

Conceptual Search:
A Generative View of Entrepreneurial Imagination

Felipe A. Csaszar
Ross School of Business, University of Michigan
Ann Arbor, MI 48109
fcsaszar@umich.edu

September 21, 2026

Forthcoming in the Strategic Entrepreneurship Journal.

**Conceptual Search:
A Generative View of Entrepreneurial Imagination**

Abstract

Research Summary: Entrepreneurs cannot pursue ventures they never imagine. This article develops a generative view of entrepreneurial imagination as the construction of an evolving candidate set through generation, provisional evaluation, and revision. Conceptual search represents venture concepts as text, transforms them through interpretable operations, and compares their promise. A large language model prototype demonstrates complementary modes: local refinement lifts a concept from sixteenth to third among thirty ventures under a validated evaluator, while global recombination reaches distant possibilities that refinement makes competitive. The framework shows how search architecture determines what enters, survives, and develops, and how AI may broaden imagination within firms while producing correlated imagination across firms. By making imagination programmable, auditable, and experimentally tractable, conceptual search opens a research program on how entrepreneurial possibilities are constructed.

Managerial Summary: Entrepreneurs can pursue only the ventures they imagine, and valuable ideas routinely go unbuilt because nobody thinks of them. AI makes imagining cheap: it can draft, critique, and rewrite venture concepts in seconds, so developing ideas becomes a process that can be designed, run, and audited. I demonstrate a system that evolves real venture concepts: step-by-step refinement lifted a middling concept to third place among the thirty ventures ranked under an AI judge shown to predict fundraising, while recombining distant ideas created novel concepts that refinement made competitive. The entrepreneur's job shifts from generating a few ideas to governing many: choosing what to search, keep, and test in the market. One caution: rivals using the same AI will imagine similar ventures; advantage comes from proprietary problems, data, and evidence.

Keywords: Entrepreneurial imagination; entrepreneurial cognition; conceptual search; artificial intelligence; genetic algorithm

“If you want to have good ideas you must have many ideas.”
—Linus Pauling

1 Introduction

1.1 The problem of entrepreneurial imagination

It is the work of entrepreneurs to bridge the distance between what is technologically possible and what is economically realized. This is among the most consequential activities in a market economy, and also among the most difficult—not only because of capital constraints or competitive pressure, but because the space of possibilities is effectively unbounded. An entrepreneur can usually do many things; the problem is to imagine which one is worth doing. The difficulty sharpens when new technologies expand the range of feasible applications before anyone knows which business models will work. Entrepreneurs confront it through *entrepreneurial imagination*, the process by which they construct, provisionally evaluate, and revise venture concepts before or between consequential actions.

Figure 1 shows why this process matters. At any moment, the economy realizes only a small subset of the value-creating ventures that are technologically feasible. The *possible frontier* marks the boundary of what could in principle be attempted given current technology, regulation, and infrastructure. The *realized frontier* marks the boundary of what has actually been imagined, designed, and brought to market. This paper studies the upstream process that helps determine the realized frontier: which possibilities enter consideration and how they are evaluated.

The Million Dollar Homepage illustrates this imagination gap in its simplest form. In 2005, Alex Tew put a blank image with one million pixels online, offered each pixel for one dollar as advertising space, and quickly sold them all (Csaszar 2018:616). All the necessary technical elements had been available at least five years earlier, and the business model was among the simplest imaginable. No technical breakthrough was required; what had been




Figure 1: The possible and realized venture frontiers; the dashed arc marks an earlier possible frontier.

missing was someone imagining it.

The example is unusual in its success, not in the underlying problem. Entrepreneurs cannot consider the entire possible venture space (Simon 1973, Levinthal 2011). Instead, they work with a small set of representations such as product concepts, business models, and venture plans. I call each such representation a *venture concept*; together, the venture concepts under consideration form the *represented candidate set*. This set changes as entrepreneurs generate, evaluate, revise, and discard concepts.

Candidate-set construction matters for a simple reason: entrepreneurs cannot choose a possibility they never consider. Ventures that would create value if attempted may therefore never be attempted because no one conceives of them—or because the few who do cannot evaluate them well enough to act. Although this implication is self-evident, studying it in practice is difficult because rejected and unimagined concepts rarely leave a record. Gruber et al. (2008) provide empirical support for the idea: among venture-backed technology firms, founding teams that considered several target markets before entry reported higher early sales

than teams that considered only one. Candidate sets do not determine venture outcomes, but they bound what can be selected for testing.

AI changes the economics of candidate-set construction. Human attention, memory, and processing limits constrain both how many venture concepts can be considered and how fully each can be developed. But AI systems based on large language models (LLMs) can repeatedly generate, evaluate, and revise such concepts in the natural language through which entrepreneurs articulate them, lowering the marginal cost of search. Software-mediated search can also preserve the sequence of candidates, comparisons, and revisions, making a process that usually unfolds privately easier to reconstruct and study. Early evidence suggests that models can already assist with both generation and evaluation of business ideas (Csaszar et al. 2024b, 2026, Boussioux et al. 2024, Doshi et al. 2025, Terwiesch et al. 2026, Dell’Acqua et al. 2026). This evidence motivates the question this paper pursues: how does entrepreneurship change when entrepreneurial imagination becomes programmable—when the venture concepts from which action proceeds can be generated, provisionally evaluated, and revised through a designed, repeatable, and auditable process?

1.2 Approach and contributions

To pursue this question, I develop a *generative view of entrepreneurial imagination* and instantiate it in a computational prototype. The generative view holds that the ventures entrepreneurs pursue are shaped by the evolving set of venture concepts they construct through generation, provisional evaluation, and revision. I treat that construction as search over representations (Csaszar and Levinthal 2016), and introduce *conceptual search*, in which positions are venture concepts expressed in text, moves are rewrites and recombinations, and evaluation is comparative judgment of the concepts’ substantive promise. The prototype uses an LLM to conduct two complementary forms of search over real venture concepts: local refinement, which develops a single concept through successive rewrites, and global search, which explores more distant possibilities through population-based mutation and

recombination. In the prototype, automated refinement lifted a mid-ranked venture concept to third place among thirty real ventures under a previously validated evaluator, and recombination reached distant concepts that refinement then made more competitive. Together, the generative view and the conceptual-search apparatus turn a process that has been difficult to observe and vary systematically into an object of computational inquiry—giving the study of entrepreneurial imagination an instrument it has lacked (Kier et al. 2026). The prototype’s role throughout is demonstrative: it establishes that the theorized process can be built, observed, and varied; the paper’s closing research agenda organizes the resulting program.

This paper makes three contributions. First, the generative view of entrepreneurial imagination makes the evolving candidate set an explicit object of explanation. That entrepreneurs consider only a subset of possible ventures is a bounded-rationality premise as old as the behavioral tradition (March and Simon 1958, Ocasio 1997); the generative view treats that subset as a state of the venture-creation process—an evolving set that is constructed and revised and, once externalized, can be recorded and deliberately designed. Its construction is vulnerable to two kinds of error: *omission errors*, when a valuable possibility ends up absent from consideration, and *commission errors*, when a poor concept is advanced toward action. The view adds a common upstream mechanism to entrepreneurship theory: candidate-set construction explains which concepts reach consequential action while leaving intact established accounts of how venture possibilities enter attention, attract commitment, and become organized. The view also explains how *search architecture*—the representation, operators, evaluator, and rules governing search—shapes candidate-set evolution by determining which paths remain open and which alternatives become inputs to subsequent search.

Second, *conceptual search* provides an apparatus for studying search over substantively meaningful alternatives. Like NK models—which made organizational search systematic enough to be varied and observed—conceptual search allows researchers to hold some elements of a search architecture fixed while varying others and comparing the candidate sets that

result. Unlike NK models, whose positions are abstract bit strings, its positions are meaningful texts and its moves are interpretable transformations: changing a customer, problem framing, or value-creation logic is simultaneously a computational operation and a theoretical event. Venture concepts provide the application developed here; the apparatus can also apply to other types of alternatives that can be represented, transformed, and provisionally evaluated. The prototype implements local refinement and global recombination on real venture concepts, illustrating how conceptual search extends computational inquiry to settings where the meaning of positions and moves is itself part of the theory.

Third, the generative view identifies *correlated imagination*: AI can broaden the candidate set within each firm while pulling candidate sets across firms toward similar regions. This possibility arises when firms use similar models to search similar problems under similar criteria. Heterogeneity then depends increasingly on what firms bring to the search—distinct problems, proprietary context and data, and different search architectures informed by evidence from action. Thus, cheaper imagination can increase search without necessarily increasing the diversity of entrepreneurial possibilities pursued.

The argument unfolds as follows. Section 2 develops the generative view and relates it to entrepreneurship theory. Section 3 introduces conceptual search and implements local refinement and global recombination in a computational prototype. Section 4 develops the implications for entrepreneurship theory and practice.

2 The generative view of entrepreneurial imagination

The generative view starts from a simple premise: entrepreneurial choice depends on a candidate set that must first be constructed. Entrepreneurs generate, evaluate, and revise possible ventures before deciding which to pursue. This section develops that process, relates it to existing entrepreneurship theory, and explains how AI changes its economics.

2.1 Entrepreneurial imagination as offline search

Before entrepreneurs commit serious resources, they typically imagine some version of the venture. The intuition runs through the field’s foundations: Schumpeter’s entrepreneur envisages new combinations before carrying them out (Schumpeter 1934), Penrose’s entrepreneur acts on an “image” of the productive possibilities open to the firm (Penrose 1959), and Shackle’s entrepreneur chooses among rival imagined futures (Shackle 1979, Lachmann 1976). Entrepreneurial imagination spans a heterogeneous literature, recently mapped by Kier et al. (2026). Rather than proposing a general definition for that literature, the generative view uses a narrower analytical definition: *entrepreneurial imagination* is the loop of generating, provisionally evaluating, and revising venture concepts before or between consequential actions.¹ A promising concept may emerge through repeated thinking, research, discussion, comparison, and revision.

Consider a founder imagining a file-authentication service—a venture revisited in Section 3—that lets a document holder prove when a file existed and whether it has changed. Before writing code or seeking capital, the founder can compare concepts aimed at creators, archives, or regulated firms; alter the pricing model; or replace an ambitious product with a simple verification service. These changes are relatively inexpensive because they operate on a representation of the venture. Following Gavetti and Levinthal (2000), I use *offline* to denote operations over representations before consequential action, whether performed by minds or software, and *online* to denote operations in the environment that produce new evidence. Figure 2 represents this offline work as a generation–evaluation–revision loop: generation produces candidate concepts, evaluation compares them, and the resulting judgments can prompt further revision. Because comparisons among existing candidates shape what gets imagined next, generation and evaluation form a loop rather than successive stages (Finke et al. 1992).

The relevant distinction is therefore between offline work on representations and online

¹Appendix A provides a glossary of the paper’s construct definitions.




Figure 2: Entrepreneurial imagination and the transition to implementation.

action. Once an entrepreneur selects and commits to a concept, implementation may involve choosing a commercialization path, securing financing and stakeholder commitments, or building and launching the product (Gans et al. 2019). Such actions produce evidence no offline model can supply (Kerr et al. 2014, Packard et al. 2017, Foss and Klein 2012), which is why offline evaluation is necessarily provisional, as illustrated by the dashed feedback arrow in Figure 2. That evidence can in turn reopen imagination, as when an entrepreneur revises or replaces a concept after learning from implementation; lean startup methods deliberately accelerate this outer feedback loop through repeated cycles of action and learning (Ries 2011).

The remainder of this subsection examines three analytical constructs in turn: the venture concept is the representation being searched, the represented candidate set is the set of venture concepts available for choice, and offline search is the process that changes that set.

A *venture concept* is a representation of a possible venture, either held in the mind or externalized in forms such as a product concept, venture plan, or business model canvas.² A causally explicit venture concept can take the form of a *theory of value*, a conjecture about which activities and assets will create and capture value (Felin and Zenger 2017, Wuebker et al. 2023, Felin et al. 2024). Because entrepreneurs may also consider representations that are not causal, theories of value are a subset of venture concepts.

The *represented candidate set* is the evolving set of venture concepts under consideration.

²In its mental form, the construct corresponds to what Davidsson (2015) calls a *new venture idea*: an imagined future venture held in cognition. Externalization is what makes venture concepts objects that can be stored, compared, revised, and—in Section 3—searched by software (Csaszar et al. 2024a).

Its members aim to describe points in the *possible venture space*: the ventures that could in principle be attempted given current technology, regulation, and infrastructure, most of which are unknown to any actor (the feasible points of Figure 1). The relation is not one-to-one: several concepts may represent substantially similar ventures, while some concepts may be infeasible, internally inconsistent, or based on mistaken beliefs and therefore fail to correspond to any point in the possible venture space. At any moment, the candidate set records which possibilities have entered consideration in represented form. The analyses in this paper hold the possible venture space fixed to focus on how search changes the represented candidate set.

