings are intended to inform strategy, funding, or public communication. Conversations with practitioners, community leaders, and subject-matter experts can reveal important nuances that AI cannot reliably capture including implementation barriers, emerging innovations, and unintended effects.

Recommended practices include:

- **Supplement with desk research:** Review relevant policy briefs, academic publications, and practitioner reports to situate findings within broader trends and known sector challenges
- **Engage field experts:** Share early synthesis with practitioners and analysts to test interpretations, pressure-test assumptions, and fill gaps
- **Conduct interviews or focus groups:** Where possible, invite feedback from end users or community-based organizations to ensure that the findings reflect lived experience and are relevant for decision-making
- **Adjust narratives as needed:** Use expert feedback to refine language, reframe conclusions, or highlight overlooked issues that require further investigation

Annex 2: A Detailed Replicable AI-enabled Method for PEST Analysis

8

This section outlines a step-by-step approach for conducting a PEST analysis using generative AI. It is designed to support teams that want to apply AI in a structured, responsible way to surface and assess external trends shaping a given field or strategy area.

The method is modular and adaptable and can be applied in a range of contexts. This five-step framework moves from scoping and longlisting to trend validation and thematic synthesis. At each stage, it emphasizes the importance of pairing AI-enabled inputs with human judgment to ensure relevance, rigor, and strategic utility.

Figure 4: 5-step methodological workflow for PEST analysis

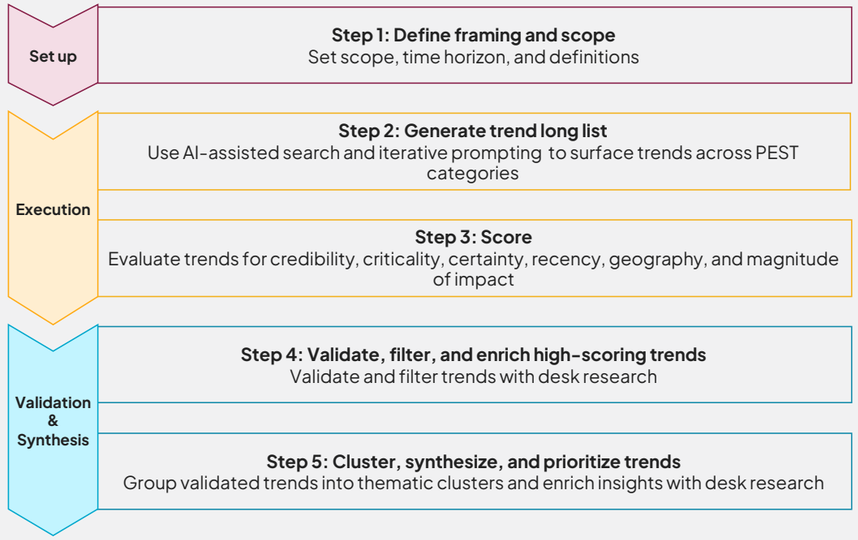

Step 1: Define framing and scope

Similar to the landscape scan, an AI-led PEST analysis process requires clear upfront framing to allow AI to focus on trends that are both strategically relevant and geographically grounded. This framing guided every subsequent step, from prompt development to trend synthesis.

Clarifying purpose and scope upfront ensures that trend identification remains focused, inclusion criteria are well-defined, and outputs are both relevant and actionable.

Before developing AI prompts or gathering data, it is important to clarify:

- **Geography:** What is the spatial focus of the analysis? This could include specific cities, states, regions, or multi-state groupings
- **Timeframe:** What period should the data reflect? This may include only current initiatives or a defined launch window (e.g., 2020–2025)
- **Organization types:** Which types of actors should be included? Options may include nonprofits, community-based organizations, CDFIs, government programs, philanthropic funders, private enterprises, etc.
- **Strategic relevance:** Clarify what types of trends will be considered relevant to the analysis (i.e. those with implications for wealth-building, household financial security, etc.)

Example from the 2025 wealth-building landscape analysis:

The purpose, audience, format, and research question were clarified at the outset to ensure the AI had the full context needed to generate strategically useful outputs:

- **Purpose:** To surface external forces that may impact household financial resilience, particularly those affecting asset accumulation, debt reduction, risk protection, and financial empowerment
- **Audience:** Funders, policymakers, practitioners, and ecosystem leaders interested in impactful or scalable wealth-building solutions
- **Geographic scope:** Trends must either affect three or more New England states or represent national shifts with demonstrated or likely regional relevance
- **Time horizon:** Focused on near-future trends shaping the 2025–2030 period
- **Strategic relevance:** Only trends with a plausible influence on wealth-building were included.

This framing ensured that the analysis stayed focused on trends with both to the overall mission and utility for funders seeking to anticipate and adapt – not just react.

Step 2: Generate trend long list

The second step involves constructing a structured prompt to generate an initial pool of trends across the four PEST categories: Political, Economic, Social, and Technological. The goal is not to reach precision at this stage, but to cast a wide net and surface directional changes that may warrant further investigation.

Generative AI can support this step by synthesizing web-based signals and institutional data into short, structured trend descriptions. Prompts should clearly define the purpose, scope, and criteria for what qualifies as a trend – ensuring outputs are usable and reduce the burden of downstream cleaning.

A strong longlist prompt typically includes the following components:

- **Context and purpose:** Describes the objective of the scan and the problem it aims to support (e.g., identifying shifts that may affect funding strategies, policy design, or program effectiveness)
- **Guiding research question:** Frames the type of trends to be identified and how they will be used
- **Scope constraints and definition of terms:** Establishes geographic relevance, time horizon, thematic focus and defines key relevant terms.
- **Trend definition:** Clarifies what constitutes a valid "trend" (i.e., observable developments or signals – not categories or topics)
- **Output format:** Instructs the AI to return 1–2 sentence descriptions, categorized by PEST domain, and written in plain language for strategic use

To improve quality and reduce repetition, prompts can be run in separate batches by domain (e.g., generate 20 Political trends only), or further segmented by sub-sector or geography.

