FounderPulse-AI: Design and Pilot Evaluation of Bounded Intelligence for Startup Founders Under Entrepreneurial Strain

Stavros Sindakis[1,2], Christina Kovatsi[3]

[1] Akio Morita School of Business, Anaheim University, USA; ssindakis@anaheim.edu

[2] School of Social Sciences, Hellenic Open University, Greece; sindakis.stavros@ac.eap.gr

[3] Areimanio Tutoring School and Independent Researcher, 126 Lavriou Ave., Glika Nera 153 54, Greece; kovatsi.christina@gmail.com

*Corresponding Author: ssindakis@anaheim.edu / sindakis.stavros@ac.eap.gr

DOI: https://doi.org/10.30212/JITI.202604.013

Submitted: Jun 21, 2026 Accepted: Aug 04, 2026

ABSTRACT

Intelligent decision support systems increasingly assist users under uncertainty, but limited research explains how they can support startup founders under entrepreneurial strain without crossing into coaching, clinical, or advisory authority. This study designed and pilot-evaluated FounderPulse-AI, a retrieval-grounded, prompt-governed, nonclinical decision support artifact for Greek startup founders. Using design science research, the study combined a qualitative needs analysis with 20 founders, co-founders, and ecosystem actors and a four-week comparison involving 100 founders and co-founders assigned to FounderPulse-AI or a structured reflection journal. FounderPulse-AI produced stronger gains in reflective clarity, perceived decision quality, perceived usefulness, trust in AI, usability, boundary clarity, escalation appropriateness, and intention to continue use, while privacy concerns decreased more in the journal condition. The main contribution is a bounded AI founder support architecture and a Responsible AI Founder Support Model that connects intelligent interaction, decision recovery, privacy sensitivity, and human escalation.

Keywords: Responsible AI, Intelligent decision support, Startup founders, Design science, Entrepreneurial strain

1. Introduction

Intelligent decision support systems are increasingly used to help users interpret complex information, compare alternatives, and make decisions under uncertainty. Recent research emphasizes that such systems are most valuable when they combine decision support with transparency, explainability, user trust, and attention to the decision process rather than producing only automated outputs [1]. These requirements are especially relevant for startup founders, who frequently make consequential decisions under resource scarcity, time pressure, incomplete information, and personal responsibility for venture outcomes. Greece provides an appropriate setting for this inquiry because recent ecosystem reports indicate more than 555 million euros invested in more than 90 Greek startups in 2024, while OECD evidence shows that nearly 60% of Greeks report fear of failure as a barrier to starting a business [2], [3]. This combination of ecosystem expansion and persistent perceived entrepreneurial risk makes Greece a useful context for examining intelligent support during moments of founder decision strain.

The practical problem addressed in this study concerns the timing and structure of founder support. A recent systematic literature review of early-stage startups shows that founders face critical decisions across areas such as people, product definition, market selection, partnerships, and other domains that shape early venture development [4]. These decisions are not only technical or informational. They are also shaped by the founder's capacity to interpret pressure, compare options, manage uncertainty, and prioritize action. Entrepreneurship research further argues that stress, coping, and resilience should be studied together, as entrepreneurs build resilience through responses to adversity rather than through isolated traits or single stressors [5]. Recent evidence on entrepreneurs' recovery experiences also shows that recovery quality is associated with well-being and burnout among small business owners, suggesting that founders need mechanisms to help them restore cognitive and emotional resources during demanding work periods [6].

Existing founder-support mechanisms can help entrepreneurs learn, reflect, and adjust, but their availability is often episodic. Founders may receive support from mentors, accelerators, incubators, investors, coaches, advisers, and peers. These support actors can improve judgment and foster learning from experience, yet they are not always available when founders face immediate decision pressure. The resulting gap is not simply a lack of advice. It is the absence of structured, context-sensitive, and timely support at the moment when pressure can interfere with reflective clarity and action prioritization. Intelligent decision-support systems provide a relevant technological foundation for addressing this gap because recent research shows that AI-based decision-support systems are increasingly designed to improve transparency, interpretability, trust, and user confidence in complex decision processes [7].

The use of intelligent systems in founder-support contexts requires careful boundaries. FounderPulse-AI is not designed to diagnose, treat, monitor, or predict mental-health conditions. It is designed as a non-clinical decision-support artifact that helps founders' structure current pressure, decompose decision problems, distinguish facts from assumptions, compare immediate options, prioritize short-term actions, and identify when human support may be needed. This boundary is necessary because founder strain is well-being-adjacent even when the system is not clinical. The NIST AI Risk Management Framework states that AI risk management should address risks to individuals, organizations, and society [8]. WHO guidance on artificial intelligence for health also emphasizes that ethics and human rights should be placed at the center of AI design, deployment, and use in sensitive health-related contexts [9]. Although this study is not a health intervention, these responsible AI principles are relevant because the artifact interacts with founders during moments of pressure and vulnerability.

This study is grounded in design science research because its primary contribution is the development and evaluation of an artifact. Design science research seeks to extend human and organizational capabilities through the creation of innovative artifacts and generates knowledge by building and applying designed solutions [10]. The design science research methodology is also appropriate because it provides a process for problem identification, solution objectives, design and development, demonstration, evaluation, and communication [11]. These methodological foundations fit this study because the research not only examines founder strain. It also designs and evaluates an intelligent system that translates founder strain into structured reflection and short-term managerial action.

The research gap addressed in this study is an integration and operationalization gap within intelligent technology and innovation research. Prior studies explain early startup decisions, entrepreneurial stress, recovery, resilience, and AI based decision support, but they do not show how these streams can be translated into a founder specific intelligent support artifact that remains useful, transparent, privacy sensitive, and bounded during moments of entrepreneurial strain [4], [5], [6], [7], [8], [9]. This gap is important because founder support requires more than generic advice or emotional encouragement. It requires a structured workflow that helps founders clarify the immediate decision, separate facts from assumptions, compare feasible options, prioritize short-term action, and recognize when human support is needed. The focus of this research is therefore the design and pilot evaluation of FounderPulse-AI as a nonclinical decision-support system that converts founder strain into reflective clarity and actionable decision recovery, while preserving founder agency, human escalation, and responsible AI boundaries.

Accordingly, the study examines how a human-centered intelligent decision-support system can be designed to help startup founders regain reflective clarity and perceived decision quality amid uncertainty. It also explores which design principles are required for founder support to remain useful, bounded, transparent, privacy sensitive, and safe. Finally, it examines whether FounderPulse-AI produces stronger pilot evidence than a structured digital reflection journal across decision recovery, venture resilience, usability, responsible AI evaluation conditions, and intention to continue use.

The study followed a design science and pilot field evaluation approach. It first conducted a needs analysis with Greek startup founders and ecosystem support actors to identify decision-making pressures, expectations for support, privacy concerns, and acceptable AI boundaries. It then developed FounderPulse-AI as a rule-constrained, retrieval-augmented, prompt-governed conversational decision-support system. The artifact included modules for founder strain input, decision decomposition, reflective reframing, action prioritization, and safety-aware escalation. The artifact was evaluated through a four-week pilot comparing FounderPulse-AI with a structured digital reflection journal.

This study contributes to intelligent technology and innovation research by developing and pilot-evaluating FounderPulse-AI as a bounded intelligent decision-support artifact for startup founders operating under entrepreneurial strain. At the artifact level, the study specifies how to integrate retrieval grounding, prompt governance, structured decision decomposition, reflective reframing, action prioritization, privacy-sensitive input design, and safety-aware escalation into a non-clinical founder support architecture. At the evaluation level, it compares adaptive AI-supported reflection with a structured digital reflection journal, showing where intelligent interaction adds value beyond nonadaptive reflection and where the journal condition remains preferable, especially in relation to privacy concerns. At the model level, it advances the Responsible AI Founder Support Model, which explains how founder strain can be translated into reflective clarity, perceived decision quality, cautious venture resilience, and continued use conditions without positioning AI as a substitute for mentors, advisers, clinicians, or founder judgment.

2. Literature Review

2.1 AI-Based Decision Support Systems

AI-based decision support systems are increasingly developed to assist users who must make decisions under uncertainty, complexity, and information overload. Recent decision support research emphasizes that AI systems must not only produce accurate outputs but also support transparency, explainability, user trust, and user acceptance [1]. This point is directly relevant to FounderPulse-AI because founders are unlikely to rely on an intelligent system during decision pressure unless they understand the system's role, limits, and reasoning cues. In this study, decision support is therefore not defined as automated advice delivery. It is defined as structured assistance that helps founders organize decision inputs, clarify assumptions, compare alternatives, and identify immediate actions while preserving human agency.

The AI decision support literature also suggests that trust cannot be treated as a secondary adoption outcome. Kővári argues that decision support systems that use AI must balance accuracy with visibility and explainability, as users need confidence in both the output and the process by which it is generated [1]. This insight matters for founder support because entrepreneurial decisions often involve incomplete evidence, personal risk, and irreversible resource commitments. FounderPulse-AI, therefore, cannot be designed as a black box that produces confident recommendations. It must instead generate reflective prompts, disclose its non-clinical and advisory limits, and make its decision decomposition logic visible to the user.

2.2 Retrieval Grounded Conversational AI

Conversational AI provides the interaction layer for FounderPulse-AI, but a generic chatbot architecture would be insufficient for this study. Retrieval-augmented generation offers a more suitable foundation by combining language generation with external knowledge retrieval, improving grounding, domain adaptation, and knowledge updating for knowledge-intensive tasks [12]. Gao and colleagues explain that retrieval-augmented generation integrates retrieval, generation, and augmentation processes, enabling generative outputs to be linked to external knowledge resources rather than relying solely on a model's internal parameters [12]. This makes retrieval-augmented generation relevant for founder support, as the artifact should rely on pre-approved decision prompts, founder support resources, escalation guidance, and responsible AI rules.

This design logic does not remove the need for human oversight. Retrieval can improve grounding, but it cannot guarantee that every output will be appropriate for a founder's context. The value of retrieval augmentation in FounderPulse-AI is therefore bounded. It supports response control, reduces free-form speculation, and allows the system to draw from predefined support materials. It does not authorize the system to diagnose distress, prescribe psychological interventions, or replace mentors and advisers. For this reason, this study treats conversational AI as an interaction mechanism and retrieval augmentation as a grounding mechanism, while reserving final judgment for the founder and escalation decisions for human support channels.

2.3 Entrepreneurial Strain as a Design Requirement

Entrepreneurial strain is reviewed in this study as a design requirement rather than as the main theoretical contribution. Early-stage founders face critical decisions across people, product definition, market segment selection, partnerships, and other domains that shape venture development [4]. These decisions require more than just factual information, as founders must also interpret uncertainty, prioritize scarce resources, and make trade-offs under time pressure. This decision context explains why an intelligent founder support system must begin with decision decomposition rather than general encouragement.