Offline search names the offline work introduced above: the manipulation of venture concepts to examine possible actions and consequences before learning through action itself (Craik 1943, Gavetti and Levinthal 2000, Nickerson and Zenger 2004). Researchers have long conceptualized entrepreneurs as searching offline through analogy, planning, and venture-concept development (Grégoire et al. 2010, Berglund et al. 2020, Osterwalder et al. 2010). Because thinking through a possibility is usually cheaper and more reversible than acting on the world, offline search can range from a single comparison to repeated cycles of generation, provisional evaluation, and revision (Newell and Simon 1972, Simon 1969, Gavetti and Levinthal 2000). Offline search changes the candidate set; I call this process candidate-set construction. Within each cycle—each pass through the loop—generation proposes candidates through processes such as retrieval, analogical reasoning, and recombination, and which process is used affects the quality of the ideas produced (Martins et al. 2015, Frederiks et al. 2019). Provisional evaluation retains some and rejects others, and revision uses those comparisons to alter a candidate or redirect the next attempt. The number and depth of these cycles are bounded by *cognitive scarcity*—the attention, memory, and processing limits that constrain both the number of alternatives considered and the detail with which each is developed. These limits vary across individuals and groups. For example, imaginativeness predicts the ability to produce good venture ideas beyond knowledge, experience, and motivation (Kier and McMullen 2018), and some founding teams examine several target

markets before entry whereas others examine only one (Gruber et al. 2008).

Candidate-set construction can fail through omission or commission (Sah and Stiglitz 1986). An *omission error* occurs when a valuable possibility is never generated or is wrongly discarded; what imagination fails to supply, choice cannot select (Shackle 1979). A *commission error* occurs when provisional evaluation advances a concept that is infeasible or would create less value than a rejected alternative. Failures after a concept has been selected—during resource mobilization, organization, or execution—are implementation failures rather than errors in candidate-set construction. Imagination is the binding constraint on venture outcomes when many plausible alternatives exist and search is narrow or costly enough that valuable possibilities go unrepresented. When alternatives are few or readily surfaced, the constraint shifts downstream: outcomes turn on who pursues a concept, what resources and incentives shape action, and how well it is executed (Kaplan et al. 2009, Gompers et al. 2020, Baumol 1990). These determinants belong chiefly to implementation—the third box in Figure 2—rather than to this paper’s focal mechanism.

The generative view therefore isolates candidate-set construction as a distinct, fallible part of venture creation.

2.2 The generative view and entrepreneurship theory

Entrepreneurship and adjacent strategy fields offer well-developed traditions explaining how venture possibilities enter attention, attract commitment, and become organized. The generative view does not replace these mechanisms. It identifies candidate-set construction as a cross-cutting process through which venture concepts enter, survive, and change. The paragraphs that follow place each tradition in Figure 2, moving left to right, then outward: the loop, the commitment that ends it, action and its feedback, and the surrounding resources. Each illuminates part of this process; none makes the evolving candidate set its primary object.³

³Because candidate-set construction concerns actors’ representations rather than the ontological status of opportunities, the generative view is compatible both with accounts in which opportunities preexist action

The discovery and opportunity-evaluation traditions concern the loop’s two boxes in Figure 2: “Generation” and “Provisional Evaluation.” Recognition, shaped by prior knowledge and structural alignment, brings candidate concepts into consideration (Shane 2000, Grégoire et al. 2010); evaluation, shaped by resource fit, decision rules, expertise, and uncertainty, influences which survive (Åstebro and Elhedhli 2006, Haynie et al. 2009, Wood and Williams 2014, Scott et al. 2020). Both explain a single pass: why an actor notices or favors a possibility. The generative view studies what repeated passes accomplish: how the set of rival concepts evolves as “Revision” carries each evaluation back into generation, reshaping the candidates available to the next round.

Within the same offline region, the theory-based and problem-solving traditions address what venture concepts contain and how the search over them is organized. A theory of value, as defined above, guides attention and experimentation (Felin and Zenger 2017); problem formulation and governance organize the search for solutions (Nickerson and Zenger 2004). The generative view studies the state this search produces: the evolving candidate set, which also carries concepts not necessarily causally articulated—analogs, fragments, sketches—that later cycles may sharpen into theories of value.

The loop ends at the “Commit to Action” arrow, where the judgment and ownership-competence traditions take over. Judgment governs non-routine resource commitments under irreducible uncertainty (Klein 2008, Foss and Klein 2012); ownership competence concerns what to own, how to govern it, and when to transfer it (Foss et al. 2021). The generative view is their pre-stage: the loop constructs and provisionally sorts the alternatives over which judgment is exercised. A richer candidate set can improve the alternatives available to judgment but can neither remove the uncertainty that makes judgment necessary nor give anyone an owner’s control over resources and stake in the outcome.

Commitment leads into the online region, whose “Implementation” box is home to the creation, effectuation, bricolage, and pragmatist traditions. Creation theory explains how

and with those in which they emerge through it, and with accounts in which the two intertwine, as when discovery and creation feed each other in a virtuous cycle (Zahra 2008).

opportunities emerge and become specified through action, experimentation, and stakeholder engagement (Alvarez and Barney 2007, Berglund et al. 2020); effectuation explains action that starts from available means rather than fixed goals (Sarasvathy 2001); bricolage, making do by recombining resources at hand (Baker and Nelson 2005); and pragmatism, the revision of beliefs through intervention and feedback (Zellweger and Zenger 2023). The generative view concerns only the offline region; thus, these traditions affect it solely through the dashed “Feedback from Action” arrow, by which evidence from action re-enters generation and changes the represented alternatives.

The resource-based and dynamic-capabilities traditions, finally, describe the environment surrounding the entire figure. Resources shape what entrepreneurs know and which concepts appear feasible, and so flow into the loop; dynamic capabilities govern how resources are deployed and reconfigured, and so operate chiefly in implementation (Barney 1991, Teece et al. 1997). These traditions explain what is feasible and what can be realized; the generative view explains how the alternatives connecting the feasible to the realized evolve.

The generative view therefore connects traditions associated with different parts of Figure 2 through a common process: the evolving candidate set. If that set is as old as entrepreneurship itself, why formulate this view now? Although entrepreneurs have always generated and revised alternatives, much of that process occurred within human cognition, leaving sparse traces and resisting controlled variation. AI-mediated search externalizes these operations, allowing their rules to be specified and systematically varied. Once external, the process becomes open to systematic study. In other words, AI offers the possibility of making entrepreneurial imagination a computational theory. Moreover, AI can affect its scale, speed, and direction. The next subsection examines how AI changes the economics of running this loop.

2.3 AI and the economics of entrepreneurial imagination

AI lowers the cost of cognition—including the operations entrepreneurial imagination runs on. Anyone who has asked a chatbot to draft a plan, critique it, and rewrite it has watched tasks that once consumed hours of human attention completed in seconds for cents. These tasks match the loop in Figure 2 one-to-one, with drafting as generation, critiquing as provisional evaluation, and rewriting as revision. A technology proficient at each operation could, in principle, run the loop itself. The possibility of running any of this in software—let alone the whole loop—is new, because the necessary language models became broadly available only with the launch of ChatGPT in 2022 (Ouyang et al. 2022) and have improved steadily since (HAI 2026).

In practice, early evidence indicates that generation and evaluation do work on venture concepts. The most direct evidence concerns AI operating alone, as it does in this paper. On generation, a preregistered blind comparison found that investors rated model-generated business plans above the entrepreneurs’ own (Csaszar et al. 2024b), and model-generated product ideas have likewise outperformed students’ (Terwiesch et al. 2026). On evaluation, studies have shown that evaluations of business plans by investors and by LLMs are highly correlated (Csaszar et al. 2024b, Doshi et al. 2025), and LLMs have outperformed individual human evaluators at predicting the funds ventures later raise (Csaszar et al. 2026).⁴ Evidence from human–AI collaboration points in a similar direction: human-guided AI solutions surpassed crowd submissions on average quality though not extreme novelty (Boussieux et al. 2024), and AI-enabled individuals have matched human teams on product-development tasks (Dell’Acqua et al. 2026).

An LLM running the loop of Figure 2 end to end would offer three benefits. It would permit *more and better search*: lower marginal costs allow more candidates, fuller development of each,

⁴Model evaluations of business plans correlated at 0.52 with investor scores (Csaszar et al. 2024b), and aggregated model judgments show a similar pattern (Doshi et al. 2025). In a fully prospective crowdfunding tournament, the strongest model’s venture ranking correlated more strongly with realized fundraising than any individual human evaluator’s (Csaszar et al. 2026)—a setting where experts distinguish quality only in some sectors (Scott et al. 2020).

and more cycles before action, relaxing the cognitive scarcity described earlier (Shepherd and Majchrzak 2022, Csaszar 2025); and because models draw on more fields than any individual, search could reach combinations that a single person would be unlikely to generate or evaluate. It would make search *programmable*: rules for generation, evaluation, and retention could be specified, varied, and rerun, turning candidate-set construction into a designed process that could aim at specific targets, such as a chosen degree of exploration or exploitation. And it would make search *auditable*: an external record of every candidate, comparison, and retention decision would allow researchers to reconstruct what the search considered and why concepts survived or disappeared.

Against these potential benefits stand three potential problems, no less conjectural than the benefits. The first concerns the quality of *generation*. Models may articulate familiar patterns fluently yet falter beyond them (McBride 2026)—and novel ventures live beyond them, which may demand causal theorizing about what does not yet exist (Felin and Holweg 2024).

The second problem concerns the accuracy of *evaluation*. Even if generation delivers, candidates multiply, making progress depend on the judgments that select among them (Spaulding 2016, Gennai et al. 2026). The difficulty is that venture concepts are not easily verifiable. In domains such as math and programming, it is possible to verify that a solution is correct (e.g., the theorem is proved or the program passes specified tests), which has recently enabled automated scientific discovery in some areas (Romera-Paredes et al. 2024, Novikov et al. 2025). In contrast, a venture concept is closer to a work of art: no procedure can execute a business plan and report whether it would succeed. Thus, evaluation is fallible; for example, it may reward polish or irrelevant attributes over substance (see, e.g., research on reward hacking and AI’s self-preference; Skalse et al. 2022, Panickssery et al. 2024). Such biases can compound into commission and omission errors, so the predictive validity of the loop’s candidates in Figure 2 must be demonstrated, not assumed (Chatterji et al. 2026:17). The third problem concerns *heterogeneity*. Even when generation and evaluation succeed,

shared foundation models could broaden each entrepreneur’s candidate set while making those sets more alike across firms—the possibility of “correlated imagination” developed in Section 4.3.

These potential benefits and problems are conjectures, and both concern the same missing object: a working search loop over venture concepts. Argument alone cannot settle claims about an artifact that does not yet exist, nor can evaluating components in isolation, because candidate-set construction is a loop whose elements interact. For conjectures of this kind, understanding advances by building the object (Simon 1969, Romme 2003, Zhang and Ertug 2025). AI research has long proceeded this way; systems that search over meaningful representations—from Logic Theorist through computational models of scientific discovery to the automated-discovery systems above—are the closest antecedents of the present work (Newell and Simon 1956, Langley et al. 1987). The next section therefore constructs a first loop over real venture concepts under a previously validated evaluator. The goal is not to adjudicate these conjectures, but to give the theory developed in this paper a working instance: a system in which venture concepts, the generation–evaluation–revision loop, and retention rules take observable and manipulable form. The working instance thus provides a concrete basis for Section 4 to revisit these conjectures and develop a research agenda on how programmable imagination changes entrepreneurship.

3 Evolving venture concepts through conceptual search

The generative view treats entrepreneurial imagination as search over venture concepts. This section makes that argument operational by building and running conceptual search: a system that generates, evaluates, and recombines real venture concepts. The search runs in natural language, which allows for flexibility and interpretability. Flexibility, because text can represent arbitrarily complex ventures; and interpretability, because all intermediate concepts and operations remain readable and inspectable by humans. Searching in text also keeps the

process close to how it is practiced: natural language is the medium in which entrepreneurs articulate ideas through analogy, metaphor, and narrative (Gartner 2007, Cornelissen and Clarke 2010). The exercise in this section is a deliberately small tutorial, not an effort to exhaust the method’s scale or potential. To keep the exposition uncluttered, method details stay out of the main text: Appendix C documents every procedure, parameter, and measurement; Appendix D reproduces the prompts; and the complete code and search records are available at <https://github.com/fcsaszar/conceptual-search>. I first describe the search inputs and evaluation; then local search, which refines a single concept; then global search, which recombines a population of concepts; and finally their integration, which does what neither does alone.

3.1 Search inputs and evaluation

The ventures come from Csaszar et al. (2026), a separate study of thirty real crowdfunding campaigns: U.S. technology-category Kickstarter projects with funding goals of \$10,000 or more that later raised between ten dollars and nearly half a million. Each is represented by an anonymized summary of about five hundred words covering the product, team, and risks. The campaigns launched after the training cutoff of every large language model that study evaluated, including the models used here. Models also ran without internet access, and anonymization removes names and other identifying details, so evaluation must rest on each concept’s substance rather than on any reputation attached to it. The thirty ventures were therefore new to the system in the same sense that a new pitch is new to an investor.