This step is intentionally expansive. Overgeneration is encouraged, as many trends will later be deprioritized or excluded through scoring and validation.

Example from the 2025 wealth-building landscape analysis:

In this application, the following full-context prompt was used:

"You are supporting a strategic foresight and ecosystem analysis. Your task is to conduct a PEST analysis (Political, Economic, Social, Technological) focused on the financial resilience and wealth-building ecosystem in New England."

The prompt further specified:

- **Time horizon:** 2025–2030
- **Geographic scope:** Trends must affect at least three New England states or reflect national dynamics with regional relevance
- **Wealth-building relevance:** Trends must relate to decreasing liabilities, increasing assets, sustaining wealth, or empowering financial resilience
- **Output format:** Return at least 20 trends per category, each in 1–2 sentences, grounded in observable developments

Similar to the landscape methodology above, prompt quality plays a critical role in generating relevant, usable trend outputs.

Key techniques include:

- **Use clear, direct language:** Phrase instructions affirmatively and include illustrative examples to reduce ambiguity
- **Embed definitions and logic:** Include key terms, inclusion criteria, and output expectations directly in the prompt
- **Segment the search space:** Break prompts into smaller parts by geography, topic, or time frame to increase relevance and reduce noise
- **Compare across models:** For complex or sensitive domains, testing across multiple AI models can improve reliability and minimize bias

Step 3: Score

After generating a long list of trends, the next step is to evaluate which ones are credible, relevant, and useful for deeper analysis. This is done using a prompt that applies a structured scoring rubric that tests inclusion criteria against each trend. Scoring helps identify high-quality signals and filter out vague, dated, or off-scope entries before synthesis.

Illustrative criteria:

AI is asked to assess each trend using a binary score – 1 (meets criterion) or 0 (does not) – across six dimensions

1. **Credibility** – The trend is supported by at least two reputable sources
2. **Strategic relevance** – The trend relates to the thematic focus of the scan (e.g., wealth-building, economic resilience, inclusion)
3. **Certainty** – The trend shows signs of persistence or momentum (e.g., funding activity, policy movement, adoption)
4. **Geographic scope** – The trend affects multiple regions or aligns with national dynamics – not just local or isolated cases.
5. **Recency** – The trend is backed by 2025 sources (with 2024 sources used only to corroborate)
6. **Magnitude of impact** – The trend has the potential to influence a significant population, structural condition, or system – either through scale (e.g., affecting millions of people or multiple sectors) or depth (e.g., altering foundational drivers such as access to capital, labor market structure, or trust in institutions)

Example from the 2025 wealth-building landscape analysis:

In the scan, AI was used to score each trend across the six inclusion criteria listed above. Prompts were crafted to return:

- A breakdown of 1 or 0 scores for each criterion
- A total score (e.g., 6/6)
- A recommended maturity level (Emerging, Gaining Traction, or Showing Uptake)
- A short rationale or source reference for each scoring decision

Batches of 5-10 trends were scored in this way, and human reviewers then audited a sample of outputs to confirm consistency, fill citation gaps, and adjust misclassifications.

Step 4: Validate, filter, and enrich high-scoring trends

After initial AI-assisted scoring, a rigorous human validation process was conducted to ensure the final set of trends was credible, relevant, and strategically usable. AI scores served as a preliminary filter – not a final judgment. Each trend was manually reviewed to confirm alignment with the inclusion criteria, with special attention given to entries flagged for unclear or borderline rationale.

This validation process involved desk research to verify sources, cross-check claims, and – when necessary – replace weak or outdated references with stronger, more current evidence. Trends that lacked sufficient backing or failed to demonstrate clear strategic relevance were removed or reclassified.

Only trends that fully satisfied all six criteria were retained for final analysis. For each validated trend, AI was prompted to draft short executive summaries and justification bullets aligned to the inclusion criteria. These were then reviewed, edited, and supplemented through desk research to ensure accuracy, clarity, and readiness for use in strategic briefings.

Step 5: Cluster, synthesize, and prioritize trends

Flag: Use AI as a source of ideas in this step but maintain human ownership of the process to ensure nuanced and reliable outcomes.

The final step transforms a set of validated trends into strategic insight by grouping trends into thematic clusters and synthesizing implications. While previous steps focused on surfacing and verifying directional shifts, this step focuses on sense-making: identifying the patterns and priorities most relevant for decision-makers.

Rather than analyzing trends individually, this step focuses on surfacing cross-cutting dynamics – areas where multiple trends converge or where the same force manifests across domains (e.g., technological disruption reinforcing economic precarity).

For each cluster, trends were summarized using plain-language insights: what is changing, why it matters, who it affects, and what the implications could be for funders, policymakers, or community organizations. Additional desk research was used to:

- Fill gaps in context or implications
- Add specificity to trend examples (e.g., citing policy names or data points)
- Clarify causal linkages between trends and wealth-building conditions

Dalberg

About Dalberg Advisors

Dalberg Advisors is a strategic advisory firm that combines the best of private sector strategy skills and rigorous analytical capabilities with deep knowledge and networks across emerging and frontier markets. Dalberg Advisors works collaboratively across the public, private, and philanthropic sectors to fuel inclusive growth and help clients achieve their goals. For more information, visit dalberg.com

Dalberg