Entrepreneurship research also shows that stress, coping, and resilience should be examined together because entrepreneurial resilience develops through responses to adversity [5]. This perspective supports the inclusion of reflective reframing in FounderPulse-AI. Reframing is not used as therapy. It is used as a structured decision-support function that helps founders distinguish controllable from partially controllable or uncontrollable pressures. This distinction can help the founder allocate attention to decisions that can be acted on rather than to undifferentiated strain.

Recent research on entrepreneurs' daily recovery experiences further strengthens this design requirement. Le Moal and colleagues show that daily recovery experiences among small business owners are associated with higher well-being and lower burnout, and that recovery dynamics among business owners differ from those of employees [6]. This finding is relevant because founders often have blurred work-life boundaries and a high level of personal involvement in the venture. FounderPulse-AI therefore focuses on decision recovery and reflective clarity rather than broad well-being improvement. The system is designed to help founders interrupt decision overload, restore structure, and convert pressure into a short set of next actions.

2.4 Responsible AI Boundaries for Well-being Adjacent Founder Support

The responsible AI literature is central to this paper because FounderPulse-AI operates in a well-being-adjacent context. NIST states that the AI Risk Management Framework was developed to help manage risks to individuals, organizations, and society associated with artificial intelligence, and to support the integration of trustworthiness considerations into AI design, development, use, and evaluation [8]. This framework supports this study's emphasis on boundedness, transparency, privacy sensitivity, and escalation. These safeguards are not peripheral design features. They define whether an intelligent founder support system can be considered responsible.

WHO guidance on artificial intelligence for health is also relevant because it highlights the ethical and governance challenges that arise when AI systems are used in sensitive contexts related to human well-being [9]. Although FounderPulse-AI is not a health system, its interaction with founders during moments of strain requires caution. WHO's guidance on large multimodal models further explains that generative AI systems can accept varied inputs and produce diverse outputs, underscoring the need for governance and boundary-setting when used in sensitive domains [13]. This supports the decision to position FounderPulse-AI as a non-clinical system that structures reflection and action rather than one that evaluates or treats mental-health conditions.

Recent reviews of large language models in mental-health applications reinforce this boundary. Guo and colleagues identify potential uses of large language models in mental health contexts, but they also raise concerns about reliability, accuracy, interpretability, data quality, and ethical risks [14]. These concerns support a conservative design stance. FounderPulse-AI should not imitate therapeutic dialogue, infer diagnoses, monitor symptoms, or encourage emotional dependency. Its appropriate role is to help founders formulate a decision problem, organize available information, generate bounded alternatives, and seek human support when risk signals appear.

2.5 Design Science Research and Artifact Evaluation

Design science research provides the methodological bridge between the technological problem and the artifact contribution. Hevner and colleagues argue that design science research develops knowledge through the building and application of artifacts, and that such artifacts should address relevant problems while being evaluated for utility [10]. This foundation fits this study because FounderPulse-AI is not only a conceptual model. It is a designed system that must be specified, demonstrated, and evaluated in a founder support context.

The design science research methodology further supports this approach by providing a process for problem identification, solution objectives, design and development, demonstration, evaluation, and communication [11]. This process aligns with the planned structure of this study. The needs analysis identifies founder decision pressure and acceptable AI boundaries. The artifact development phase translates these requirements into system modules and design principles. The pilot evaluation examines whether the system improves reflective clarity, perceived decision quality, venture resilience, and usability compared with a structured digital reflection journal.

Design science also helps differentiate this study from research on entrepreneurship psychology. The goal is not to develop another explanatory model of founder stress. The goal is to create and evaluate an intelligent artifact that operationalizes decision decomposition, reflective reframing, action prioritization, and escalation in a responsible AI architecture. This distinction locates the contribution in intelligent technology and innovation rather than in founder well-being alone.

2.6 Research Gap and Study Positioning

The literature review identifies an integration gap rather than an isolated gap within any single stream. AI decision-support research establishes the importance of transparency, explainability, usefulness, user confidence, and decision-process quality [1]. Retrieval-augmented generation research explains how generative outputs can be grounded in external knowledge resources and constrained by approved support materials [12]. Entrepreneurship research shows that early-stage founders face consequential decisions under uncertainty, and that strain, coping, recovery, and resilience are connected to how entrepreneurs respond to adversity [4], [5], [6]. Responsible AI research establishes the need for privacy sensitivity, human oversight, transparency, accountability, and appropriate system boundaries, especially when AI systems operate in contexts related to well-being [8], [9], [13], [14], [15]. Design science research provides the methodological logic for converting these requirements into an artifact that can be demonstrated and evaluated [10], [11].

The unresolved problem is that these streams have not been sufficiently integrated into a founder-specific intelligent support artifact that is both useful and deliberately bounded. Existing research explains decision support, retrieval grounding, entrepreneurial strain, responsible AI, and artifact evaluation, but it does not show how these elements can be operationalized together in a non-clinical AI system for founders under decision pressure. FounderPulse-AI addresses this gap by translating founder strain into structured decision framing, reflective clarity, action prioritization, and appropriate human escalation within a responsible AI architecture. Table 1 summarizes how each literature stream informs the artifact design and evaluation logic.

Table 1. Literature streams and research positioning

Literature stream	Core insight	Limitation in prior work	Relevance to FounderPulse-AI	Key references
AI-based decision support systems	AI-based decision support systems can improve decision-making when accuracy, transparency, explainability, usefulness, and user trust are addressed together.	Prior work often examines AI decision support in general organizational, technical, or sectoral settings rather than in founder-specific decision-making contexts.	FounderPulse-AI applies AI-based decision support to startup founders by supporting decision decomposition, reflective clarity, option comparison, and action prioritization while preserving founder agency.	[1], [7]
Retrieval-augmented conversational AI	Retrieval-augmented generation can ground language model outputs in external knowledge sources and reduce reliance on unsupported	Prior work often focuses on technical retrieval-augmented generation performance, while less attention is given to bounded, user-facing,	FounderPulse-AI uses retrieval grounding to connect responses to approved decision prompts, startup support resources, responsible AI rules,	[12]

Literature stream	Core insight	Limitation in prior work	Relevance to FounderPulse-AI	Key references
	generation.	safety-aware applications in entrepreneurial support.	privacy notices, and escalation guidance.	
Entrepreneurial decision pressure	Early-stage founders face critical decisions across product, people, market, partnerships, financing, and strategic positioning.	Prior entrepreneurship work identifies decision domains, but it does not explain how intelligent systems can provide timely and structured decision support during founder strain.	FounderPulse-AI addresses this gap by converting founder pressure into a structured decision frame that separates facts, assumptions, constraints, missing information, options, and tradeoffs.	[4]
Entrepreneurial strain, coping, and recovery	Entrepreneurial resilience develops through responses to adversity, and recovery experiences are associated with well-being and burnout among entrepreneurs and small business owners.	Prior work explains founder strain, coping, resilience, and recovery, but it does not operationalize these insights as responsible intelligent system design requirements.	FounderPulse-AI translates this literature into non-clinical design functions, including reflective reframing, attention to controllable pressures, short-term action prioritization, and escalation to human support.	[5], [6]
Responsible AI and trustworthy system design	Responsible AI requires attention to risk management, human oversight, privacy, transparency, accountability, safety, and appropriate system boundaries.	Prior responsible AI guidance is often general, while founder support requires concrete safeguards for a well-being-adjacent but non-clinical use context.	FounderPulse-AI incorporates boundedness, transparency, privacy sensitivity, human escalation, prompt governance, and non-clinical positioning as core design principles rather than optional	[8], [9], [13], [15]

Literature stream	Core insight	Limitation in prior work	Relevance to FounderPulse-AI	Key references
			safeguards.	
LLM risks in well-being adjacent contexts	LLM applications near mental health and well-being contexts raise concerns about reliability, accuracy, interpretability, privacy, ethical risk, and overreliance.	Prior work warns of clinical and well-being-related risks, but it offers limited guidance for entrepreneurial decision-support systems that deliberately avoid clinical functions.	FounderPulse-AI responds by avoiding diagnosis, therapy, symptom monitoring, and emotional dependency while supporting structured reflection and timely redirection to human support.	[14]
Design science research	Design science research produces knowledge through the creation, demonstration, and evaluation of artifacts that address relevant organizational and technological problems.	Prior entrepreneurship and AI support research often remains conceptual or explanatory rather than artifact-based and empirically evaluated.	FounderPulse-AI is positioned as a designed artifact developed from user needs, translated into system modules, and evaluated through a pilot comparison and structured reflection.	[10], [11], [16], [17]

3. Theoretical and Design Framework

The theoretical and design framework translates the literature synthesis into an artifact logic. FounderPulse-AI is not positioned as an autonomous adviser, a therapeutic agent, or a substitute for founder judgment. It is positioned as a bounded decision support system that helps founders move from unstructured entrepreneurial strain toward reflective clarity, perceived decision quality, cautious venture resilience, and informed continued use. This framing is consistent with AI decision support research, which emphasizes that intelligent systems should support decision process quality, transparency, explainability, and user confidence [1]. It is also consistent with responsible AI guidance, which requires risks to individuals, organizations, and society to be governed, mapped, measured, and managed throughout system design and use [8].

The framework begins with founder strain as the triggering condition. Founder strain refers to the perceived pressure that emerges when founders face urgent decisions under uncertainty, limited resources, and personal responsibility for venture outcomes. Early-stage startup research shows that founders face critical decisions across product, people, market, partnerships, and other domains of venture development [4]. Entrepreneurship research further argues that stress, coping, and resilience should be examined together because resilience develops through responses to adversity [5]. In this study, founder strain is not treated as a clinical condition. It is treated as a decision context that can reduce reflective clarity and make action prioritization more difficult.

Reflective clarity is the first central mechanism in the framework. It refers to the founder's perceived ability to separate facts from assumptions, distinguish controllable from uncontrollable pressures, and formulate the decision problem in actionable terms. This mechanism is important because FounderPulse-AI is designed to support reflection before recommendation. AI decision-support research indicates that users need visibility into the logic and limits of AI-supported decisions to develop confidence in system outputs [1]. For this reason, the system should not provide immediate answers as if they were superior to the founder's judgment. It should guide the founder through a structured process that makes the decision situation more visible and manageable.

Perceived decision quality is the second central mechanism. It refers to the founder's judgment that a decision has been clarified, compared, and prioritized in a way that supports action. The framework assumes that FounderPulse-AI can improve perceived decision quality by guiding founders through decision decomposition, option comparison, and short-term action planning. This assumption is consistent with design science research, which treats artifacts as designed solutions that should improve human or organizational capabilities in a relevant problem context [10]. It is also consistent with the design science research methodology, which requires the artifact to be demonstrated and evaluated rather than only described conceptually [11].

Venture resilience is the framework's entrepreneurial outcome. It refers to the founder's perceived capacity to maintain adaptive action, continue learning, and respond constructively to venture pressure. Recent evidence on entrepreneurs' recovery experiences shows that recovery is associated with well-being and burnout outcomes among small business owners [6]. This framework does not claim that FounderPulse-AI improves mental health. It proposes that structured decision support may help founders regain reflective clarity and maintain adaptive action during pressure-intensive periods. This makes venture resilience an appropriate non-clinical outcome for evaluating the artifact.