Search requires a common measure of performance. An LLM evaluator—DeepSeek V3.2 in the main analyses—judges which of two venture concepts is more promising. In this section, a concept’s *performance* is its win rate against the thirty originals—the *reference set*—or, for an original, against the other twenty-nine. Pairwise comparison is more reliable than absolute scoring when differences are subtle (David 1988, Zheng et al. 2023), and each pair is judged in both orders to offset any positional bias.

The win rate reads like a percentile: the median original wins about half its matches against the other twenty-nine, the strongest (**dreamie**) wins 96 percent, and the weakest wins none (Table B.1 in Appendix B lists all thirty). Every performance axis in the figures that follow is this win rate, and every claim of improvement or placement in this section is relative to this evaluator’s judgment. Performance here is evaluator-relative, not a realized venture outcome. Prior validation nonetheless suggests that the evaluator tracks realized outcomes: in the fully prospective tournament presented in Csaszar et al. (2026), the best human expert achieved a correlation of 0.48 between predicted and realized performance rankings, while the same version of DeepSeek used here achieved a correlation of 0.62.

To make the performance effects of the search methods clearly visible, search starts from the weaker half of that ranking. The bottom fifteen ventures (ranks sixteen through thirty in Table B.1) are the *seeds*; all thirty remain the measuring stick, so evolved concepts must beat ventures the search never saw. The seeds average a 28 percent win rate, against 72 percent for the withheld top half. The illustrations of local and global search below use this material in slightly different ways. Local search starts from the single strongest seed, a concept in the middle of the overall ranking: starting from mid-table makes the effect of refinement easy to measure. Global search starts from all fifteen seeds and asks how far this bad-to-mediocre material can be lifted toward the reference set’s best.

3.2 Local search

Local search uses hill-climbing: start from an incumbent concept, try small variations, and keep what improves performance the most (Kauffman 1993). Each round, the system proposes five targeted changes to the incumbent, rewrites the full concept around each change, and lets the evaluator rank the incumbent alongside its five variants; the winner becomes the next round’s incumbent, and the process repeats for twelve rounds. These moves are like those a founder could make when tightening a pitch after each investor meeting, but performed here in software, five variants at a time, with every comparison recorded.




Figure 3: Performance of the local-search incumbent (seeded with `clockchain`), round by round. Ticks at right mark the thirty reference-set ventures (dark: the withheld top half; light: the seeds).

Every generated text undergoes a *factual-fidelity check* against its source material. The key constraint is that a rewrite may not claim something that does not exist. Proposing a pilot, target, or partnership is fair; claiming that one already exists is not. Without this check, unsupported accomplishments could masquerade as improvements that the evaluator would erroneously reward.

The local-search run shown here is seeded with `clockchain`, the file-authentication venture of Section 2, which anchors a file’s digital fingerprint to a public time-stamped ledger—the strongest seed, sixteenth of thirty at a 47 percent win rate. Figure 3 plots the incumbent’s performance round by round, and Table 1 lists each accepted change. None changes the product radically. The accepted revisions chiefly alter the venture’s design and evidence plan—a narrower launch segment, a small user test with an explicit success threshold, a public demonstration of the core claim—and all remain proposals rather than accomplishments.

The curve in Figure 3 is not monotone because a round’s winner is chosen within its round while the curve measures it against the reference set, so a variant can win its round yet score slightly lower on the common scale. After twelve rounds, `clockchain` wins 88 percent of its matches and would place third among the thirty originals, behind only the two strongest.

Round	Accepted change	Performance
0	<i>Seed concept (clockchain), before any revision.</i>	47%
1	Propose a plan to target a specific, narrow initial user segment (e.g., digital illustrators on a single platform) for the beta launch to focus marketing and development.	62% ↑
2	Propose a small-scale, planned test with a handful of digital illustrators from the target platform to gather initial workflow feedback before the beta, framing it as a step to refine the tailored interface.	68% ↑
3	Propose a planned small-scale public demonstration during the crowdfunding campaign, using a sample file to show the technology’s verification capability and build trust.	80% ↑
4	Add a proposed, simple pricing structure (e.g., ‘We are considering a freemium model with X free verifications per month and a subscription for Y verifications’) to the Revenue Model and Sustainability section to address a key backer question.	82% ↑
5	In the ‘Revenue Model and Sustainability’ section, propose a specific, hypothetical pricing tier (e.g., ‘5 free verifications/month, then \$5/month for unlimited’) to make the freemium model more tangible for backers, while still noting it is subject to change based on the planned test.	80% ↓
6	Add a brief, hypothetical example of a digital illustrator using the tool to prove authorship of a piece of art posted on the named social platform, to make the value proposition more tangible.	82% ↑
7	Include a mock-up or wireframe screenshot of the proposed beta interface tailored for digital illustrators in the project description to visualize the solution.	78% ↓
8	In the ‘Roadmap & Delivery Plan’ section, specify the proposed timing for the small-scale test with digital illustrators (e.g., ‘planned for the month following successful funding’) and the public demonstration (e.g., ‘planned for the third week of the campaign’) to add temporal clarity and show proactive planning.	87% ↑
9	Define a clear, manageable test size for the ‘small-scale, planned test’ (e.g., ‘5 test users’) to show structured planning.	82% ↓
10	Define a clear, simple metric for the success of the proposed small-scale test with 5 digital illustrators (e.g., ‘percentage of test users who successfully completed the fingerprinting workflow without assistance’) to show how feedback will be measured.	87% ↑
11	Specify a candidate public ledger technology (e.g., a particular blockchain or a public timestamping service such as OpenTimestamps) to be used in the technical details to enhance credibility and transparency.	87% =
12	Detail the exact, simple metric for the small-scale test (e.g., ‘target: 80% of 5 test users complete the workflow unassisted’) in the Roadmap.	88% ↑

Table 1: The accepted revision behind each round of the local search in Figure 3, with the concept’s performance after each round.

Table 1 shows how the accepted revisions changed the represented venture; Appendix C.4 complements the table with exact excerpts from the seed and final concept.

Thus, twelve rounds of fidelity-checked local refinement lifted a mid-table venture concept from 47 to 88 percent—from sixteenth to third place among the thirty originals. Local search implemented exploitation (March 1991), but automated end to end.

3.3 Global search

A well-known disadvantage of local search is that it can stall at a *local peak*: a concept that none of the nearby rewrites on offer improves, even when better concepts exist elsewhere (Levinthal 1997). Global search, as implemented here, addresses this limitation with a *genetic algorithm*, a procedure that mimics natural selection: maintain a population of candidate concepts, let the stronger ones survive and reproduce with variation, and repeat until the population moves into better regions of the space (Holland 1975, Goldberg 1989, Koza 1992).⁵

Global search maintains a population of fifteen concepts—initially the fifteen seeds—each containing six written components: problem, customer, solution, delivery model, revenue logic, and distinctiveness. Search proceeds in *generations*. Each generation opens with a tournament in which every concept meets every other under the evaluator, giving each concept a *tournament score*. The search also tracks each concept’s *semantic distance*, a term borrowed from creativity research (see, e.g., Beaty and Johnson 2021). Here it denotes the distance in meaning between a concept’s text and the nearest of the fifteen initial seeds. Semantic distance therefore measures novelty relative to the seeds: how far imagination has traveled from where search began. The *selection rule* averages percentile scores derived

⁵The logic of this algorithm resonates with a long social-science lineage: variation and selective retention in creative thought (Campbell 1960), organizing (Weick 1979), economic change (Nelson and Winter 1982), and intraorganizational strategy making (Burgelman 1991). Unlike biological variation, variation in social systems can be directed by judgment and acquired improvements can be inherited—an “unabashedly Lamarckian” logic (Nelson and Winter 1982:11). The genetic algorithm also has a Michigan history: John Holland developed it there (Holland 1975), and his student John Koza extended it to genetic programming (Koza 1992); Holland’s advisor Arthur Burks had edited von Neumann’s (1966) posthumous *Theory of Self-Reproducing Automata*, an antecedent of the algorithm (see Dyson 1998); Burks, Robert Axelrod, Michael Cohen, and Holland formed the BACH complexity group; Axelrod’s (1984) tournaments anticipate this paper’s round-robin ranking; and Cohen is credited with pointing March to Holland (1975), a central inspiration of March (1991).

from tournament and distance ranks, so a concept ranks highly by winning matches and by occupying ground the seeds did not—a rule that protects newborn combinations long enough to have a chance to survive.

Three operators then build the next generation of fifteen—five concepts each—and the cycle repeats for twelve generations. *Elitism* carries five concepts ranked highly under the selection rule unchanged into the next generation, giving the population a memory; good material is not lost while search explores elsewhere. At most two elites may come from any one seed’s *family*—the concepts that have that seed among their recorded ancestors—so no family monopolizes survival. A mutation thus stays in its parent’s family, while a crossover child inherits the families of both parents and, if retained as an elite, counts against the cap of each. *Mutation* instructs the model to change one component of a well-ranked concept—a new revenue logic, say—while preserving the other five. *Crossover* builds a child from two parents descended from different seeds; its prompt assigns each component to one parent as its source. It operationalizes Schumpeter’s entrepreneur as a maker of new combinations (Schumpeter 1934, Bruderer and Singh 1996). An important constraint is that a child may inherit its parents’ designs but never their accomplishments; the prompt recasts any inherited claim of achievement as a proposal, after which the result undergoes the fidelity check. What crosses over is a problem framing, a solution architecture, or a revenue logic.

Every element of this procedure reflects a design choice: the size of the population, the mix of operators, the rule deciding which concepts survive and reproduce, the cap on any one family, the number of generations. Designing a genetic algorithm is a craft—which settings work best depends on the specifics of the problem, and volumes have been written on these choices (Goldberg 1989, Mitchell 1996, Eiben and Smith 2003). This paper picks simple, pedagogical values; what matters here is that every choice is explicit, so alternative configurations can be tested and compared to a clear baseline.

Figure 4 traces the complete ancestry—the *lineage*—of the concept with the highest performance in the final generation. Four of the fifteen seeds enter the line through eight




Figure 4: Ancestry of the best concept in global search’s final generation. Percentages report performance against the reference set.

crossovers and eleven single-component mutations spread across eleven of the twelve generations; the final step is an elite copy. Early blends scored as low as 15 percent; later operators matured the line to 40 percent, above all but two seeds. The result exhibits combinatorial novelty: not a variation on any single parent, but a new arrangement whose recorded ancestry traces to multiple seed concepts and whose text visibly combines material from several of them (Boden 2004). Appendix C.6 provides additional detail on how evolution took place, by placing exact passages from its most legible ancestral sources beside the resulting text.

Figure 5 compares the two searches by plotting every concept the prototype generated against the two measures already defined—performance as win rate against the reference set, and semantic distance as distance in meaning to the nearest seed. Local search climbs within a narrow band near its seed; global search occupies a region several times wider, powered by crossovers joining distant parents.




Figure 5: Every concept generated by the prototype, positioned by semantic distance and performance. The solid path is the local search; the dashed path traces global search’s generation winners; the ringed diamond is the concept whose ancestry Figure 4 shows.

The same figure shows what global search does not deliver: immediate performance. The ringed concept reaches 40 percent, and the dashed path of generation winners—winners under the search’s tournament-and-distance selection rule, not the reference-set measure—ends lower still. What happens is that recombination produces many incoherent blends, and even coherent children start without evidence, so they read as sketches beside polished pitches. Research on genetic algorithms shows that recombination typically requires vast search before delivering performance improvements; its returns generally accrue over larger populations and longer horizons than this small tutorial uses (Koza 1992). Global search’s immediate contribution is novelty—the discovery of distant concepts.

The run also illustrates the critical role of the *retention rule*—the selection rule and family cap governing elite survival. By generation 4, a mutation of **clockchain** had reached 78 percent, and elite slots kept it alive through generation 7; by generation 10 it was gone. The loss happened because the selection rule weighted tournament and distance ranks equally and allowed no family more than two elite slots; as distant hybrids rose in the tournament, this

valuable variant was crowded out. Note that the rule that discarded a 78-percent concept was also the rule that kept newborn blends alive. This trade-off shows that alternative rules would trace different performance–distance frontiers; the retention rule is thus a design lever.

In sum, global search’s main effect is the creation of novelty, not immediate performance. It reached concepts several times farther from the seeds than local search did while its final generation topped out at only 40 percent.

3.4 Combining searches

So far, the analyses have shown that local search can refine concepts, while global recombination can supply distant starting points. It is therefore natural to ask whether their benefits can be combined by using global search to create novel concepts and local search to polish them. To test this combination, the pipeline gives each concept in global search’s final generation twelve rounds of local search; I call this last stage *polish*.

Figure 6 shows that the combination is indeed possible and that it greatly raises the performance of the raw concepts global search produced. All fifteen concepts improved, and every polished descendant outranks every seed. The best descendant—the polished concept with four seed ancestors shown in Figure 4, which entered polish at 40 percent—reaches 73 percent and outranks twenty-four of the thirty originals. In sequence, the two searches move the seed median from 33 percent to a polished median of 63 percent, a gain of 30 percentage points.