The framework also includes perceived usefulness, perceived ease of use, trust in AI, and privacy concern as adoption and evaluation conditions. Technology acceptance research identifies perceived usefulness and perceived ease of use as central determinants of user acceptance [18]. Research on trust in AI further shows that trust influences acceptance through perceived usefulness and user attitudes toward AI technologies [19]. These constructs are necessary because an intelligent founder support system can only have practical value if founders consider it useful, understandable, trustworthy, and acceptable to use during real decision pressure. Privacy concerns are also necessary because FounderPulse-AI may process sensitive descriptions of venture uncertainty, personal pressure, or strategic problems. Responsible AI guidance, therefore, supports privacy sensitivity, human oversight, transparency, accountability, and risk control as design conditions rather than optional features [8], [15].

FounderPulse-AI is grounded in retrieval-augmented interaction because unrestricted conversational generation would be unsuitable for this context. Retrieval-augmented generation combines retrieval mechanisms with generative language models, enabling outputs to be linked to external knowledge resources and better supporting knowledge-intensive tasks [12]. In this study, retrieval grounding is used to connect founder support outputs to approved decision prompts, responsible AI rules, escalation guidance, and startup support resources. This does not eliminate all risk. It reduces the likelihood of unsupported advice and strengthens the traceability of system responses. The framework, therefore, treats retrieval grounding as a technical safeguard that supports transparency, boundedness, and response control.

The framework translates these theoretical foundations into six design principles. The first principle is boundedness. FounderPulse-AI must define what it can and cannot do. It should support decision reflection and action prioritization, but it should not diagnose, treat, monitor, or predict mental health conditions. The second principle is transparency. The system must explain that its outputs are reflective prompts and decision-support suggestions, rather than professional advice. The third principle is privacy sensitivity. The system must minimize unnecessary data collection and avoid asking for sensitive information unless it is necessary for the support task. The fourth principle is actionability. The system must convert reflections into short, concrete, and time-bound actions. The fifth principle is human escalation. The system must redirect founders toward mentors, trusted peers, advisers, or professional support when the user's language suggests that system-based support is insufficient. The sixth principle is non-clinical positioning. The system must avoid therapeutic imitation, diagnostic labeling, and emotional dependency.

The Responsible AI Founder Support Model (Figure 1) therefore specifies both a causal sequence and a design constraint. Founder strain initiates the support process, while reflective clarity and perceived decision quality represent the primary decision recovery mechanisms. Venture resilience is treated as a cautious entrepreneurial outcome rather than a clinical or psychological claim. Continued use is shaped by perceived usefulness, perceived ease of use, trust in AI, privacy concern, and usability. FounderPulse-AI supports this sequence through structured input, decision decomposition, reflective reframing, action prioritization, and escalation, while preserving the non-clinical boundary of the system.

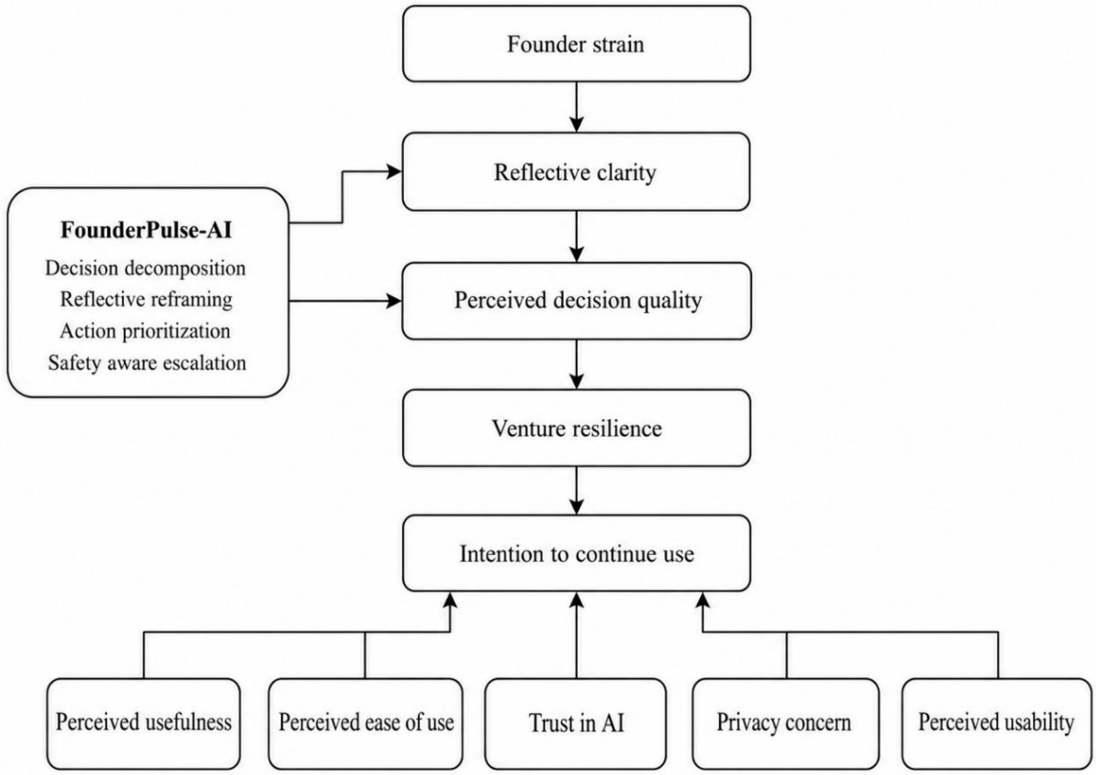

Figure 1. Responsible AI founder support model

4. FounderPulse-AI System Architecture

FounderPulse-AI operationalizes the framework through a human-centered, intelligent decision-support architecture for startup founders experiencing decision strain. The system is not designed as a clinical, therapeutic, diagnostic, monitoring, legal, financial, or investment advisory tool. Its purpose is to help founders clarify a decision problem, separate facts from assumptions, identify controllable and uncontrollable pressures, compare short-term options, prioritize action, and recognize when human support is needed. This boundary follows responsible AI guidance on risk governance, human oversight, privacy, transparency, and accountability [8], [15].

The architecture is organized around a rule-constrained, retrieval-augmented, and prompt-governed conversational workflow. The prototype used a conversational AI interface connected to a curated support corpus containing founder decision prompts, reflective questioning scripts, action prioritization templates, privacy notices, responsible AI rules, and escalation guidance. Retrieval selected relevant support materials before response generation, while prompt governance constrained response format, prohibited clinical claims, preserved founder agency, and triggered escalation messages when out-of-scope or high-risk language appeared. The system produced decision summaries, clarified assumptions, provided action options, issued boundary reminders, and provided feedback prompts. Figure 2 summarizes the architecture by showing the central workflow, retrieval grounding, prompt governance, safety-aware escalation, and system constraints.

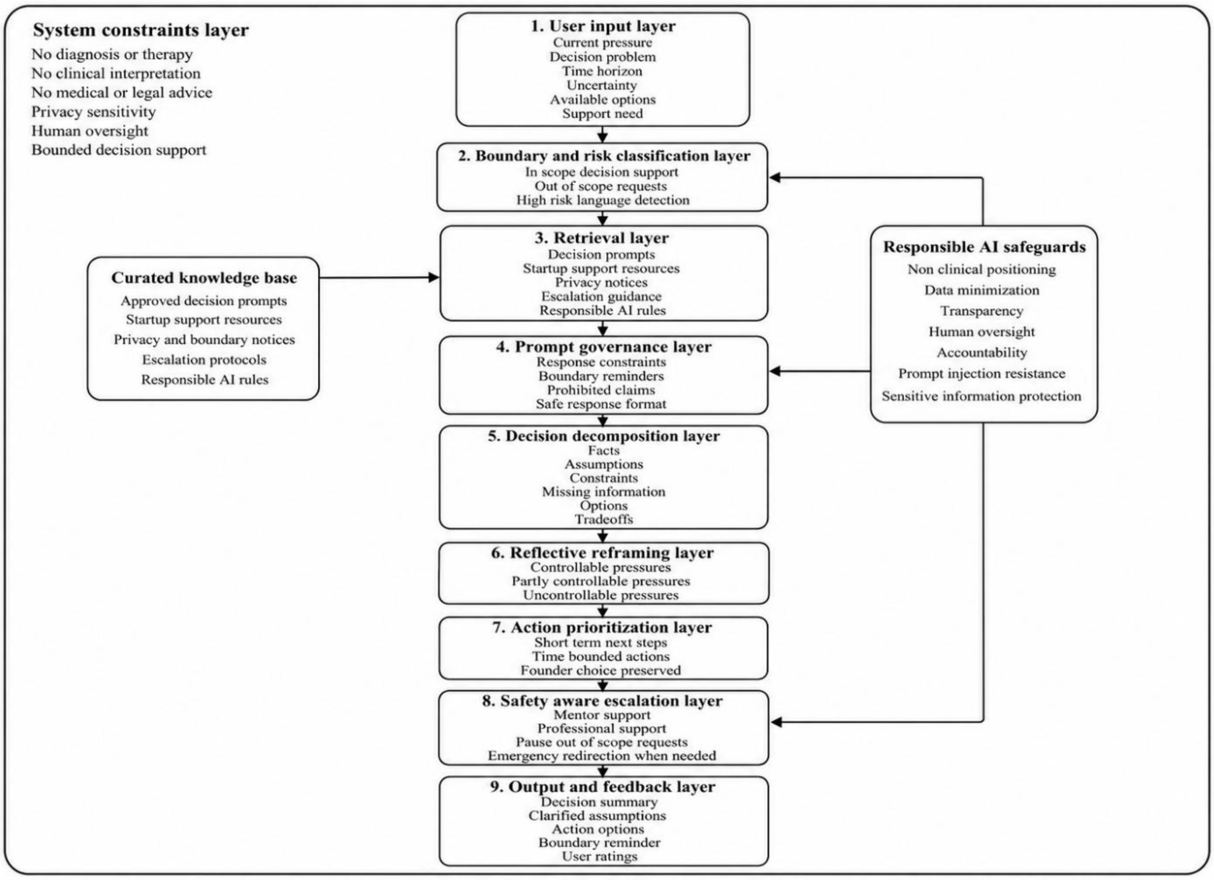

Figure 2. FounderPulse-AI system architecture

The first architectural layer is the user input layer. The founder enters a short description of the current pressure, the decision that needs attention, the time horizon, the main uncertainty, and the type of support requested. The system does not ask for medical history, diagnosis, private family details, or unnecessary personal data. This limitation reflects the principle of data minimization, which requires personal data to be adequate, relevant, and limited to what is necessary for the stated purpose [20]. The input layer, therefore, captures only information required for decision support and avoids collecting sensitive information that is not needed for reflection or action prioritization.

The second architectural layer is the boundary and risk classification layer. This layer checks whether the user request fits the system's non-clinical scope. Requests about venture decisions, prioritization, team conflict, investor communication, customer pressure, resource constraints, and time management remain within scope. Requests that appear to seek diagnosis, therapy, crisis intervention, medical interpretation, or emotional dependency are treated as out of scope. When out-of-scope language is encountered, the system does not continue the conversation as if it were a mental health assistant. It provides a boundary message and redirects the founder toward appropriate human support. This logic follows the NIST generative AI profile, which treats generative AI risk management as a companion to the broader AI Risk Management Framework and highlights the need to address risks specific to generative systems [21].

The third architectural layer is the retrieval layer. This layer retrieves relevant support materials from a curated knowledge base. The knowledge base includes founder decision prompts, reflective questioning scripts, action prioritization templates, non-clinical strain language, mentor escalation guidance, privacy statements, and system boundary rules. Retrieval-augmented generation research shows that retrieval mechanisms can support knowledge-intensive tasks by connecting model outputs to external databases and domain-specific information [12]. In this system, retrieved content functions as a control mechanism. It limits the response space and supports traceability between user needs, retrieved support logic, and generated outputs.

The fourth architectural layer is the prompt governance layer. This layer converts the user input and retrieved materials into a controlled prompt sequence. The prompt sequence instructs the model to avoid diagnosis, avoid therapeutic claims, avoid unsupported certainty, avoid medical advice, preserve founder agency, and produce short decision-support outputs. This layer is necessary because large language model applications face security and safety risks such as prompt injection, which can manipulate model behavior and bypass safety instructions [22]. FounderPulse-AI reduces this risk by separating user input from system instructions, constraining the allowed response format, and preventing user text from overriding boundary rules.

The fifth architectural layer is the decision decomposition layer. This layer transforms the founder's initial description into a structured decision frame. It separates the problem statement, known facts, assumptions, missing information, constraints, options, tradeoffs, and immediate deadlines. This layer operationalizes the system's decision-support function. It does not tell the founder what to do. It helps the founder see the decision more clearly. This design follows AI decision-support research, which emphasizes that decision systems should support transparency, explainability, and user confidence rather than merely producing outputs [1].

The sixth architectural layer is the reflective reframing layer. This layer helps the founder distinguish controllable, partially controllable, and uncontrollable pressures. It then prompts the founder to redirect attention toward actions that can be taken within the next short decision window. This function draws on entrepreneurship research showing that stress, coping, and resilience should be understood together because resilience develops through responses to adversity [5]. In the system architecture, reframing is not presented as a form of therapy. It is implemented as a decision-support function that helps founders convert an undifferentiated strain into a clearer action context.

The seventh architectural layer is the action prioritization layer. This layer generates a brief ranked list of actions. Each action must be specific, feasible, time-bounded, and linked to the decision frame created earlier in the workflow. The system avoids long generic advice and instead produces short next steps that the founder can review, accept, reject, or modify. This layer preserves human agency because the founder remains responsible for selecting and adapting actions. The system's role is to improve structure and clarity, not to automate entrepreneurial judgment. System outputs were treated as reflective decision support and were not presented as verified professional, legal, financial, investment, or clinical advice.

The eighth architectural layer is the safety-aware escalation layer. This layer monitors whether the conversation includes high-risk indicators such as references to self-harm, inability to function, severe distress, coercion, abuse, or requests for clinical guidance. When such language appears, the system does not attempt to continue decision support. It pauses the workflow, states its limitations, and directs the founder to human support, emergency support, a qualified professional, or trusted contacts, depending on the severity of the signal. This design is consistent with responsible AI principles that require human oversight, safety, accountability, and risk management in AI systems [8], [15].

The ninth architectural layer is the output and feedback layer. The final output comprises a concise decision summary, a list of clarified assumptions, a set of immediate action options, a reflective question for the founder, and a reminder of the system's non-clinical boundary. The founder can then rate perceived usefulness, reflective clarity, trust in AI, privacy concern, usability, and intention to continue use. These feedback items support the study's evaluation logic, as the artifact must be assessed not only by the coherence of its responses but also by the extent to which founders perceive it as useful, understandable, trustworthy, usable, and safe.

Table 2 summarizes how each FounderPulse-AI module operationalizes a responsible design principle, controls a specific risk, and produces an expected user benefit. In this sense, the architecture is not an unrestricted chatbot with appended warnings, but a layered decision-support workflow in which each module links a functional role to the responsible design logic. Privacy and security safeguards are therefore embedded within the artifact's design rather than treated as external compliance additions. The system should not store identifiable founder narratives unless storage is required by the study protocol and the participant has provided informed consent. It should also avoid collecting unnecessary personal, medical, family, financial, legal, investor-related, or confidential venture information. These safeguards respond to data minimization requirements and to generative AI risks associated with sensitive information disclosure [20], [23].

Table 2. FounderPulse-AI modules and responsible design principles

FounderPulse-AI module	Operational function	Responsible design principle	Risk controlled	Expected user benefit
User input layer	Captures the founder's current pressure, decision problem, time horizon, uncertainty, available options, and support needs.	Privacy sensitivity and data minimization	Unnecessary collection of personal, clinical, confidential, legal, investor, or strategic information	The founder can describe the decision situation without being asked to disclose irrelevant or excessive information.
Boundary and	Determines whether	Boundedness	System overreach into	The founder

FounderPulse-AI module	Operational function	Responsible design principle	Risk controlled	Expected user benefit
risk classification layer	the request is within the system's non-clinical decision support scope or outside its permitted use.	and non-clinical positioning	therapy, diagnosis, clinical interpretation, crisis response, or professional advice	receives clear guidance about what the system can and cannot support.
Retrieval layer	Retrieves approved decision prompts, startup support resources, responsible AI rules, privacy notices, and escalation guidance.	Grounding and response traceability	Unsupported generation, outdated advice, free-form speculation, and weak alignment with approved system logic	The founder receives responses that are better connected to predefined support materials and system boundaries.
Prompt governance layer	Applies system instructions that constrain tone, content, response format, prohibited claims, and escalation behavior.	Transparency, technical robustness, and safety	Prompt injection, unsafe outputs, clinical overclaiming, excessive certainty, and user instructions that attempt to bypass safeguards	The founder receives bounded and consistent responses that support reflection rather than replace judgment.
Decision decomposition layer	Separates the decision into facts, assumptions, constraints, missing information, options, tradeoffs, and immediate deadlines.	Actionability and explainability	Vague pressure, cognitive overload, unclear assumptions, and premature action under uncertainty	The founder can see the decision situation more clearly and identify what needs attention first.
Reflective reframing layer	Distinguishes controllable, partly controllable, and uncontrollable pressures.	Human agency and reflective clarity	Undifferentiated strain, misplaced attention, and excessive focus on issues outside the founder's control	The founder can redirect attention toward issues that can be acted upon within the current

FounderPulse-AI module	Operational function	Responsible design principle	Risk controlled	Expected user benefit
				decision window.
Action prioritization layer	Produces short, feasible, and time-bound next steps linked to the decomposed decision frame.	Actionability and founder agency	Generic advice, long recommendations, unrealistic plans, and loss of founder control	The founder receives practical options that can be accepted, rejected, modified, or discussed with a mentor.
Safety-aware escalation layer	Detects high-risk language, clinical requests, severe distress cues, or out-of-scope support needs and redirects the user to human or professional support.	Human oversight, safety, and accountability	Harmful dependency on the system, delayed human support, unsafe continuation of AI interaction, and inappropriate response to severe distress.	The founder is redirected when system-based decision support is insufficient or inappropriate.
Output and feedback layer	Provides a decision summary, clarified assumptions, action options, reflective questions, boundary reminders, and user ratings.	Transparency, usability, and continuous evaluation	Ambiguous outputs, weak user understanding, missing evaluation data, and poor design feedback	The founder receives a concise decision-support output, and the study captures usability, usefulness, trust in AI, privacy concerns, and intention to continue use.
System constraints layer	Maintains cross-cutting constraints across all modules, including non-clinical use, data	Responsible AI governance	Inconsistent behavior across modules and misalignment between the artifact, ethical safeguards, and	The founder experiences the system as a bounded decision support tool rather than an

FounderPulse-AI module	Operational function	Responsible design principle	Risk controlled	Expected user benefit
	minimization, no diagnosis or therapy, human oversight, and privacy sensitivity.		evaluation criteria	unrestricted chatbot.

The resulting architecture converts the Responsible AI Founder Support Model into an operational artifact by linking founder input, boundary classification, retrieval grounding, prompt governance, decision decomposition, reflective reframing, action prioritization, escalation, and feedback. Each layer has a distinct function in the decision support workflow, but the layers operate together to preserve boundedness, traceability, privacy sensitivity, and human agency. Through this architecture, FounderPulse-AI is positioned as a responsible intelligent decision-support system rather than an unrestricted chatbot or psychological intervention.

5. Research Methods

This study adopted a design science research approach with a pilot field evaluation. The approach was appropriate because the study developed and evaluated an artifact intended to address a relevant organizational and technological problem [10]. The design science research methodology provided the procedural logic by connecting problem identification, solution objectives, artifact design, demonstration, evaluation, and communication [11].

The study followed three connected phases. The first phase identified founder decision support needs and responsible AI boundaries through a qualitative needs analysis. The second phase translated these needs into the FounderPulse-AI prototype. The third phase evaluated the prototype through a four-week pilot, comparing it with a structured digital reflection journal. The phases were carried out between January 2025 and April 2026 within the Greek startup ecosystem. Phase 1 was conducted through semi-structured online video interviews with founders, co-founders, and ecosystem support actors. Phase 2 was conducted by the research team through iterative artifact development, prompt specification, support corpus preparation, boundary rule design, and internal testing. Phase 3 was delivered through a web-based prototype platform over four consecutive weeks, with baseline and final questionnaires, weekly use logs, and post-use feedback. This sequence ensured that the artifact was grounded in founder needs, specified through system modules and responsible design principles, demonstrated in use, and evaluated against both intelligent support outcomes and non-AI structured reflection.

The first phase used a qualitative needs analysis with 20 participants, including 12 Greek startup founders or co-founders and 8 ecosystem support actors. The ecosystem support actors included 3 startup mentors, 2 incubator managers, 2 accelerator staff members, and 1 entrepreneurship coach.

Participants were recruited through startup ecosystem networks, incubator and accelerator contacts, founder communities, and professional networks. Data were collected between January 2025 and March 2025 through semi-structured online interviews. Each interview lasted approximately 45–60 minutes and followed an interview guide covering founder decision pressure, existing support gaps, AI acceptability, privacy concerns, system boundaries, and escalation expectations. Interviews were audio-recorded and transcribed for analysis. The qualitative data were coded through reflexive thematic analysis, which treats coding and theme development as an interpretive process rather than as a mechanical extraction of themes from the data [24]. The coded results were used to derive design requirements for FounderPulse-AI and to connect participant needs with system modules, responsible AI safeguards, and evaluation criteria.