3.5 Robustness

Could local search alone deliver improvements similar to those delivered by global plus local search? To answer this question, I ran the same twelve-round procedure starting from each of the fifteen seeds. Local-only search achieved higher performance, with final concepts reaching a median of 68 percent and a best of 90 percent, compared with 63 and 73 percent for the polished slate. But it stayed near the seeds—median semantic distance was 0.10 (maximum




Figure 6: Global search’s final generation before (open circles) and after (filled circles) local polish, on the same scale as the thirty originals (tick marks).

0.15), versus 0.35 for the recombined-and-polished slate, whose concepts occupied regions no local run visited. Local refinement and global recombination are therefore complements rather than rivals: refinement raises performance wherever search stands, while recombination can supply starting points outside the regions reached by refinement. In the market, greater semantic distance may increase the chance of reaching a less crowded position—a “blue ocean”—because it moves search away from known venture concepts (Kim and Mauborgne 2005).

Because DeepSeek both generated and judged the concepts, I re-scored the end-of-search concepts—the polished slate and the local-only slate, alongside the thirty originals—using two open-weight models from other developers. Both placed the polished slate’s median above the local-only slate’s, and the focal descendant remained near the top of the field, reinforcing the conclusion that global recombination followed by polish can produce distant, high-performing concepts across evaluators, although exact rankings vary (Appendix C.9).

The exercise in this section thus establishes that venture-concept search can be built and systematically varied, with each mechanism playing a distinct, observable role: local refinement lifted a mid-ranked concept from 47 to 88 percent; global recombination reached concepts more than three times as distant as refinement alone; and the two searches in sequence turned bottom-half seeds into a slate whose best outranks twenty-four of the thirty originals. The prototype is also auditable: its record preserves each candidate, comparison, fidelity decision, and retention choice, allowing researchers to trace how selected concepts emerged, locate where promising alternatives disappeared, and rerun the process under different rules. Overall, the existence of cheap, designable, and inspectable imagination raises the question to which the rest of the paper now turns: what changes for entrepreneurship theory and practice?

4 What changes when imagination becomes programmable

The preceding sections moved from theory to demonstration: Section 2 developed the generative view of candidate-set construction, and Section 3 instantiated and varied that process in software. This section traces what follows across four levels: the loop, whose architecture shapes which possibilities can be reached and which survive; the entrepreneur, who must govern the resulting abundance; the firm, where shared systems can press toward convergence; and the field, which gains an instrument for studying all three.

4.1 The loop: How search architecture shapes the candidate set

Search architecture shapes candidate-set development in two ways: (i) representations and operators open paths from one concept to another, and (ii) retention rules determine which of those paths remain open. What search ultimately reaches depends on both.

First, representations and operators shape the paths available to search. Prior knowledge helps determine where search begins (Shane 2000, Grégoire et al. 2010); representations

specify what can vary, and operators the moves between concepts. In Section 3, local rewrites improved concepts near their seeds, whereas recombination opened routes to regions no local run visited. If the operators provide no path to a possibility, even perfect evaluation cannot recover it—the concept is never judged—so coverage depends not only on how much search is performed but on which paths the architecture makes available.

Second, retention determines which paths remain open long enough to develop. Retained concepts receive further revision and recombination; rejected ones do not. When recombination is possible and developmental capacity is limited, the candidate set works like an ecology: a route to a strong concept may pass through one that initially appears weak, so a concept’s value may lie in what it can later contribute, not only in its current merit. The prototype’s retention rule illustrates the point, discarding a 78-percent `clockchain` variant while preserving weak, distant recombinations that later refinement made competitive. Retention can thus cause omission errors by ending lineages that would have become valuable, and commission errors when weak concepts survive to advance toward action.

The broader implication is that search should not be judged only by the accuracy of its individual evaluations. Opportunity-evaluation research asks whether a focal judgment is correct (Åstebro and Elhedhli 2006, Haynie et al. 2009, Wood and Williams 2014, Scott et al. 2020); the additional question is what candidate set repeated retention produces. A rich set of operators cannot compensate for premature pruning, and generous retention cannot reach possibilities for which no path exists. Search regimes—whether run by founders, investors, or models—should therefore be judged by the quality, coverage, and diversity of the concepts they allow to develop. The next subsection turns to the entrepreneur who must govern this expanded search.

4.2 The entrepreneur: Governing an abundance of possibilities

Cheaper search can create an abundance of possibilities. As the cost of generation, provisional evaluation, and revision falls, a loop once bounded by a founder’s attention can become a

persistent process that maintains many concepts and lineages (HAI 2026, Csaszar 2025). Search may even display a Jevons paradox: lowering its cost may increase the total amount of search rather than reduce it (Jevons 1866). This abundance changes the entrepreneur’s task in three ways.

First, cheaper search eases the exploration–exploitation trade-off. Conventionally, exploration and exploitation compete for the same scarce attention (March 1991). Software-mediated search weakens this constraint: both activities can be purchased partly with computation rather than the entrepreneur’s time, and models draw on knowledge spanning more domains than any person commands—relaxing the limits of attention and experience that bind exploration most. If provisional evaluation is accurate enough to keep volume from becoming noise, exploration can expand without a matching claim on human attention, as when the prototype’s global search reached distant regions that local search then refined. The trade-off remains—computation, reliable evaluation, and opportunities for development are still limited—but its terms change.

Second, entrepreneurial judgment extends from choosing among concepts to governing the search that produces them. Pauling’s “have many ideas” advice in the epigraph has a less quoted continuation: “Most of them will be wrong, and what you have to learn is which ones to throw away.”⁶ AI makes the first part cheap, making the second more consequential. Judgment theories ask whether and when to commit resources under uncertainty (Klein 2008, Foss and Klein 2012). Entrepreneurs must now also decide what to search, how to represent it, how much computation to allocate to exploration and refinement, which generators and evaluators to trust, what novelty to preserve, and when to stop. These are the choices whose effects the preceding subsection traced. Choices once left largely implicit in human thought become explicit and revisable. Judgment still governs the “Commit to Action” arrow in Figure 2, but it must now also govern the process that constructs the alternatives available

⁶Both the epigraph and its continuation are as recounted by Francis Crick in his 1995 lecture “The Impact of Linus Pauling on Molecular Biology.”

for commitment.⁷

Third, evidence from action becomes the scarce complement to cheap conjecture. Computation can multiply venture concepts much faster than the world can reveal customer behavior, stakeholder commitment, or implementation feasibility. Capital, attention, distribution, and trust remain scarce, and an abundance of concepts may intensify competition for them. Offline search therefore does not replace action; its value lies in directing scarce real-world tests toward the uncertainties computation cannot resolve.⁸ Deciding which uncertainties to address offline and which to test in the world becomes a central entrepreneurial task, making theories of entrepreneurial action more, not less, important (Alvarez and Barney 2007, Berglund et al. 2020, Sarasvathy 2001, Baker and Nelson 2005, Zellweger and Zenger 2023). Whether firms do this in ways that make their searches more similar or more distinct is the next question.

4.3 The firm: Correlated imagination and the sources of differentiation

At the firm level, programmable imagination creates two opposing forces: one toward convergence and another toward differentiation.

First, more imagination within firms may produce less diversity across firms. When firms approach similar problems using the same foundation models, public information, representations, evaluators, and retention rules, each may consider more alternatives while their surviving candidate sets may increasingly occupy the same regions. I call this *correlated imagination*: AI can broaden candidate sets within firms while making candidate sets across firms more alike. This outcome does not require models to generate identical ideas: even when generation differs, shared evaluators can repeatedly favor the same attributes and steer

⁷In this same issue, Zott and Amit (2026) propose a new area judgment must also govern: after commitment, how much of the business model AI may adapt on its own.

⁸Also in this issue, Li et al. (2026) make a parallel point about AI agents that act autonomously: their intelligence grows from the operational experience that only execution can generate.

later generation in similar directions. Findings that generative AI can increase individual creativity while making outputs more alike show the same pattern in simpler settings (Doshi and Hauser 2024, Terwiesch et al. 2026, Lu et al. 2026); whether repeated conceptual search amplifies or counteracts it is an empirical question. Research on this question can measure breadth within each firm’s candidate set separately from overlap across firms, because greater imagination by each organization need not produce greater variety in an industry.

Second, differences may increasingly come from what firms bring to the search. Proprietary problems, data, and feedback from action, combined with firm-specific representations, evaluators, and retention rules, can send otherwise similar systems along different trajectories. Action-generated evidence matters especially because rivals cannot reproduce it merely by using the same model: it reflects a firm’s own customers, experiments, assets, relationships, and history. This sharpens resource-based and dynamic-capabilities perspectives. Resources and capabilities shape not only which ventures firms can execute or regard as feasible, but also which possibilities their searches generate, preserve, and develop (Barney 1991, Teece et al. 1997, Camuffo et al. 2026). As capable foundation models diffuse, durable differences may depend increasingly on proprietary inputs, accumulated evidence, and organizational routines rather than on access to the model alone.

4.4 The field: A research agenda for entrepreneurial search

Conceptual search lets researchers observe and test candidate-set construction. It preserves each candidate, transformation, comparison, retention decision, and abandoned branch. Researchers can hold starting concepts and information fixed, change one part of the search process, and compare the trajectories that result. Questions about how venture concepts should be searched can therefore become subjects of cumulative experimentation rather than argument alone.

The prototype samples this possibility narrowly: one platform and product category, thirty ventures, fifteen seeds, and twelve-step horizons. Its interpretation has limits as well. Rankings

are evaluator-relative; stronger articulation need not imply stronger venture substance (Kirsch et al. 2009); semantic distance does not establish economic or transformational novelty; and evidence validating model judgment on these tasks need not transfer to harder settings. These boundaries, together with the possibilities the method creates, motivate six questions: three about the loop, then one each about the entrepreneur, the firm, and the wider field.

1. *When is provisional evaluation accurate enough to guide repeated search?* Following Csaszar et al. (2026), prospective benchmarks can compare provisional evaluations with outcomes realized later and identify what drives accuracy; in that study, expert-level reasoning was the only measured factor reliably associated with higher accuracy. Studies could compare general and specialized models with human and hybrid evaluators across industries, venture stages, representations, and degrees of novelty.

2. *How should the rest of the search process be designed?* Question 1 concerns the evaluator; this question concerns the loop’s other design choices, such as the supplied representation, operators, population size, novelty incentives, retention rules, and sequence of exploration and refinement. Genetic search alone offers many adjustable settings and can be combined with other search methods (Michalewicz and Fogel 2004). Researchers could compare designs by the quality, coverage, diversity, robustness, and computational cost of the candidate sets they produce. The aim is not one universally best algorithm, but an account of which designs work in which settings.

3. *Can search change the representation it searches?* Question 2 allows researchers to supply different representations from the outside. This question asks whether search can revise its representation from within by creating new categories, dimensions, or possibilities rather than only recombining existing elements (Boden 2004, Rindova and Martins 2024). This requires reliable methods for detecting transformation rather than recombination and, if current systems rarely achieve it, methods intended to produce it (Csaszar and Levinthal 2016, Rindova and Martins 2021). The central question is whether, and under what conditions, conceptual search can systematically expand and transform the space it searches.

4. *How should entrepreneurs combine conceptual search with human judgment and action?* Computational experiments can test which findings about representation, evaluation, retention, and sequence transfer to unaided human search. Human–AI studies can examine which interfaces make search histories understandable and which teaching practices help entrepreneurs learn from the process and become better at generating, evaluating, revising, and deciding when to commit to action.

5. *When does shared AI create correlated imagination, and what sustains difference?* Question 2 studies one search; this question studies many searches across rival firms, where the outcome of interest is overlap across candidate sets rather than the coverage of any one set. Studies could vary the models, information, and search architectures rival firms use, along with the feedback each obtains from action. The aim is not to assume that overlap is harmful, but to identify when shared search systems homogenize candidate sets and what sustains distinct trajectories.

6. *How can conceptual search improve computational models in organization theory?* Conceptual search offers a new infrastructure for computational models in which positions, moves, and evaluations are substantively meaningful. The searched objects need not be venture concepts: the same apparatus can operate over representations of firm strategies, organization designs, or other types of alternatives. It can extend NK-style models beyond abstract bit strings and fixed operators (Levinthal 1997). It can also extend evolutionary and history-friendly models by making the generation and transformation of innovations endogenous (Nelson and Winter 1982, Malerba et al. 2016). Researchers can then compare search processes across representations, organizations, and institutional settings, preserve candidate histories, and run interpretable counterfactuals over paths not taken (for complementary simulation-based approaches, see Choudhury et al. 2026, Tranchero et al. 2026).

Across these questions, particular models and algorithms will be superseded; the research objects will not. What endures are candidate sets and their evolution, the rules that shape them, and the feedback from action that tests both the concepts and the rules.