The second phase developed the FounderPulse-AI prototype as a rule-constrained, retrieval-augmented, prompt-governed conversational decision-support system. Prototype development took place between February and March 2025. The research team created the support corpus, specified the prompt governance rules, defined out-of-scope and escalation triggers, prepared privacy and boundary notices, and conducted internal testing before the pilot evaluation. Retrieval-augmented generation was suitable because it connects generative outputs with external knowledge resources and can improve grounding for knowledge-intensive tasks [12]. The prototype included a curated knowledge base containing decision prompts, reflective questioning scripts, non-clinical strain language, action prioritization templates, privacy notices, responsible AI rules, and escalation guidance. The prototype was implemented as a web-based conversational interface using a retrieval-augmented large language model and delivered to participants via secure online access. These materials were used to ground system responses, constrain interaction, and reduce unrestricted generation.

The prototype operationalized six design principles. These principles were boundedness, transparency, privacy sensitivity, actionability, human escalation, and non-clinical positioning. They followed responsible AI guidance that emphasizes risk management, human oversight, transparency, privacy, accountability, and trustworthy system behavior [8], [15]. The prototype, therefore, included a boundary notice at the start of use, a restricted input structure, non-clinical response rules, retrieval grounding, prompt constraints, risk signal detection, and escalation responses. The system did not provide diagnosis, therapy, clinical interpretation, or medical advice.

The third phase evaluated FounderPulse-AI through a four-week field pilot with 100 Greek startup founders and co-founders from March to April 2026. Participants were eligible if they were active founders or co-founders of a Greek startup and were currently involved in venture decision-making. Participants were excluded if they were not actively involved in startup decision-making or if they sought clinical, therapeutic, legal, or financial advisory support rather than bounded decision reflection. Fifty participants used FounderPulse-AI, and 50 participants used a structured digital reflection journal. Participants were allocated to two equal conditions after eligibility screening. Allocation was designed to preserve baseline comparability across age, founder role, entrepreneurial experience, prior AI tool use, recent high-pressure decisions, baseline reflective clarity, baseline perceived decision quality, baseline venture resilience, expected usefulness, expected ease of use, trust in AI, and privacy concern. Baseline equivalence was assessed before interpreting the outcomes. The FounderPulse-AI condition was delivered through a web-based conversational interface, while the structured reflection journal was delivered through an online digital journal template. Baseline and final questionnaires were administered through online survey forms, weekly use logs were collected through participant self-report entries submitted electronically, and post-use feedback was collected through an online feedback questionnaire. The comparison condition was appropriate because both groups received structured reflection support, while only the FounderPulse-AI group received adaptive conversational interaction, retrieval grounding, prompt governance, and safety-aware escalation. This design allowed the evaluation to examine whether the intelligent system added value beyond structured reflection alone. Because the study was a pilot evaluation, the findings were interpreted as evidence of feasibility, acceptability, and usability rather than as definitive evidence of effectiveness [17]. Figure 3 shows the pilot evaluation design, including recruitment, baseline measurement, group assignment, four-week use, weekly logs, final measurement, post-use feedback, and integrated evaluation.

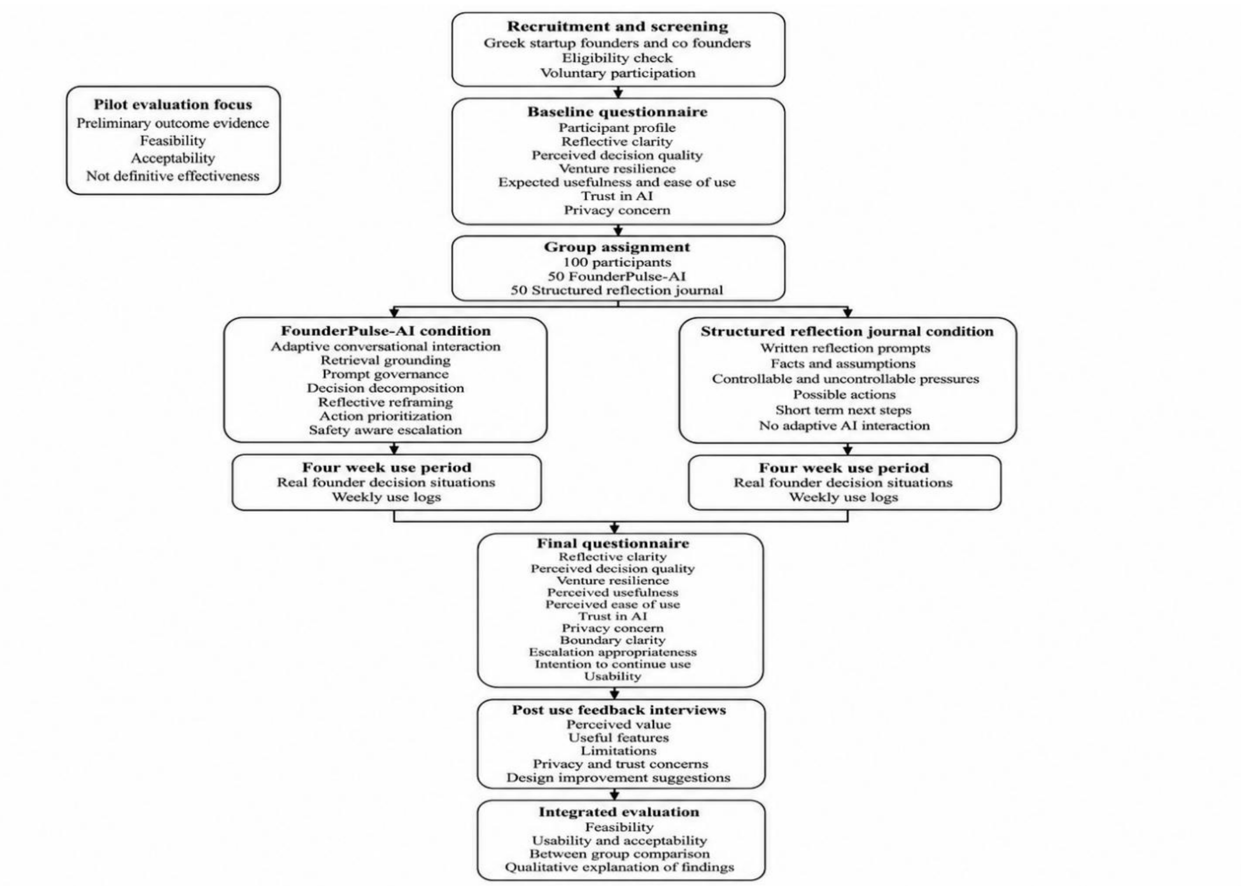

Figure 3. Pilot evaluation design

Participants in the FounderPulse-AI condition were instructed to use the prototype when they faced a current founder decision or pressure situation during the four weeks. Each use involved entering a brief decision description, reviewing the structured decision frame, considering the action options, and rating the support's usefulness. Participants in the comparison condition used a structured digital reflection journal with equivalent prompts covering decision description, facts and assumptions, controllable and uncontrollable pressures, possible actions, and next steps. The journal condition did not include conversational interaction, retrieval grounding, adaptive feedback, or safety-aware escalation. Both groups completed baseline and final questionnaires, and all eligible participants were retained through the final measurement. The study also collected weekly use logs and post-use qualitative feedback.

The primary outcomes were reflective clarity and perceived decision quality. Reflective clarity captured the founder's perceived ability to understand the decision problem, distinguish facts from assumptions, identify controllable and uncontrollable pressures, and formulate next steps. Perceived decision quality captured the founder's judgment that the decision process became clearer, more structured, and more actionable after using the assigned support tool. These constructs were developed from the system's design requirements and reviewed for alignment with the artifact's purpose before pilot testing.

The secondary outcome was venture resilience. Venture resilience captured the founder's perceived capacity to continue adaptive action, maintain learning, and respond constructively to venture pressure. This measure was framed as an entrepreneurial action construct rather than as a clinical resilience measure. This choice was consistent with entrepreneurship research that connects stress, coping, and resilience through responses to adversity [5], and with recent evidence showing that recovery experiences are associated with well-being and burnout among business owners [6]. The study did not assess mental health outcomes.

The pilot evaluation measured primary outcomes, secondary outcomes, system evaluation variables, responsible AI evaluation conditions, and feasibility indicators. The system evaluation variables included perceived usefulness, perceived ease of use, trust in AI, usability, privacy concern, boundary clarity, escalation appropriateness, and intention to continue use. Perceived usefulness and perceived ease of use were included because technology acceptance research identifies them as central determinants of user acceptance [18]. Trust in AI was incorporated because it shapes how users interpret and accept AI-supported outputs [19]. Usability was assessed using the System Usability Scale (SUS), a widely used standardized measure of perceived usability [25]. Privacy concern was included because founder inputs may involve sensitive venture, strategic, or personal pressure information, making privacy sensitivity a central requirement of responsible AI design [15], [20].

Table 3 summarizes each construct, its role in the study, source basis, timing, example item focus, and analysis. It also aligns the measurement strategy with the artifact's intended contribution: reflective clarity and perceived decision quality assess decision recovery, venture resilience captures cautious entrepreneurial adaptation, adoption variables assess practical use, and boundary clarity and escalation appropriateness evaluate whether responsible AI safeguards were experienced as functional system features.

Table 3. Measurement constructs and evaluation timing

Construct	Role in the study	Source basis	Timing	Example item focus	Analysis
Reflective clarity	Primary outcome	Study-specific scale derived from FounderPulse-AI design requirements	Baseline and final questionnaire	Ability to clarify the decision problem, separate facts from assumptions, identify missing information, and distinguish controllable from uncontrollable pressures	Internal consistency, baseline to final change, between-group comparison, and effect size
Perceived decision quality	Primary outcome	Study-specific scale aligned with AI-based decision support logic	Baseline and final questionnaire	Perceived improvement in option comparison, tradeoff recognition, action selection, and decision structure	Internal consistency, baseline to final change, between-group comparison, and effect size
Venture resilience	Secondary outcome	Study specific entrepreneurial adaptation informed by resilience and entrepreneur recovery literature.	Baseline and final questionnaire	Perceived ability to continue adaptive action, recover focus, learn from pressure, and maintain venture progress	Internal consistency, baseline to final change, between-group comparison, and cautious interpretation
Perceived usefulness	Technology acceptance outcome	Adapted from the Technology Acceptance Model measurement logic	Baseline expectation and final questionnaire	The extent to which the tool helps the founder handle decision pressure more effectively	Internal consistency, final group comparison, and adoption analysis