4.5 Conclusion: The construction of possibilities

This paper began with a blank image of one million pixels. The technical ingredients needed to sell those pixels had existed for years; what had been missing was a represented possibility connecting them to a venture. The generative view makes the construction of such possibilities an explicit object of entrepreneurship theory; conceptual search makes that theory computational.

A classic entrepreneurial problem is which possibility to pursue. The generative view adds an earlier one: how the possibilities available for pursuit came to exist in the first place. AI does not solve that problem, but it makes candidate-set construction increasingly scalable, programmable, and auditable. Entrepreneurial imagination, once largely hidden inside cognition and conversation, can now become the subject of detailed computational inquiry. Imagination is only one force separating the possible and realized frontiers, but making it programmable gives researchers a new way to study how entrepreneurial possibilities are constructed and entrepreneurs a new way to expand what gets imagined—and, potentially, what gets realized.

References

- Alvarez, S. A., J. B. Barney. 2007. Discovery and creation: Alternative theories of entrepreneurial action. *Strategic Entrepreneurship Journal* **1**(1–2) 11–26.
- Åstebro, T., S. Elhedhli. 2006. The effectiveness of simple decision heuristics: Forecasting commercial success for early-stage ventures. *Management Science* **52**(3) 395–409.
- Axelrod, R. M. 1984. *The Evolution of Cooperation*. Basic Books, New York.
- Baker, T., R. E. Nelson. 2005. Creating something from nothing: Resource construction through entrepreneurial bricolage. *Administrative Science Quarterly* **50**(3) 329–366.
- Barney, J. 1991. Firm resources and sustained competitive advantage. *Journal of Management* **17**(1) 99–120.
- Baumol, W. J. 1990. Entrepreneurship: Productive, unproductive, and destructive. *Journal of Political Economy* **98**(5) 893–921.
- Beaty, R. E., D. R. Johnson. 2021. Automating creativity assessment with SemDis: An open platform for computing semantic distance. *Behavior Research Methods* **53**(2) 757–780.
- Berglund, H., M. Bousfiha, Y. Mansoori. 2020. Opportunities as artifacts and entrepreneurship as design. *Academy of Management Review* **45**(4) 825–846.
- Boden, M. A. 2004. *The Creative Mind: Myths and Mechanisms*. 2nd ed. Routledge, London, UK.
- Boussiou, L., J. N. Lane, M. Zhang, V. Jacimovic, K. R. Lakhani. 2024. The crowdless future? Generative AI and creative problem-solving. *Organization Science* **35**(5) 1589–1607.
- Bruderer, E., J. V. Singh. 1996. Organizational evolution, learning, and selection: A genetic-algorithm-based model. *Academy of Management Journal* **39**(5) 1322–1349.
- Burgelman, R. A. 1991. Intraorganizational ecology of strategy making and organizational adaptation: Theory and field research. *Organization Science* **2**(3) 239–262.
- Campbell, D. T. 1960. Blind variation and selective retentions in creative thought as in other knowledge processes. *Psychological Review* **67**(6) 380–400.
- Camuffo, A., A. Gambardella, S. Kazemi, A. Pandey. 2026. Beyond black boxes: Designing and testing agentic AI systems for strategy. *Strategy Science* **11**(1) 137–156.
- Chatterji, A., F. A. Csaszar, J. Evans, T. Felin, J. Hullman, K. R. Lakhani, M. Sako, T. Zenger. 2026. Can AI do strategy? A dialogue and debate. *Strategy Science* **11**(1) 16–30.
- Chen, J., S. Xiao, P. Zhang, K. Luo, D. Lian, Z. Liu. 2024. M3-embedding: Multi-linguality, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. arXiv:2402.03216.
- Choudhury, P., Y. Du, M. Lawrence. 2026. Learning-by-twinning: Organizational digital twins and the future of search. *Strategic Entrepreneurship Journal* **20**(4) xxx–xxx.
- Cornelissen, J. P., J. S. Clarke. 2010. Imagining and rationalizing opportunities: Inductive reasoning and the creation and justification of new ventures. *Academy of Management Review* **35**(4) 539–557.
- Craik, K. J. W. 1943. *The Nature of Explanation*. Cambridge University Press, Cambridge, UK.
- Csaszar, F. A. 2018. What makes a decision strategic? Strategic representations. *Strategy Science* **3**(4) 606–619.
- Csaszar, F. A. 2025. Unbounding rationality: Why AI is a fundamental issue for strategy. Chapter in a forthcoming book edited by Barney, J. B. and Zenger, T. R. Available at SSRN: <https://ssrn.com/abstract=5454634>.

- Csaszar, F. A., N. Hinrichs, M. Heshmati. 2024a. External representations in strategic decision making: Understanding strategy’s reliance on visuals. *Strategic Management Journal* **45**(11) 2191–2226.
- Csaszar, F. A., H. Ketkar, H. Kim. 2024b. Artificial intelligence and strategic decision-making: Evidence from entrepreneurs and investors. *Strategy Science* **9**(4) 322–345.
- Csaszar, F. A., D. A. Levinthal. 2016. Mental representation and the discovery of new strategies. *Strategic Management Journal* **37**(10) 2031–2049.
- Csaszar, F. A., A. Peterson, D. Wilde. 2026. The strategic foresight of LLMs: Evidence from a fully prospective venture tournament. arXiv:2602.01684.
- David, H. A. 1988. *The Method of Paired Comparisons*. 2nd ed. Oxford University Press, New York.
- Davidsson, P. 2015. Entrepreneurial opportunities and the entrepreneurship nexus: A re-conceptualization. *Journal of Business Venturing* **30**(5) 674–695.
- Dell’Acqua, F., C. Ayoubi, H. Lifshitz, R. Sadun, E. Mollick, L. Mollick, Y. Han, J. Goldman, H. Nair, S. Taub, K. R. Lakhani. 2026. The cybernetic teammate: A field experiment on generative AI and teamwork. *Organization Science* **37**(4) 1217–1242.
- Doshi, A. R., J. J. Bell, E. Mirzayev, B. S. Vanneste. 2025. Generative artificial intelligence and evaluating strategic decisions. *Strategic Management Journal* **46**(3) 583–610.
- Doshi, A. R., O. P. Hauser. 2024. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances* **10**(28) eadn5290.
- Dyson, G. 1998. *Darwin Among the Machines: The Evolution of Global Intelligence*. Perseus Books, Reading, MA.
- Eiben, A. E., J. E. Smith. 2003. *Introduction to Evolutionary Computing*. Springer, New York.
- Felin, T., A. Gambardella, E. Novelli, T. Zenger. 2024. A scientific method for startups. *Journal of Management* **50**(8) 3080–3104.
- Felin, T., M. Holweg. 2024. Theory is all you need: AI, human cognition, and causal reasoning. *Strategy Science* **9**(4) 346–371.
- Felin, T., T. R. Zenger. 2017. The theory-based view: Economic actors as theorists. *Strategy Science* **2**(4) 258–271.
- Finke, R. A., T. B. Ward, S. M. Smith. 1992. *Creative Cognition: Theory, Research, and Applications*. MIT Press, Cambridge, MA.
- Foss, N. J., P. G. Klein. 2012. *Organizing Entrepreneurial Judgment: A New Approach to the Firm*. Cambridge University Press, Cambridge, UK.
- Foss, N. J., P. G. Klein, L. B. Lien, T. Zellweger, T. Zenger. 2021. Ownership competence. *Strategic Management Journal* **42**(2) 302–328.
- Frederiks, A. J., B. G. Englis, M. L. Ehrenhard, A. J. Groen. 2019. Entrepreneurial cognition and the quality of new venture ideas: An experimental approach to comparing future-oriented cognitive processes. *Journal of Business Venturing* **34**(2) 327–347.
- Gans, J. S., S. Stern, J. Wu. 2019. Foundations of entrepreneurial strategy. *Strategic Management Journal* **40**(5) 736–756.
- Gartner, W. B. 2007. Entrepreneurial narrative and a science of the imagination. *Journal of Business Venturing* **22**(5) 613–627.
- Gavetti, G., D. A. Levinthal. 2000. Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly* **45**(1) 113–137.

- Gennai, J., R. Katila, M. Kumar, I. Sinapuelas. 2026. Everything everywhere all-at-once? AI and strategy trade-offs in the age of abundance. *Strategic Entrepreneurship Journal* **20**(4) xxx–xxx.
- Goldberg, D. E. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, MA.
- Gompers, P. A., W. Gornall, S. N. Kaplan, I. A. Strebulaev. 2020. How do venture capitalists make decisions? *Journal of Financial Economics* **135**(1) 169–190.
- Grégoire, D. A., P. S. Barr, D. A. Shepherd. 2010. Cognitive processes of opportunity recognition: The role of structural alignment. *Organization Science* **21**(2) 413–431.
- Gruber, M., I. C. MacMillan, J. D. Thompson. 2008. Look before you leap: Market opportunity identification in emerging technology firms. *Management Science* **54**(9) 1652–1665.
- HAI. 2026. The AI index 2026 annual report. Institute for Human-Centered AI, Stanford University. Available at: <https://hai.stanford.edu/ai-index/2026-ai-index-report> (accessed August 15, 2026).
- Haynie, J. M., D. A. Shepherd, J. S. McMullen. 2009. An opportunity for me? The role of resources in opportunity evaluation decisions. *Journal of Management Studies* **46**(3) 337–361.
- Holland, J. H. 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI.
- Jevons, W. S. 1866. *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-Mines*. 2nd ed. Macmillan, London, UK.
- Kaplan, S. N., B. A. Sensoy, P. Strömberg. 2009. Should investors bet on the jockey or the horse? Evidence from the evolution of firms from early business plans to public companies. *Journal of Finance* **64**(1) 75–115.
- Kauffman, S. A. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York.
- Kerr, W. R., R. Nanda, M. Rhodes-Kropf. 2014. Entrepreneurship as experimentation. *Journal of Economic Perspectives* **28**(3) 25–48.
- Kier, A. S., J. S. McMullen. 2018. Entrepreneurial imaginativeness in new venture ideation. *Academy of Management Journal* **61**(6) 2265–2295.
- Kier, A. S., N. A. Thompson, T. H. Chiles, J. P. Cornelissen. 2026. Entrepreneurial imagination: A pluralistic scoping review of a sprawling literature. *Journal of Business Venturing* **41**(1) 106555.
- Kim, W. C., R. Mauborgne. 2005. *Blue Ocean Strategy*. Harvard Business School Press, Boston, MA.
- Kirsch, D., B. Goldfarb, A. Gera. 2009. Form or substance: The role of business plans in venture capital decision making. *Strategic Management Journal* **30**(5) 487–515.
- Klein, P. G. 2008. Opportunity discovery, entrepreneurial action, and economic organization. *Strategic Entrepreneurship Journal* **2**(3) 175–190.
- Koza, J. R. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA.
- Lachmann, L. M. 1976. From Mises to Shackle: An essay on Austrian economics and the Kaleidic society. *Journal of Economic Literature* **14**(1) 54–62.
- Langley, P., H. A. Simon, G. L. Bradshaw, J. M. Zytkow. 1987. *Scientific Discovery: Computational Explorations of the Creative Processes*. MIT Press, Cambridge, MA.
- Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Science* **43**(7) 934–950.

- Levinthal, D. A. 2011. A behavioral approach to strategy: What’s the alternative? *Strategic Management Journal* **32**(13) 1517–1523.
- Li, Y., M. Zeng, A. Li. 2026. The intelligence loop and value creation in the age of AI agents. *Strategic Entrepreneurship Journal* **20**(4) xxx–xxx.
- Lu, J. G., G. G. Zhao, A. M. Zheng. 2026. Generative AI use in entrepreneurship: An integrative review and an empowerment–entrapment framework. *Strategic Entrepreneurship Journal* **20**(4) xxx–xxx.
- Malerba, F., R. R. Nelson, L. Orsenigo, S. G. Winter. 2016. *Innovation and the Evolution of Industries: History-Friendly Models*. Cambridge University Press, Cambridge, UK.
- March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science* **2**(1) 71–87.
- March, J. G., H. A. Simon. 1958. *Organizations*. John Wiley & Sons, New York.
- Martins, L. L., V. P. Rindova, B. E. Greenbaum. 2015. Unlocking the hidden value of concepts: A cognitive approach to business model innovation. *Strategic Entrepreneurship Journal* **9**(1) 99–117.
- McBride, R. 2026. Mean articulation machines. *Strategy Science* **11**(1) 31–54.
- Michalewicz, Z., D. B. Fogel. 2004. *How to Solve It: Modern Heuristics*. 2nd ed. Springer, Berlin, Germany.
- Mitchell, M. 1996. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA.
- Nelson, R. R., S. G. Winter. 1982. *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge, MA.
- Newell, A., H. A. Simon. 1956. The logic theory machine: A complex information processing system. *IRE Transactions on Information Theory* **2**(3) 61–79.
- Newell, A., H. A. Simon. 1972. *Human Problem Solving*. Prentice-Hall, Englewood Cliffs, NJ.
- Nickerson, J. A., T. R. Zenger. 2004. A knowledge-based theory of the firm: The problem-solving perspective. *Organization Science* **15**(6) 617–632.
- Novikov, A., N. Vũ, M. Eisenberger, E. Dupont, P.-S. Huang, A. Z. Wagner, S. Shirobokov, B. Kozlovskii, F. J. R. Ruiz, A. Mehrabian, M. P. Kumar, A. See, S. Chaudhuri, G. Holland, A. Davies, S. Nowozin, P. Kohli, M. Balog. 2025. AlphaEvolve: A coding agent for scientific and algorithmic discovery. arXiv:2506.13131.
- Ocasio, W. 1997. Towards an attention-based view of the firm. *Strategic Management Journal* **18**(Summer Special Issue) 187–206.
- Osterwalder, A., Y. Pigneur, T. Clark. 2010. *Business Model Generation: A Handbook for Visionaries, Game Changers, and Challengers*. Wiley, Hoboken, NJ.
- Ouyang, L., J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. Christiano, J. Leike, R. Lowe. 2022. Training language models to follow instructions with human feedback. arXiv:2203.02155.
- Packard, M. D., B. B. Clark, P. G. Klein. 2017. Uncertainty types and transitions in the entrepreneurial process. *Organization Science* **28**(5) 840–856.
- Panickssery, A., S. R. Bowman, S. Feng. 2024. LLM evaluators recognize and favor their own generations. *Advances in Neural Information Processing Systems*, vol. 37. 68772–68802.
- Penrose, E. T. 1959. *The Theory of the Growth of the Firm*. Wiley, New York.