Construct	Role in the study	Source basis	Timing	Example item focus	Analysis
Perceived ease of use	Technology acceptance outcome	Adapted from the Technology Acceptance Model measurement logic	Baseline expectation and final questionnaire	The extent to which the tool is clear, understandable, and easy to use without unnecessary effort	Internal consistency, final group comparison, and adoption analysis
Usability	System evaluation outcome	System Usability Scale	Final questionnaire	Overall usability, simplicity, consistency, and ease of interaction	System Usability Scale score and final group comparison
Trust in AI	Adoption and risk condition	Adapted from trust in automation and AI trust measurement research [26]	Baseline and final questionnaire	Confidence that the system supports reflection within its stated limits and does not overstate its authority	Internal consistency, baseline to final change, and adoption analysis
Privacy concern	Responsible AI risk condition	Adapted from information privacy concern research [27]	Baseline and final questionnaire	Concern about entering sensitive venture, personal, strategic, or confidential information	Internal consistency, baseline to final change, and adoption and risk analysis
Boundary clarity	Responsible AI evaluation condition	Study-specific scale derived from responsible AI and non-clinical design principles	Final questionnaire	Clarity that the system is not clinical, does not provide therapy, and does not replace mentors, advisers, or professionals.	Internal consistency, descriptive analysis, and relationship with trust and intention to continue use

Construct	Role in the study	Source basis	Timing	Example item focus	Analysis
Escalation appropriateness	Responsible AI evaluation condition	Study specific scale derived from safety-aware escalation requirements	Final questionnaire and post-use feedback	Perceived appropriateness of redirecting high-risk or out-of-scope situations toward human support	Descriptive analysis and qualitative explanation
Intention to continue use	Practical adoption outcome	Adapted from technology acceptance research	Final questionnaire	Willingness to use the tool again for bounded founder decision reflection	Internal consistency, adoption analysis, and exploratory prediction by usefulness, trust, usability, and privacy concern
Use frequency	Feasibility indicator	Study specific behavioral log	Weekly log	Number of uses per week and type of decision situation addressed	Descriptive analysis, adherence assessment, and relationship with perceived value
Perceived value of support	Explanatory evaluation outcome	Study specific post-use assessment	Weekly log and post-use feedback interview	Participant's explanation of how the tool helped or failed to help with decision reflection	Thematic analysis and integration with quantitative findings

The qualitative data from the needs analysis and post-study feedback were analyzed through reflexive thematic analysis [24]. The analysis focused on founder decision pressure situations, perceived usefulness of intelligent support, trust and privacy concerns, perceived limitations, escalation expectations, and differences between structured journaling and AI-supported interaction.

The qualitative findings were used to explain the quantitative results rather than duplicate them. This allowed the study to show not only whether FounderPulse-AI was perceived as useful, but also why founders considered it useful, limited, trustworthy, or risky.

The quantitative analysis proceeded in several stages. First, descriptive statistics summarized participant characteristics, system use, and baseline measures. Second, internal consistency was examined for multi-item constructs. Third, baseline-to-final differences were calculated within each condition. Fourth, between-group differences were examined using independent-samples t-tests on change scores. Fifth, 95% confidence intervals, p-values, and Cohen's d effect sizes were reported because the study was a pilot evaluation, and practical significance was important. This analysis focused on feasibility, usability, baseline-to-final change, between-group comparisons, effect sizes, and qualitative explanations.

Participants received information about the study purpose, voluntary participation, data use, privacy protections, and the non-clinical status of FounderPulse-AI before participation. Informed consent was obtained before data collection. The system instructed participants not to enter medical information, confidential venture secrets, legal documents, credentials, investor documents, or unnecessary personal data. This approach followed the principle of data minimization, which requires personal data to be adequate, relevant, and limited to what is necessary for the stated purpose [20]. It also responded to generative AI risk guidance concerning prompt injection and sensitive information disclosure in large language model applications [22], [23]. When a participant used high-risk language or requested clinical support, the system paused the decision-support workflow and redirected the participant to human or professional support.

6. Results

6.1 Sample and Data Integrity

The pilot evaluation included 100 eligible Greek startup founders and co-founders. Participants were evenly distributed across the two study conditions: 50 assigned to FounderPulse-AI and 50 to the structured reflection journal. The sample included 38 founders and 62 co-founders. The mean participant age was 39.37 years. The startup stages represented in the sample were seed, early growth, pre-seed, and growth. The most frequent sectors were fintech, health technology, tourism technology, clean technology, education technology, creative technology, software-as-a-service, retail technology, agri-food technology, and logistics technology. All participants were marked as eligible in the dataset.

The two study conditions were balanced at baseline. The FounderPulse-AI and structured reflection journal groups were similar in age, entrepreneurial experience, prior AI tool use, recent high-pressure decisions, baseline reflective clarity, baseline perceived decision quality, baseline venture resilience, expected usefulness, expected ease of use, trust in AI, and privacy concern. This baseline balance supports a cautious comparison between the intervention and comparison conditions. A study dataset audit showed complete participant matching across the participant profile, baseline item, final item, score, and weekly log sheets. No missing values were observed in the main score variables used for the pilot evaluation. The Phase 1 needs analysis was based on 20 qualitative participants, including 12 Greek startup founders or co-founders and 8 ecosystem support actors. The needs analysis and post-use feedback were treated as complementary qualitative evidence. They were used to derive system requirements and to explain pilot patterns, whereas the 100-participant sample describes the pilot evaluation rather than the full qualitative evidence base.

Internal consistency was acceptable across the main multi-item constructs. Baseline reflective clarity, perceived decision quality, venture resilience, expected usefulness, expected ease of use, trust in AI, and privacy concern all showed acceptable to strong reliability. Final reflective clarity, perceived decision quality, venture resilience, perceived usefulness, perceived ease of use, trust in AI, privacy concern, boundary clarity, escalation appropriateness, intention to continue use, and recoded usability also showed acceptable to strong reliability. This supports the use of the computed scale scores for pilot evaluation, while recognizing that reflective clarity, perceived decision quality, boundary clarity, and escalation appropriateness remain study-specific constructs that require further validation in later studies. Reliability was interpreted cautiously because Cronbach's alpha is affected by item number, dimensionality, and scale context, and it should not be treated as the sole indicator of measurement quality [28].

6.2 Needs Analysis and Design Requirements

The Phase 1 coded needs analysis identified founder pressure as a decision support problem rather than a general well-being problem. Participants and ecosystem actors emphasized pressure related to team conflict, investor communication, customer acquisition, market-entry uncertainty, time overload, and urgent prioritization. These findings supported the decision to design FounderPulse-AI around structured decision decomposition rather than motivational support.

Table 4 presents the traceability between the coded needs analysis, the interpreted design requirements, the system responses, and the evaluation criteria. This traceability is important because the study evaluates FounderPulse-AI as a design science artifact rather than as a speculative conversational agent. It shows that the artifact requirements were derived from founder decision pressure, support gaps, privacy concerns, and escalation expectations rather than from generic chatbot functionality. The needs analysis therefore justifies the treatment of decision decomposition, action prioritization, privacy sensitivity, non-clinical positioning, and safety-aware escalation as core system functions.

Table 4. Design requirements derived from needs analysis

Empirical need from Phase 1	Representative coded evidence	Interpreted design requirement	FounderPulse-AI system response	Evaluation criterion
Founders faced urgent decision pressure under uncertainty	Team conflict, investor communication, customer acquisition	The system must help founders convert vague pressure into	The user input layer captures current pressure, decision time	Participants report improved reflective clarity and clearer decision framing.

Empirical need from Phase 1	Representative coded evidence	Interpreted design requirement	FounderPulse-AI system response	Evaluation criterion
	pressure, market entry uncertainty, time overload, and urgent prioritization	structured decision problem.	horizon, uncertainty, and support need.	
Founders lacked immediate access to mentors, advisers, or other ecosystem support actors during periods of decision-making pressure.	Mentor not immediately available and difficulty preparing for mentor meetings	The system must provide structured support when human support is not immediately available.	The conversational workflow provides guided decision decomposition, reflective reframing, and action prioritization.	Participants report that the tool helped them prepare clearer questions or next steps before seeking human support.
Founders struggled to separate facts, assumptions, constraints, and missing information.	Facts versus assumptions separation and decision decomposition appeared as valued support functions in post-use feedback.	The system must make the decision situation more explicit and analytically manageable.	The decision decomposition layer separates facts, assumptions, constraints, missing information, options, and tradeoffs.	Participants report stronger reflective clarity and perceived decision quality.
Founders needed short-term action priorities rather than broad motivational advice.	Action prioritization was the most frequently valued FounderPulse-AI feature in post-use feedback.	The system must convert reflection into feasible, time-bound managerial action.	The action prioritization layer produces short, feasible, time-bound next steps for the next decision window.	Participants report that the output was actionable, specific, and useful for immediate venture decisions.
Founders experienced pressure that was partly	Market entry uncertainty, time overload, and investor	The system must help founders distinguish controllable,	The reflective reframing layer classifies pressures and redirects	Participants report a better ability to focus on actionable issues rather than

Empirical need from Phase 1	Representative coded evidence	Interpreted design requirement	FounderPulse-AI system response	Evaluation criterion
controllable and partly outside their control.	communication were coded as pressure situations with mixed controllability.	partly controllable, and uncontrollable pressures.	attention toward action within the founder's control.	undifferentiated strain.
Founders worried that AI support could overreach into therapy, diagnosis, or professional advice.	Need for clear boundaries, professional support for clinical requests, and pausing out-of-scope requests.	The system must remain bounded and explicitly non-clinical .	Boundary notices and prompt governance rules state that the system does not diagnose, treat, monitor, or provide therapy.	Participants report a clear understanding of what the system can and cannot do.
Founders hesitated to enter sensitive business, personal, legal, investor, or strategic information.	Need a clear data warning; avoid confidential venture data; allow only minimal input; avoid personal data; and be concerned about storage.	The system must follow data minimization and privacy-sensitive design principles.	The input layer requests only decision-relevant information and warns participants not to enter confidential or unnecessarily sensitive information.	Participants report clearer data entry boundaries, while privacy concern is monitored as a responsible AI evaluation condition
Founders could use high-risk or out-of-scope language under severe pressure.	Redirect severe distress, professional support for clinical requests, pause out-of-scope requests, recommend mentor contact, and human oversight	The system must detect out-of-scope requests and redirect users to human support.	The boundary and risk classification layer and safety-aware escalation layer pause decision support and redirect users to appropriate human or professional support.	Escalation responses are judged appropriate, bounded, and clear.

Empirical need from Phase 1	Representative coded evidence	Interpreted design requirement	FounderPulse-AI system response	Evaluation criterion
Founders preferred AI support only when it preserved their judgment and did not replace human mentoring.	Mentor integration, prepared mentor conversation, combined with mentor support, and trusted because the system was bounded	The system must preserve the founder agency and position AI as support rather than authority.	System outputs are framed as reflective prompts, decision summaries, and action options that the founder can accept, reject, or modify.	Participants report that the system supported their judgment rather than replacing it.
Founders needed transparency about how the system produced responses.	Trusted because transparent, trusted because bounded, and trusted but cautious	The system must make its support logic understandable and traceable.	Retrieval grounding connects outputs to approved decision prompts, startup support resources, responsible AI rules, privacy notices, and escalation guidance.	Participants report that outputs were understandable, bounded, and consistent with the stated system purpose.
Founders needed reassurance that the system was safe enough for well-being-adjacent but non-clinical use.	Clear boundary needs, escalation expectations, and non-clinical support limits	The system must integrate responsible AI safeguards into the core workflow rather than add them after design.	Prompt governance, privacy warnings, non-clinical boundary rules, and escalation logic are embedded across the workflow.	Participants report perceived safety, clarity of boundaries, and trust in AI support.
Founders differed in their levels of trust, familiarity with AI, and willingness to continue using it.	Trust varied between trusted because transparent, trusted because bounded, neutral trust, and trusted but cautious.	The evaluation must assess adoption conditions and decision-support outcomes.	The feedback layer collects perceived usefulness, ease of use, usability, trust in AI, privacy concern, and intention to continue use.	The evaluation shows whether the artifact is usable, acceptable, trusted, and practically viable.

The needs analysis also supported the system's responsible AI boundaries. Privacy-related codes emphasized clear data warnings, avoidance of confidential venture data, minimal input, avoidance of personal data, and concern about storage. Escalation-related codes emphasized redirecting severe distress, providing professional support for clinical requests, pausing out-of-scope requests, recommending mentor contact, and maintaining human oversight. These findings support the design decision to make privacy sensitivity, non-clinical positioning, and safety-aware escalation core system functions rather than optional safeguards.

6.3 Pilot Evaluation Outcomes

Table 5 reports the pilot evaluation results, including baseline and final means, change scores, between-group differences, confidence intervals, p-values, and effect sizes. The pattern was qualified rather than uniformly favorable. FounderPulse-AI produced stronger improvements in outcomes aligned with adaptive intelligent support and responsible system design, while venture resilience improved more modestly, and privacy concerns decreased more in the structured journal condition. The following subsections interpret these results by distinguishing decision recovery, adoption value, responsible AI safeguards, and feasibility evidence.

Table 5. Pilot evaluation results

Outcome	FounderPulse-AI baseline	FounderPulse-AI final	FounderPulse-AI change	Journal baseline	Journal final	Journal change	Between-group difference	95% CI	p	Cohen d
Reflective clarity	3.44	4.52	1.08	3.46	3.96	0.50	0.59	0.38 to 0.80	< .001	1.11
Perceived decision quality	3.65	4.50	0.85	3.68	4.22	0.54	0.31	0.10 to 0.51	.004	0.60
Venture resilience	3.79	4.26	0.46	3.89	4.14	0.25	0.22	0.00 to 0.43	.050	0.40
Perceived usefulness	4.28	5.43	1.16	4.26	4.58	0.32	0.84	0.41 to 1.28	< .001	0.77
Perceived	4.29	5.41	1.11	4.13	5.12	0.99	0.12	-	.496	0.14

Outcome	FounderPulse-AI baseline	FounderPulse-AI final	FounderPulse-AI change	Journal baseline	Journal final	Journal change	Between-group difference	95% CI	p	Cohen d
ease of use								0.23 to 0.47		
Trust in AI	3.38	5.29	1.91	3.63	4.09	0.46	1.45	1.00 to 1.91	< .001	1.26
Privacy concern	4.08	3.67	-0.41	3.93	2.94	-0.98	0.57	0.10 to 1.05	.018	0.48
Boundary clarity	Not applicable	5.69	Not applicable	Not applicable	5.22	Not applicable	0.47	0.22 to 0.72	< .001	0.73
Escalation appropriateness	Not applicable	5.35	Not applicable	Not applicable	4.61	Not applicable	0.74	0.44 to 1.04	< .001	0.99
Intention to continue use	Not applicable	5.09	Not applicable	Not applicable	3.99	Not applicable	1.10	0.80 to 1.40	< .001	1.43
Usability	Not applicable	77.00	Not applicable	Not applicable	65.70	Not applicable	11.30	7.03 to 15.57	< .001	1.05

6.4 System Evaluation and Adoption Outcomes

FounderPulse-AI demonstrated a clear advantage in perceived usefulness, with a mean improvement of 1.16 compared with 0.32 for the structured journal. The between-group difference was 0.84, and the effect size was moderate to strong, indicating that founders perceived the adaptive intelligent system as more useful than static structured reflection for managing founder decision pressure. By contrast, the two conditions did not meaningfully differ in ease of use. FounderPulse-AI improved by 1.11, while the journal improved by 0.99, yielding a small and statistically not meaningful between-group difference. This pattern suggests that the intelligent system did not impose a usability burden relative to structured journaling. Its distinctive contribution lay in perceived usefulness, reflective clarity, decision quality, and trust.

Trust in AI also improved more strongly in the FounderPulse-AI condition. The mean increase in trust was 1.91 in the FounderPulse-AI condition and 0.46 in the journal condition, yielding a between-group difference of 1.45 and a large effect size. This pattern suggests that trust in AI-supported reflection was associated with clear boundaries, transparency, and clear limits on system authority.

FounderPulse-AI further outperformed the journal condition on boundary clarity, escalation appropriateness, intention to continue use, and usability. The final mean for boundary clarity was 5.69 for FounderPulse-AI and 5.22 for the journal, while the final mean for escalation appropriateness was 5.35 and 4.61, respectively. Intention to continue use was also higher in the FounderPulse-AI condition, with a mean of 5.09 compared with 3.99 in the journal condition. Usability followed the same pattern, with FounderPulse-AI achieving a mean score of 77.00 compared with 65.70 for the journal.

Privacy concern revealed the clearest responsible AI tradeoff. Although concern decreased in both conditions, the reduction was stronger in the structured journal condition. FounderPulse-AI therefore provided stronger decision support and adoption value, while the journal created lower disclosure demands. This pattern indicates that wider deployment of adaptive AI support requires clearer storage notices, stronger user control, minimal input options, and visible data minimization safeguards.

6.5 Weekly Use and Feasibility Evidence

The weekly logs provide additional evidence of feasibility. Participants in the FounderPulse-AI condition used the assigned tool more frequently than participants in the structured journal condition. FounderPulse-AI participants recorded 569 uses over the four weeks, averaging 11.38 per participant. Journal participants recorded 405 uses, averaging 8.10 per participant. FounderPulse-AI participants also spent more time with the tool, with 91.56 minutes per participant compared with 58.26 minutes in the journal condition.

FounderPulse-AI also received stronger weekly ratings for helping with next steps and perceived value. The mean helped next steps score was 4.52 for FounderPulse-AI and 3.69 for the journal. The mean weekly value score was 5.45 for FounderPulse-AI and 4.33 for the journal. Discomfort was slightly higher in the FounderPulse-AI condition, with a mean of 2.37 compared with 1.92 for the journal. This combination reinforces the pilot’s central interpretation. FounderPulse-AI was more engaging and more valuable for decision support, but it also created stronger sensitivity around privacy, boundaries, and comfort.

Participants in both groups reported seeking human support at the same frequency across the weekly logs. Each condition recorded 35 instances of human support seeking. This pattern is important because it suggests that FounderPulse-AI did not replace human support. Instead, the qualitative evidence indicates that it helped some founders prepare more focused mentor conversations. This is consistent with the paper's non-replacement logic and supports the claim that responsible intelligent systems should complement rather than replace human support infrastructures.

6.6 Qualitative Explanations and Redesign Priorities

Table 6 integrates the Phase 1 coded needs analysis with post-use feedback and explains why the quantitative findings followed a qualified pattern. FounderPulse-AI users most often valued action prioritization, decision decomposition, adaptive prompts, preparation for mentor conversations, and bounded system behavior. Journal users valued separating facts from assumptions and self-led reflection, but they also described the journal as repetitive and less engaging. These patterns indicate that FounderPulse-AI added value through adaptiveness and action translation, while the journal remained useful as a lower disclosure reflection mode.

The privacy codes clarify the main responsible AI tradeoff. FounderPulse-AI users requested a stronger privacy dashboard, clearer storage information, a minimal input mode, and more visible control over the information they entered. Journal users reported lower AI-related risk because the journal was self-led and did not involve adaptive generation. This contrast indicates that privacy concerns are a central design condition for founder-facing AI because founders may disclose sensitive venture, strategic, or personal information during use.

The feedback also identified redesign priorities. FounderPulse-AI should retain decision decomposition, action prioritization, boundary clarity, and safety-aware escalation. It should also strengthen its curated knowledge base with startup- and sector-specific resources. A mentor handoff function should convert the founder's decision frame, assumptions, options, and questions into a concise summary for discussion with a mentor, adviser, or trusted peer. These priorities strengthen usefulness while preserving human agency, privacy sensitivity, and non-clinical boundaries.

Table 6. Qualitative themes, explanatory patterns, and design implications

Theme	Evidence base	Main condition pattern	Interpretation	Design implication
Decision pressure required a structured decomposition.	Phase 1 interviews showed recurring pressure on team conflict and investor communication (4 mentions each), followed by customer-acquisition pressure, market-entry uncertainty, time	Both conditions benefited from structure, but FounderPulse-AI participants emphasized decision decomposition while journal participants	Founders needed support that made pressure analytically manageable rather than motivational. The results explain why reflective clarity improved in both groups and	Retain decision decomposition as a core module and make the separation of facts, assumptions, constraints, missing information, options, and

Theme	Evidence base	Main condition pattern	Interpretation	Design implication
	overload, and urgent prioritization (2 mentions each). Post-use feedback showed that decision decomposition was useful for three FounderPulse-AI participants and that facts versus assumptions was useful for six journal participants.	emphasized facts versus assumptions.	why the AI condition showed stronger gains.	tradeoffs visible in the user interface.
Action prioritization was the clearest added value of FounderPulse-AI	Post-use feedback showed that action prioritization was the most useful feature for six FounderPulse-AI participants, whereas one journal participant emphasized an action list.	The theme was concentrated in the FounderPulse-AI condition.	Adaptive AI support was most valuable when it translated reflections into short, actionable next steps. This explains the stronger FounderPulse-AI results for perceived decision quality, usefulness, and intention to continue use.	Preserve the action prioritization layer and require each output to include a concise set of time-bound next steps that the founder can accept, reject, modify, or discuss with a mentor.
Adaptive prompts increased perceived value but created a need for stronger	Adaptive prompts were coded as most useful by two FounderPulse-AI participants. Stronger domain resources	FounderPulse-AI users valued adaptation, while some journal users missed adaptive feedback.	The findings suggest that adaptive interaction adds value beyond static reflection, but the	Expand the curated knowledge base with sector-specific decision prompts, founder support scenarios,