- Ries, E. 2011. *The Lean Startup*. Crown Business, New York.
- Rindova, V. P., L. L. Martins. 2021. Shaping possibilities: A design science approach to developing novel strategies. *Academy of Management Review* **46**(4) 800–822.
- Rindova, V. P., L. L. Martins. 2024. The imagination advantage: Why and how strategists combine knowledge and imagination in developing theories. *Strategy Science* **9**(4) 499–514.
- Romera-Paredes, B., M. Barekatin, A. Novikov, M. Balog, M. P. Kumar, E. Dupont, F. J. R. Ruiz, J. S. Ellenberg, P. Wang, O. Fawzi, P. Kohli, A. Fawzi. 2024. Mathematical discoveries from program search with large language models. *Nature* **625**(7995) 468–475.
- Romme, A. G. L. 2003. Making a difference: Organization as design. *Organization Science* **14**(5) 558–573.
- Sah, R. K., J. E. Stiglitz. 1986. The architecture of economic systems: Hierarchies and polyarchies. *American Economic Review* **76**(4) 716–727.
- Sarasvathy, S. D. 2001. Causation and effectuation: Toward a theoretical shift from economic inevitability to entrepreneurial contingency. *Academy of Management Review* **26**(2) 243–263.
- Schumpeter, J. A. 1934. *The Theory of Economic Development: An Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle*. Harvard University Press, Cambridge, MA.
- Scott, E. L., P. Shu, R. M. Lubynsky. 2020. Entrepreneurial uncertainty and expert evaluation: An empirical analysis. *Management Science* **66**(3) 1278–1299.
- Shackle, G. L. S. 1979. *Imagination and the Nature of Choice*. Edinburgh University Press, Edinburgh, UK.
- Shane, S. 2000. Prior knowledge and the discovery of entrepreneurial opportunities. *Organization Science* **11**(4) 448–469.
- Shepherd, D. A., A. Majchrzak. 2022. Machines augmenting entrepreneurs: Opportunities (and threats) at the nexus of artificial intelligence and entrepreneurship. *Journal of Business Venturing* **37**(4) 106227.
- Simon, H. A. 1969. *The Sciences of the Artificial*. MIT Press, Cambridge, MA.
- Simon, H. A. 1973. The structure of ill structured problems. *Artificial Intelligence* **4**(3–4) 181–201.
- Skalse, J., N. H. R. Howe, D. Krashenninikov, D. Krueger. 2022. Defining and characterizing reward gaming. *Advances in Neural Information Processing Systems*, vol. 35. 9460–9471.
- Spaulding, S. 2016. Imagination through knowledge. A. Kind, P. Kung, eds., *Knowledge Through Imagination*. Oxford University Press, Oxford, UK, 207–226.
- Teece, D. J., G. Pisano, A. Shuen. 1997. Dynamic capabilities and strategic management. *Strategic Management Journal* **18**(7) 509–533.
- Terwiesch, C., L. Meincke, K. Girotra, E. Mollick, G. Nave, K. T. Ulrich. 2026. AI and its impact on creativity and diversity: An empirical study of LLM-generated product ideas. arXiv:2607.27553.
- Tranchoero, M., C. Brenninkmeijer, A. Murugan, A. Nagaraj. 2026. Simulating strategic interactions with AI agents. *Strategic Management Journal* (forthcoming).
- von Neumann, J. 1966. *Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana, IL. A. W. Burks, ed.
- Weick, K. E. 1979. *The Social Psychology of Organizing*. 2nd ed. Addison-Wesley Publishing Company, Reading, MA.
- Wood, M. S., D. W. Williams. 2014. Opportunity evaluation as rule-based decision making. *Journal of Management Studies* **51**(4) 573–602.

- Wuebker, R., T. Zenger, T. Felin. 2023. The theory-based view: Entrepreneurial microfoundations, resources, and choices. *Strategic Management Journal* **44**(12) 2922–2949.
- Zahra, S. A. 2008. The virtuous cycle of discovery and creation of entrepreneurial opportunities. *Strategic Entrepreneurship Journal* **2**(3) 243–257.
- Zellweger, T., T. Zenger. 2023. Entrepreneurs as scientists: A pragmatist approach to producing value out of uncertainty. *Academy of Management Review* **48**(3) 379–408.
- Zhang, S. X., G. Ertug. 2025. Organization research as an applied science: Lessons from fields that shape practice and policy. *Organization Science* **36**(5) 2028–2039.
- Zheng, L., W. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. P. Xing, H. Zhang, J. E. Gonzalez, I. Stoica. 2023. Judging LLM-as-a-judge with MT-bench and Chatbot Arena. arXiv:2306.05685.
- Zott, C., R. Amit. 2026. Whither business model innovation strategy in the age of AI? *Strategic Entrepreneurship Journal* **20**(4) xxx–xxx.

Appendix A: Construct definitions

Table A.1 collects the paper’s constructs and identifies where each is defined or developed.

Construct	Definition
Possible and realized frontiers (§1.1)	The boundary of what could in principle be attempted given current technology, regulation, and infrastructure; the boundary of what has been imagined, designed, and brought to market. The feasible points bounded by the former constitute the <i>possible venture space</i> (§2.1).
Venture concept (§2.1)	A representation of a possible venture, held in mind or externalized as, for example, a product concept, venture plan, or business model canvas.
Theory of value (§2.1)	A causally explicit venture concept: a conjecture about how activities and assets create and capture value.
Represented candidate set (§2.1)	The evolving set of venture concepts under consideration; the paper thereafter uses <i>candidate set</i> as shorthand.
Candidate-set construction (§2.1)	The process by which venture concepts enter, change, remain in, or leave the represented candidate set.
Offline search and online action (§2.1)	Operations over representations before consequential action; implementation and learning through action in the environment.
Entrepreneurial imagination (§2.1)	The loop of generating, provisionally evaluating, and revising venture concepts before or between consequential actions.
Generative view (§1.2)	The view that the ventures entrepreneurs pursue are shaped by the evolving candidate set they construct through generation, provisional evaluation, and revision.
Cognitive scarcity (§2.1)	The attention, memory, and processing limits that constrain how many alternatives are considered and how fully each is developed.
Omission error (§2.1)	A valuable possibility is never generated or is wrongly discarded.
Commission error (§2.1)	Provisional evaluation advances a concept that is infeasible or would create less value than a rejected alternative.
Conceptual search (§3)	A programmable form of offline search in which positions are venture concepts expressed in text, moves are rewrites and recombinations, and evaluation is comparative judgment of substantive promise.
Search architecture (§1.2)	The representation, operators, evaluator, and rules that govern search.
Auditability (§2.3)	The capacity to reconstruct, from an external record, what a search considered and why concepts survived or disappeared.
Correlated imagination (§4.3)	The possibility that AI broadens candidate sets within firms while making candidate sets across firms more similar.

Table A.1: The paper’s constructs and their definitions.

Appendix B: Reference set

Table B.1 lists the 30 ventures used to evaluate the search outputs, ordered by the pairwise evaluator. The ruler separates the withheld top half from the 15 ventures used as seeds. Names are pseudonyms (as used in Csaszar et al. 2026); all dates are in 2025.

Venture (Performance)	Brief summary URL	Goal Raised	Start End
1 dreamie (96.5%)	Bedside sleep device with audio, sunrise alarms, and private listening kickstarter.com/projects/ambient/dreamie	\$10,000	Oct 7
2 fraction (94.8%)	Modular, repairable vacuum with self-diagnostics and UVC disinfection kickstarter.com/projects/deglace/fraction-the-worlds-first-modular-vacuum-built-to-last	\$15,000	Sep 30
3 prolo-ring (86.2%)	Wearable ring giving cursor control, shortcuts, and air gestures kickstarter.com/projects/prolo/prolo-ring-precision-control-for-keyboard-power-users	\$428,684	Nov 14
4 ubo-pod (84.5%)	Open-source Raspberry-Pi device for voice-and-vision assistants kickstarter.com/projects/ubopod/ubo-pod-hackable-personal-ai-assistant	\$10,000	Oct 15
5 pictomic (84.5%)	Computer-controlled macro rail that automates precise focus stacking kickstarter.com/projects/pictomic/pictomic-macro-photography-rail-slider	\$481,926	Nov 18
6 humanecheck (84.5%)	App that scans barcodes to grade animal welfare and suggest alternatives kickstarter.com/projects/496302772/humanecheck-scan-barcodes-for-animal-welfare-grades	\$25,000	Oct 8
7 lifesummary (72.4%)	AI service answering questions about insurance and estate documents kickstarter.com/projects/797876113/lifesummaryai-your-lifes-most-important-documents	\$37,200	Nov 7
8 kujietool (72.4%)	Compact desktop milling machine with professional-grade accuracy kickstarter.com/projects/564580014/kujietool-the-compact-mill-for-makers	\$49,900	Sep 8
9 toolkaiser (70.7%)	Software that scans hand tools and exports precisely scaled outline files kickstarter.com/projects/gridmat/toolkaiser	\$2,998	Nov 9
10 robotin-r2 (66.4%)	Autonomous appliance that deep-vacuums, machine-washes, and dries carpets kickstarter.com/projects/robotin/robotin-r2-worlds-first-robot-carpet-cleaner	\$25,000	Oct 8
11 cyberdesk (63.8%)	Case-management platform for cybersecurity investigations and reporting kickstarter.com/projects/cyberdefensedesk/cyberdefensedesk	\$13,172	Nov 7
12 cop-rights (63.8%)	Free app with state-specific guidance on drivers' rights at traffic stops kickstarter.com/projects/jersey/what-cops-can-and-cannot-do	\$35,000	Oct 15
13 pickup (49.1%)	App for finding quality pickup basketball games and tracking performance kickstarter.com/projects/pickupapp/pickup-the-basketball-app	\$8,773	Nov 14
14 zulni (48.3%)	Journaling companion with stress-aware prompts and optional biofeedback kickstarter.com/projects/fwa-zulni/zulni-journaling-with-biofeedbacks-for-calm-and-clarity-app	\$10,000	Sep 14
15 soloist (46.6%)	Loop-pedal-style app where musicians record ideas others grow into tracks kickstarter.com/projects/soloistapp/soloist-infinite-loop-pedal	\$138,547	Nov 2
16 clockchain (46.6%)	Tool anchoring a file's fingerprint to a public time-stamped ledger kickstarter.com/projects/315841049/clockchain-proving-whats-real-online	\$10,000	Oct 1
17 hply (44.8%)	Donation app connecting givers with vetted nonprofits and impact updates kickstarter.com/projects/546557384/hply-smarter-giving-made-simple	\$12,869	Nov 1
18 angelry (39.7%)	Heirloom jewelry of leather cultivated in bioreactors from cow skin cells kickstarter.com/projects/1687452435/angelry-the-worlds-first-cultivated-leather-jewelry	\$10,000	Sep 23
19 lockguard (37.9%)	Portable home-security device installing without tools or Wi-Fi kickstarter.com/projects/lockguard/lockguard-smart-security-that-travels-with-you	\$74,373	Nov 22
20 sale-finder (36.2%)	App mapping local yard and estate sales live, with routes and rewards kickstarter.com/projects/salearoundthecorner/the-sale-around-the-corner-map-it-find-it-flip-it	\$20,000	Sep 15
21 genaix (36.2%)	Learning platform for human-centered AI education with an AI mentor kickstarter.com/projects/1327323563/genaix-20-building-the-future-of-human-centered-ai-learning	\$2,245	Nov 14
22 bugout-battery (34.5%)	Rugged off-grid battery-and-solar system for emergencies and RVs kickstarter.com/projects/bugoutbattery/not-a-toy-bugoutbattery-all-in-one-off-grid-solar-power	\$15,000	Sep 24
23 rent-a-bee (32.8%)	Hive-sponsorship subscription linking supporters with local beekeepers kickstarter.com/projects/1252482382/rent-a-bee-bridging-nature-and-technology	\$25	Nov 8
24 hooper (31.0%)	App where basketball players find games, check in, and build rankings kickstarter.com/projects/1748255141/hooper-app-to-find-games-rank-up-and-connect-hoopers	\$15,500	Oct 8
25 remotion-ai (20.7%)	AI video analysis delivering sprint-coaching feedback on athletes' form kickstarter.com/projects/remotionai/remotion-ai	\$15,600	Nov 7
26 viaia (19.0%)	AI travel assistant for trustworthy local services in Venezuela kickstarter.com/projects/viaia/viaia-connecting-travelers-with-local-experiences	\$10,000	Sep 25
27 lumivisor (19.0%)	Transparent LED film for helmet visors that displays expressions kickstarter.com/projects/lumivisor/lumivisor-flexible-led-film-for-helmet-visors	\$282	Nov 9
28 voyagx (13.8%)	Social travel app for planning trips, splitting costs, meeting companions kickstarter.com/projects/voyagxapp/voyagx-where-travelers-connect	\$10,000	Oct 7
29 kebo (3.5%)	Spring-loaded bottle opener using kinetic energy to eject caps cleanly kickstarter.com/projects/kebottleopener/kebo	\$10,000	Oct 9
30 my-story (0.0%)	Children's storytelling system pairing a print book with an AR app kickstarter.com/projects/specialprincess/my-story-my-way	\$11,165	Nov 8
		\$44,500	Oct 14
		\$50	Nov 13
		\$35,000	Oct 7
		\$6,604	Nov 6
		\$20,000	Sep 9
		\$21,821	Nov 8
		\$25,000	Oct 9
		\$752	Nov 15
		\$50,000	Sep 22
		\$195	Nov 21
		\$25,000	Oct 6
		\$125	Nov 5
		\$24,997	Sep 30
		\$451	Nov 29
		\$60,000	Oct 14
		\$573	Nov 18
		\$10,000	Oct 15
		\$10	Nov 14
		\$10,000	Oct 15
		\$10	Nov 14
		\$13,200	Oct 7
		\$14	Nov 21
		\$15,000	Oct 15
		\$20,907	Nov 14
		\$28,000	Oct 13
		\$35	Nov 12
		\$50,000	Sep 17
		\$715	Nov 1
		\$250,000	Oct 14
		\$808	Nov 13

Table B.1: Reference set.

Appendix C: Method details

This appendix explains how concepts were generated and selected, and how performance was measured, in the conceptual-search exercise described in Section 3. More detailed implementation information, code, and search records are available at <https://github.com/csaszar/conceptual-search>.

C.1 Reference and seed sets

The reference set contains the same 30 companies and campaign summaries used by Csaszar et al. (2026): U.S. Kickstarter Technology campaigns launched in 2025, after the training cutoff of every model in that study, including DeepSeek V3.2. Each campaign had a funding goal of at least \$10,000. Names were replaced with pseudonyms and campaign pages were converted to descriptions of approximately 500 words. Table B.1 lists the projects and their outcomes, which were withheld from the search and evaluation agents.

The experimental evaluator first ranked the 30 original descriptions in a double round-robin tournament. The bottom 15 formed the global-search seed set. Local search began from `clockchain`, the highest-ranked seed.

C.2 Model calls and tournament evaluation

Generation and evaluation used DeepSeek V3.2 (`deepseek/deepseek-v3.2`) through OpenRouter. Its open weights permit local replication, its training cutoff preceded the campaigns, and the companion study found it achieved strong prospective judgment on these ventures: its rankings reached a rank correlation of 0.62 with realized fundraising, correctly ordered 72 percent of venture pairs, and captured 83 percent of attainable fundraising value in its top-five selections—above every individual expert on each measure (Csaszar et al. 2026). Calls used a sampling temperature of 0.5, except fidelity checks at 0.0; optional reasoning was disabled. Responses had a 2,000-token limit and a 120-second timeout. The client retried

Item	Calculation	Calls	Cost
Focal local search	12 rounds $\times$ [1 advisor call (to propose 5 changes) + 5 rewriter calls (one neighboring concept per change) + 5 fidelity checks + 30 ordered comparisons (1 incumbent + 5 variants, each facing the other 5) + 1 narrative-summary call] $-$ 14 duplicate comparisons + 11 retries	501	\$0.25
Global search	13 population tournaments (initial population and 12 evolved generations) $\times$ 15 concepts $\times$ 14 opponents + 43 concepts retained as elites or generation winners $\times$ 15 seeds $\times$ 2 presentation orders + 5 final mutation candidates $\times$ 15 seeds $\times$ 2 presentation orders + (1 selected global concept + 15 seeds) $\times$ 15 opponents + 12 generations $\times$ (5 mutations + 5 crossovers + 10 fidelity checks) + 5 final mutations + 5 fidelity checks + 20 narrative-summary calls $-$ 42 duplicate comparisons + 198 retries	4,836	\$2.72
Polish searches	15 concepts in global search’s final generation $\times$ 504 calls per 12-round local search $-$ 593 duplicate comparisons + 131 retries	7,098	\$3.94
Local-only comparison	14 remaining seeds $\times$ 504 calls per 12-round local search + 49 tournament entrants $\times$ 48 opponents $-$ 261 duplicate comparisons + 248 retries	9,395	\$4.72
Scoring against reference set	171 distinct concepts generated or retained in the focal local and global searches $\times$ 30 reference ventures $\times$ 2 presentation orders $-$ 30 exact-text ties + 482 retries	10,712	\$5.23
Final comparison tournament	[30 originals + 15 polished descendants + 14 distinct pre-polish parents (the fifteenth duplicated another parent) + 1 focal local final + 1 global champion + 1 best hybrid + 2 global runners-up] $\times$ 63 other entrants + 194 retries	4,226	\$1.72
Overall	501 + 4,836 + 7,098 + 9,395 + 10,712 + 4,226	36,768	\$18.58

Table C.1: Model-call and cost accounting.

API and malformed-JSON failures up to five times with exponential backoff.

Table C.1 covers the main DeepSeek run; it excludes the cross-evaluator robustness checks reported later in this appendix. The covered stages ran on August 8, 2026, required 36,768 model invocations and approximately three hours, and issued up to 40 calls concurrently. Exact duplicate texts were recorded as ties without another call; retries were additional attempts after an API or parsing failure; and cached repetitions count as invocations but were charged only once. Provider-reported stage costs total \$18.58.

Tournament scores were determined through pairwise comparison. For a pool of N concepts, the software evaluated every ordered pair, producing $N(N - 1)$ judgments. Thus,

each unordered pair appeared twice, with the positions of A and B reversed, a double round-robin design that controls for position bias (David 1988, Zheng et al. 2023). The evaluator compared seven features: (1) clarity of value proposition, (2) market size and demand, (3) feasibility and development plan, (4) revenue model and sustainability, (5) differentiation, (6) impact and evidence plan, and (7) risk awareness and mitigation. The prompt explicitly instructed it not to reward unsupported specificity. Exact duplicate texts were scored as ties without an API call; a comparison that still failed after retries also became a tie, awarding half a win to each plan.

These tournament scores governed search. Reported performance is a concept’s win rate against the 30-project reference set: a generated concept receives 60 judgments, two against each original, while an original receives 58 judgments against the other 29 originals. This common post-search measure makes concepts encountered in different pools and stages comparable.

C.3 Local search

Local search ran for 12 rounds fixed in advance. In each round, an advisor agent proposed exactly five small, targeted changes to the current incumbent. Five rewriter calls then produced one variant per proposal. The prompt requested approximately 500 words, but the software imposed no length bound. Texts therefore sometimes grew through cumulative revision, especially in local search, where each round revised the full incumbent. The retained records preserve every candidate in full. After the fidelity check described below, the incumbent and five evaluated variants entered a six-plan double round robin (30 judgments). The plan with the most wins became the next round’s incumbent. The incumbent occupied the first tie-breaking position, so a variant had to improve its win total strictly to replace it. The same 12-round procedure was later applied during polish.

Two features of this procedure explain why the reported trajectory need not ascend monotonically. First, a round’s winner is selected within the six-plan tournament, whereas

reported performance is measured separately against the reference set, so a variant can win its round yet score slightly lower on the common scale. Second, the evaluator is stochastic: on the 60 reference-set matches behind each plotted point, sampling alone contributes a standard error of roughly 6 percentage points. Setting the evaluator’s temperature to zero would reduce sampling noise but not the difference between in-run selection and external measurement, and presentation-order effects would remain; monotone ascent requires selection and measurement by the same deterministic function, as in an NK model. The same stochasticity in rewrites and judgments also makes LLM-based local search less local than its NK counterpart: the same incumbent can generate different neighborhoods, and an occasional variant may land far away. How this looseness changes classic local-search results is a question for future research.

C.4 An illustrative search artifact

Table C.2 compares five areas changed during the `clockchain` search. Except for the italicized statements marking material absent from the seed, the cells reproduce exact passages. Text retained unchanged is omitted; the complete descriptions and all intervening variants accompany the source code.

Aspect	Seed concept	Round-12 concept
Target users	The campaign calls out artists, small businesses, journalists, and anyone who wants to prove “I had this first.” Primary uses include establishing authorship, deterring copycats, and documenting the creation time of digital assets.	The proposed initial target is digital illustrators who primarily share their work on a single, major social media platform.
Illustrative use	<i>No specific user scenario was included in the seed concept.</i>	A digital illustrator could use the tool to generate a fingerprint of their new artwork file before posting it on the named social platform. Later, if a copy appears elsewhere, the illustrator can publicly prove they had the original file first by showing the time-stamped ledger entry matching their file’s fingerprint, without ever revealing the full file.
Testing and evidence	Planned steps include hiring a developer and a designer to build a smoother public interface; polishing usability; and layering in infrastructure for scale. The team references a January beta target and says cushion time has been built into the schedule.	To inform this tailoring, the team proposes a small-scale, planned test with a handful of digital illustrators from the target platform to gather initial workflow feedback before the beta . . . with a proposed target of 80% of the 5 test users completing the workflow unassisted.
Revenue model	<i>No revenue model was specified in the seed concept.</i>	As a hypothetical example to make this tangible, a proposed pricing tier could be “5 free verifications per month, then \$5 per month for unlimited.”
Ledger technology	The product creates a digital fingerprint of a user’s file and records proof of that fingerprint on a public, time-stamped ledger.	The team proposes to anchor fingerprints to a specific public ledger technology, such as a public timestamping service like OpenTimestamps or a specific, energy-efficient blockchain.

Table C.2: `clockchain` before and after local search.

C.5 Global search

Global search began with the 15 seeds and ran for 12 replacement generations, maintaining a population of 15. At each generation, the current population entered a 210-judgment double round robin. Each concept received two ranks. Its *tournament rank* came from tournament wins. Its *distance rank* came from semantic distance to the nearest member of the fixed, original 15-seed set.

Semantic distance was calculated at the level of six venture-concept components: problem, customer, solution, delivery model, revenue logic, and distinctiveness. The software extracted

the text under the corresponding Markdown headings, embedded each component with BAAI’s BGE-M3 model (chosen for its strong retrieval performance across text lengths and open weights, which permit local, reproducible inference; Chen et al. 2024), and computed cosine distance separately for each component. The distance between two plans was the mean of those six distances; a plan’s semantic distance was computed relative to its nearest original seed. The analysis fixed one revision of the embedding model for all comparisons.

Tournament and distance ranks were converted to percentile scores and averaged with equal weights. This average defined the selection score; concepts then received a selection rank, with rank 1 best. The ten best selection ranks formed the parent pool. Five concepts were copied unchanged as elites in selection-rank order, subject to a cap of two elite slots per family; the remaining ten slots held five mutants and five crossover children. Mutation changed one randomly selected component while asking the model to preserve the other five. For crossover, the first parent was sampled from the parent pool with probability proportional to the inverse of its selection rank; the second was the best-ranked eligible plan that added a different family. The model assigned each of the six components to one parent, with at least two components taken from each, and rewrote only enough connective material to form a coherent new concept. Candidate identifiers and component-level ancestry were retained throughout the run. Equal tournament totals were shuffled using the run’s seeded pseudorandom generator (seed 42).

C.6 An illustrative global-search artifact

Table C.3 makes the descendant with four seed ancestors in Figure 4 substantively interpretable. At the end of global search, before local polish, the concept was a privacy-first platform that recorded verifiable fingerprints of ethical-sourcing events on a public ledger while protecting sensitive operational data. It had the highest performance in the final generation (40 percent) and reached 73 percent after local polish. The table shows exact passages in which `clockchain`’s verification mechanism, `hply`’s charitable-impact problem, and `angelry`’s

Ancestral seed	Selected seed passage	Related global-search passage
clockchain	The product creates a digital fingerprint of a user’s file and records proof of that fingerprint on a public, time-stamped ledger. . . . users do not upload the actual file; instead, a compact fingerprint is used as a durable evidence marker.	The venture plans to build a tool that creates a digital fingerprint of a user’s charitable event data—such as a donation receipt or a nonprofit’s impact milestone—and records proof of that fingerprint on a public, time-stamped ledger. . . . users would not upload sensitive raw data; instead, a compact fingerprint would serve as durable evidence.
hply	The creators say many people want to help but hesitate because it’s hard to know which organizations are trustworthy, where the money goes, and whether it truly helps.	The central challenge is enabling people to credibly prove the authenticity and first possession of a digital record—such as a donation confirmation or an impact report—amid pervasive skepticism and the ease of generating false claims.
angelry	The creators frame conventional leather as dependent on animal harm and present a gap for ethically sourced, real leather that preserves the material’s qualities without sacrificing animal welfare.	The venture plans to target both brands seeking to validate their supply chains and consumers looking for trustworthy ethical goods. . . . verifying material provenance with confidence, and providing transparent proof of sourcing standards to end customers.

Table C.3: An interpretable global-search recombination.

ethical-sourcing domain reappear in the evolved text. These are visible resemblances, not clause-level attributions: crossover rewrites its parents, and not every ancestor remains identifiable.

C.7 Fidelity checks

Generation was allowed to propose new features, markets, revenue models, tests, partnerships, and targets, because such novelty is the object of the search. It was not allowed to turn these proposals into claims of existing facts. Every concept-generation prompt therefore required new claims to be framed as proposed, planned, hypothetical, or to be tested, and prohibited invented traction, customers, revenue, pilots, testimonials, partners, prototypes, credentials, dates, and quantified evidence.

Rewrites and mutations were compared with their sole parent by a separate fidelity auditor.

The auditor distinguished a new venture design—which may legitimately be new—from claims that the venture had already achieved something. A rejected rewrite or mutation was replaced in the tournament by an unchanged copy of that parent.

Crossover children received a narrower check because the new design was deliberately assembled from two sources. Inherited design elements were permitted, but a child could not inherit either parent’s accomplishments: as a newly proposed venture, it cannot have a history of its own. A deterministic rule prevented the model check from rejecting clearly prospective statements or descriptions of the proposed concept while still treating unqualified claims of existing accomplishments as failures. A rejected crossover was replaced by an unchanged copy of its first parent. In all cases, the generated text and evaluated fallback were saved, and the program stopped rather than silently accepting a candidate if the check could not be completed.

Across the main pipeline, the checks rejected 171 of 1,080 generated texts: 4 of 60 local-search rewrites, 16 of 60 mutations, 2 of 60 crossovers, and 149 of 900 polish rewrites.

C.8 Polish and final measurement

After global search, each of the 15 concepts in the final population received 12 rounds of local search. These runs constitute the *polish* stage.

For the local-only comparison, the same 12-round local-search procedure was run once from each of the 15 original seeds, without global search. This set comprises the focal **clockchain** run and 14 additional searches. The final concepts were measured against the same reference set as the global-search and polished concepts.

Final reference-set win rates were calculated from the recorded matches, counting only comparisons in which the opponent was one of the 30 originals. Component embeddings, reference-set win rates, lineage summaries, figures, and tables were produced after the searches ended and did not alter their decisions.

C.9 Cross-evaluator robustness

Because DeepSeek generated and evaluated the concepts, I re-scored the frozen outputs with two open-weight models from other developers: Gemma 3 4B and Qwen 2.5 72B. Each used the same pairwise prompt and both presentation orders. For each evaluator, the check reran the 30-original double round robin and compared each of the 44 distinct generated texts with every original, requiring 3,510 judgments. Rankings of the originals remained close to DeepSeek’s (Spearman $\rho = 0.94$ for Gemma and 0.82 for Qwen). The focal polished descendant outranked 29 originals under Gemma and 26 under Qwen, compared with 24 under DeepSeek. Polish improved 13 of 15 concepts under Gemma and 12 under Qwen, although exact generated-concept rankings were less stable. Both outside judges placed the polished slate’s median above the local-only slate’s. Together with the focal descendant’s strong placement, this strengthens the evidence that global recombination followed by polish can reach distant regions while retaining high performance across model evaluators. Exact generated-concept rankings still varied, and preferences shared across model families remain possible.

C.10 Search records and reproducibility

The system cached each successful model response together with the prompt and model settings. An interrupted search could therefore resume without repeating completed calls. The reported run also replayed the previously completed original-versus-original judgments, keeping the reference-set ranking fixed.

The external record preserves every prompt, generated and evaluated concept, fidelity decision, tournament result, and surviving lineage, along with the random seed, model settings, embedding-model revision, timing, and cost. These materials permit reconstruction of the reported searches. A fresh live replication may still produce different paths because model responses remain stochastic and provider implementations can change.

Appendix D: Prompts

The prompts below show the operative instructions and response formats. Bracketed placeholders stand for project descriptions and other inputs supplied at runtime. The complete prompts and inputs are included in the source code.

D.1 Shared prompt segments

Two segments recur in the prompts that follow. [SHARED CRITERIA] refers to the seven evaluation criteria in the first block below. [FIDELITY RULE] refers to the fidelity instruction in the second block. The software inserted the full segment at runtime; the bracketed labels are abbreviations used only in this appendix.

-
1. Clarity of value proposition
 2. Market size and demand
 3. Feasibility and development plan
 4. Revenue model and sustainability
 5. Differentiation
 6. Impact and evidence plan
 7. Risk awareness and mitigation
-

The shared fidelity rule allowed new proposals but prevented the model from presenting invented evidence as fact.

Treat the supplied project description(s) as the complete factual record. You may propose new features, markets, revenue models, tests, partnerships, targets, projections, or plans only when they are explicitly framed as proposed, hypothetical, planned, or to be tested. Do not invent or imply existing traction, customers, users, sales, revenue, growth, pilots, tests, testimonials, partners, endorsements, certifications, awards, prototypes, development status, launch status, team credentials, dates, or quantified evidence. Do not convert a recommendation or aspiration into an accomplished fact. Preserve any factual claims that are carried forward from a source without strengthening them.

D.2 Pairwise evaluator

Every local and global tournament used this prompt to compare two venture concepts. Each pair was judged twice with the presentation order reversed. The response identified a winner

and recorded the evaluator’s reasoning.

You are an expert venture capital analyst evaluating crowdfunding projects. Compare the two project descriptions below along these criteria:
[SHARED CRITERIA]

--- PROJECT A ---
[PLAN A]
--- PROJECT B ---
[PLAN B]

For each project, list its pros and cons as bullet points, grounded in the criteria above. Do not reward unsupported specificity: claims of traction, partnerships, testimonials, completed development, or quantified evidence should improve a project’s assessment only when they are genuine source facts. Then state which project wins overall and why.

Respond with a JSON object: {"analysis": "your analysis here", "winner": "A" or "B"}
Output only valid JSON, nothing else.

D.3 Local-search operators

At the start of each local-search round, the advisor prompt generated five small changes to the incumbent. It proposed directions for revision rather than rewriting the concept itself.

You are an expert startup advisor. Below is a crowdfunding project description:
[CURRENT PLAN]

Evaluate this project along these dimensions:
[SHARED CRITERIA]

Suggest exactly 5 concrete, actionable improvement ideas that would make this project more compelling to potential backers. Each idea should propose a small, targeted change (not a complete rewrite). Suggestions must obey these rules:
[FIDELITY RULE]

Respond with a JSON object: {"ideas": ["idea 1", "idea 2", ...]}
Output only valid JSON, nothing else.

The system then called the rewriter once for each idea. This prompt asked for a complete revised concept while preserving the source’s structure, tone, and factual record.

You are an expert startup copywriter. Below is a crowdfunding project description followed by one specific improvement idea.
[CURRENT PLAN]
[IMPROVEMENT IDEA]

Rewrite the project description incorporating the improvement idea. Keep the overall structure and tone, but integrate the improvement naturally. The updated description

should be approximately 500 words and must obey [FIDELITY RULE].

Respond with a JSON object: {"plan": "updated project description", "evidence_notes": "note identifying proposed or supported claims"}

Output only valid JSON, nothing else.

D.4 Mutation

Global search used this prompt to change one component of a selected parent while preserving the other five. The algorithm selected the component at random before the call. [COMPONENT] is its name (for example, `revenue_logic`); [COMPONENT GUIDANCE] is a one-line definition of what it covers (in this example, who pays, what they pay for, and how the venture can be sustained).

You are an expert startup advisor. Here is a crowdfunding project description:
[PARENT PLAN]

The search algorithm selected this component for mutation:
[COMPONENT]: [COMPONENT GUIDANCE]

Make one substantive change to that component only. Preserve the other five components and all source-grounded facts as closely as possible. Keep the existing Markdown headings so that the components remain machine-readable. Rewrite the description incorporating only that improvement. The result should be approximately 500 words and must obey [FIDELITY RULE].

Respond with a JSON object: {"component": "[COMPONENT]", "plan": "rewritten description"}

Output only valid JSON, nothing else.

D.5 Crossover

The crossover prompt combined two parents selected by the search algorithm. It required an explicit source for each component and at least two components from each parent, producing a new venture design rather than a lightly edited copy of either parent.

You are an expert startup strategist. Here are two crowdfunding project descriptions:
[PARENT A]
[PARENT B]

Create a NEW crowdfunding project by RECOMBINING the two parents at the level of their six components: problem, customer, solution, delivery_model, revenue_logic, and distinctiveness.

For each component, choose exactly ONE parent as its source. Take at least two components from each parent, using [SHARED CRITERIA] to judge which supplies the stronger version. Carry each chosen component’s substantive design over faithfully, but do not carry claims of existing traction, customers, prototypes, development status, testimonials, partners, awards, or other accomplishments. Re-express evidence-dependent ideas as proposed plans or targets. Rewrite only enough connective tissue to make one coherent new venture of approximately 500 words. Keep the Markdown headings; the child must not restate either parent. Obey [FIDELITY RULE].

Respond with a JSON object: {"component_map": {"problem": "A" or "B", ..., "distinctiveness": "A" or "B"}, "plan": "the new project description"}

Output only valid JSON, nothing else.

D.6 Fidelity auditor

Every rewritten, mutated, or crossed concept passed through a separate fidelity-check prompt before entering a tournament. The auditor returned a pass/fail decision and identified unsupported claims. The central distinction is between a proposal and a claimed accomplishment. For example, “We plan to test the product with 50 users” is allowed; “We tested the product with 50 users” fails unless the source plan reports that test.

The generic prompt below audited rewrites and mutations. [PARENT TEXT OR TEXTS] was the single parent plan authorized as evidence, and [CANDIDATE] was the newly generated plan being checked.

You are auditing factual fidelity in a generated crowdfunding venture concept.

--- AUTHORIZED SOURCE MATERIAL ---
[PARENT TEXT OR TEXTS]
--- GENERATED CANDIDATE ---
[CANDIDATE]

Apply [FIDELITY RULE]. Pass the candidate only if every statement about existing real-world evidence or status is supported by the authorized source. New ideas are allowed when clearly framed as proposals, plans, targets, hypotheses, or future tests. Distinguish the proposed concept---its features, target users, pricing, and positioning---from claims of accomplished evidence. Do not reject a new or detailed proposal merely because it is absent from the source.

Respond with a JSON object: {"passed": true or false, "unsupported_claims": ["claim", ...], "analysis": "brief explanation"}

Output only valid JSON, nothing else.

Crossover used a separate prompt because the child was meant to combine design com-

ponents not stated together in either parent. Its new features, target users, pricing, and positioning were therefore allowed even when absent from the sources. The child nevertheless began with no history of its own, so claims that it already had customers, revenue, pilots, partners, prototypes, testimonials, or other accomplishments were unsupported even when true of a parent.

You are auditing a proposed crowdfunding venture concept deliberately assembled by RECOMBINING components from two parent concepts.

--- AUTHORIZED SOURCE MATERIAL ---
[PARENT A]
[PARENT B]
--- RECOMBINED CANDIDATE ---
[CANDIDATE]

The recombined concept itself is allowed to be new. Do not reject its features, target users, pricing, positioning, or other design choices merely because neither parent states that exact combination.

Identify only claims that the child venture has already accomplished something in the real world. The child is a new proposal and has no history of its own; it may not inherit claims of customers, revenue, pilots, partners, prototypes, testimonials, awards, credentials, or other accomplishments from either parent. Statements framed as proposals, plans, targets, hypotheses, or future tests are allowed. Apply [FIDELITY RULE].

Respond with a JSON object: {"passed": true or false, "unsupported_claims": ["claim", ...], "analysis": "brief explanation"}
Output only valid JSON, nothing else.

After the crossover audit, a deterministic rule retained only flagged claims whose surrounding sentence asserted an existing accomplishment without prospective wording. A candidate passed if no flagged claim met that test. This prevented clearly proposed design content from being rejected while preserving the prohibition on invented evidence.