Theme	Evidence base	Main condition pattern	Interpretation	Design implication
contextual resources.	were requested by four FounderPulse-AI participants. Journal participants requested adaptive prompts once and adaptive feedback twice.		system needs deeper startup-specific retrieval resources to avoid generic outputs.	mentor preparation templates, and domain-relevant startup resources.
Human support remained central.	Phase 1 interviews included difficulty preparing for mentor meetings (6 mentions) and the mentor not being immediately available (1 mention). Post-use feedback showed that they wanted mentor integration for four FounderPulse-AI participants, prepared mentor conversations for one FounderPulse-AI participant and two journal participants, and combined with a mentor for five journal participants.	Both conditions pointed to human support, but the requested integration differed across conditions. FounderPulse-AI users wanted mentor integration, while journal users wanted the journal to connect more directly with mentoring.	Participants did not treat either tool as a replacement for human support. The findings support the paper's non-replacement logic and show that AI should prepare founders for more focused human conversations.	Add a mentor handoff option that summarizes the decision frame, assumptions, options, and questions for discussion with a mentor, adviser, or trusted peer.
Boundary clarity supported trust, but caution remained.	FounderPulse-AI feedback showed trusted because transparent for 4 participants, trusted because bounded for 4 participants, and	FounderPulse-AI trust depended on transparency and boundedness, while journal trust was linked to self-direction and	The AI condition generated stronger trust, but trust was conditional rather than unconditional. Participants trusted the system when its	Keep boundary statements, non-clinical notices, and output limitations visible throughout the workflow rather

Theme	Evidence base	Main condition pattern	Interpretation	Design implication
	trusted but cautious for 4 participants. Journal feedback showed no AI risk for four participants, neutral trust for three, wanted feedback for three, and trusted because self-led for two.	absence of AI risk.	limits were clear, and it did not overstate its authority.	than presenting them only at the start.
Privacy concern remained the main responsible AI design issue.	Phase 1 privacy codes included: need clear data warnings with 5 mentions, avoid confidential venture data with 4 mentions, allow minimal input only with 4 mentions, avoid personal data with 4 mentions, and concern about storage with 1 mention. Post-use feedback indicated moderate privacy concerns among six FounderPulse-AI participants and five journal participants. FounderPulse-AI participants also requested a better privacy dashboard 4 times and a clearer storage policy 1 time.	Privacy concerns appeared in both groups, but redesign requests were more specific in the FounderPulse-AI condition.	FounderPulse-AI provided stronger decision-support value, but participants wanted clearer control over which information was entered, stored, and used. This explains why privacy concerns decreased more in the journal condition.	Add a privacy dashboard, a clearer storage notice, a minimal-input mode, a delete-input option, and stronger warnings against entering confidential venture or personal information.

Theme	Evidence base	Main condition pattern	Interpretation	Design implication
Escalation expectations supported the safety-aware design.	Phase 1 escalation codes included: redirect severe distress (5 mentions), professional support for clinical requests (5 mentions), pause out-of-scope requests (4 mentions), recommend mentor contact (3 mentions), and human oversight needed (1 mention).	The need for escalation was established before prototype evaluation and was consistent across founder and support actor perspectives.	Participants expected the system to stop when requests exceeded the founder's decision support. This supports the safety-aware escalation layer and the non-clinical positioning of the artifact.	Retain escalation logic and distinguish between mentor escalation for business support, adviser escalation for professional issues, and emergency or qualified professional support for high-risk or clinical requests.
The journal condition remained useful but less adaptive	Journal feedback showed as too repetitive (4 mentions), less engaging (3 mentions), wanting adaptive feedback (2 mentions), and needing examples (1 mention). Journal redesign requests included shortening the journal to five mentions and combining it with mentor, also with five mentions.	The theme was concentrated in the structured reflection journal condition.	The journal reduced privacy concerns more than FounderPulse-AI, but it appeared less engaging and less adaptive. This helps explain why the journal performed less strongly on usefulness, usability, trust, and intention to continue use	Retain the journal as a valid comparison condition and interpret it as a lower-disclosure, lower-interactivity reflection method rather than a weak control.

7. Discussion

This study evaluated FounderPulse-AI as a responsible intelligent decision support artifact for startup founders operating under entrepreneurial strain. The findings should be interpreted as a qualified evaluation of the artifact rather than as evidence that AI is inherently superior to structured reflection. The central contribution is that adaptive, intelligent support can improve decision recovery, adoption value, usability, boundary clarity, escalation appropriateness, and intention to continue use when accompanied by visible governance, privacy controls, and human support pathways.

The strongest interpretation concerns decision recovery. FounderPulse-AI improved reflective clarity more than the structured journal because it helped founders move from undifferentiated pressure to a clearer decision frame. This extends AI decision-support research by showing that value in entrepreneurial contexts comes less from automated recommendations and more from scaffolding the founder's decision-making process [1]. The moderate advantage in perceived decision quality supports the same interpretation. The artifact translated founder pressure into facts, assumptions, missing information, options, trade-offs, and short-term actions. This connects early-stage startup decision research with an operational intelligent support workflow by showing how founder decision domains can be supported through decomposition and action prioritization rather than through generic encouragement [4].

The venture resilience result should remain cautious. FounderPulse-AI showed only a modest advantage over the journal, and the four-week pilot period cannot support strong claims about durable resilience, founder behavior, or venture performance. This finding is consistent with entrepreneurship research that treats resilience as a process shaped over time through stress, coping, recovery, and responses to adversity [5], [6]. The study, therefore, contributes a preliminary decision recovery mechanism rather than a general theory of entrepreneurial resilience.

The usefulness and continuation intention findings clarify where FounderPulse-AI added value. The AI condition produced stronger gains in perceived usefulness and a higher intention to continue use, while ease of use improved similarly in both conditions. This pattern is theoretically meaningful because technology acceptance research identifies usefulness and ease of use as distinct acceptance mechanisms [18]. The journal was easy to use, but less adaptive and less engaging. FounderPulse-AI, therefore, appears to have created added value through contextual guidance, adaptive response, and action-oriented support rather than through lower effort alone. For founders facing intelligent systems, simplicity is insufficient when a low-risk non-AI alternative is already available. The system must make the support process meaningfully more useful while preserving user agency and responsible boundaries.

The trust and privacy findings reveal the central responsible AI tradeoff in the study. FounderPulse-AI produced a larger increase in trust than the structured journal, consistent with research showing that trust influences acceptance of AI technologies through perceived usefulness and user attitudes [19]. The qualitative evidence indicates that this trust was conditional rather than blind. Founders trusted the system when it was transparent, bounded, careful about its limits, and unwilling to overstate its authority. At the same time, privacy concerns decreased more strongly in the journal condition, suggesting that adaptive AI support can increase perceived value while also heightening sensitivity to disclosure, storage, control, and accountability [8], [15], [20]. AI systems for founders therefore require visible data minimization, clear storage notices, deletion options, minimal input modes, and user control over what is entered and retained.

The qualitative findings clarify why the comparison matters. Structured journaling remained useful because it supported self-led reflection with lower disclosure demands. FounderPulse-AI added value by enabling adaptiveness and action translation, helping founders convert pressure into a clearer decision frame and a more immediate action path. This contrast prevents a simplistic pro-AI conclusion and supports a more precise claim about bounded intelligent decision support.

The findings also support retrieval-grounded interaction as a governance mechanism rather than a claim of autonomous intelligence. Retrieval-augmented generation research shows that external knowledge retrieval can support knowledge-intensive tasks by connecting outputs to external resources [12]. This study did not directly test retrieval accuracy, so it should not overstate technical performance. Its contribution is more specific. FounderPulse-AI used retrieval grounding to link outputs to approved decision prompts, startup support resources, responsible AI rules, privacy notices, and escalation guidance. This design reduced the risk of unrestricted conversational generation and made response control part of the artifact architecture.

Boundary clarity and escalation appropriateness are also central to the responsible AI contribution. FounderPulse-AI outperformed the journal on both outcomes, indicating that safety features can be evaluated as system functions rather than discussed only as ethical intentions. This finding translates responsible AI principles into measurable workflow features by showing how boundary messages, prompt governance, safety-aware escalation, and non-clinical positioning can be assessed through user evaluation [8], [15].

The design science contribution lies in the artifact chain that connects needs analysis, system requirements, implemented modules, pilot outcomes, and redesign implications [10], [11]. This chain shows how founder strain can be translated into a responsible intelligent support architecture and evaluated through decision recovery, adoption, usability, privacy concern, boundary clarity, escalation appropriateness, and intention to continue use.

The structured reflection journal strengthens the study's internal comparison because it was not a passive control condition. It supported reflection, reduced privacy concerns more strongly, and provided a lower-disclosure alternative. FounderPulse-AI nevertheless performed better on outcomes linked to adaptive intelligent support. This balanced comparison enhances credibility by identifying both the added value and the governance costs of founder-facing AI.

In practice, FounderPulse-AI should be positioned as a preparatory layer for human mentoring rather than as an autonomous support service. Its strongest use is to help founders clarify the decision frame, assumptions, options, and questions before seeking support from mentors, advisers, or trusted peers. Deployment should therefore prioritize mentor handoff, privacy dashboards, minimal-input modes, deletion controls, and clear boundary notices over unrestricted conversational capability.

The limitations are those of a pilot artifact evaluation. The study supports feasibility, usability, acceptability, and preliminary outcome interpretation, but it does not establish definitive effectiveness [17]. It relies on perceived decision quality and perceived venture resilience rather than objective venture performance. The four-week period limits conclusions about longer-term resilience, founder behavior, or venture outcomes. The Greek startup context provides empirical depth, but the model should also be tested in other entrepreneurial ecosystems. FounderPulse-AI was evaluated as a bounded pilot artifact rather than a fully deployed commercial system, so future work should examine performance under real-world scaling, extended retrieval resources, privacy-preserving deployment, and mentor-integrated workflows. Future research should also examine longer observation periods, objective decision follow-through, retrieval quality, prompt robustness, and safeguards against the disclosure of sensitive information and prompt injection [22], [23].

8. Conclusion

This paper focused on a specific problem in intelligent technology. It examined how a startup founder support system can provide adaptive decision support during entrepreneurial strain while remaining bounded, transparent, privacy-sensitive, and subordinate to human judgment. The pilot evidence shows that FounderPulse-AI added value beyond structured journaling, mainly by improving reflective clarity, perceived decision quality, perceived usefulness, trust, usability, boundary clarity, escalation appropriateness, and intention to continue use. The weaker venture resilience result narrows the claim. FounderPulse-AI should be interpreted as a decision recovery artifact, not as a resilience intervention or mental health technology.

The main contribution is the design and pilot evaluation of a bounded founder support architecture that integrates retrieval grounding, prompt governance, decision decomposition, reflective reframing, action prioritization, privacy-sensitive input design, and safety-aware escalation. The Responsible AI Founder Support Model adds a conceptual contribution by explaining how founder strain can be translated into reflective clarity, perceived decision quality, cautious venture resilience, and continued use conditions without replacing mentors, advisers, clinicians, or founder judgment.

Future research should move from pilot feasibility to stronger longitudinal and technical validation. The next stage should test FounderPulse-AI across different entrepreneurial ecosystems, longer observation periods, and more diverse founder profiles. It should also examine objective decision follow-through, mentor handoff quality, retrieval quality, response traceability, prompt robustness, privacy-preserving deployment, and safeguards against overreliance. These priorities align with recent work on generative AI governance, RAG trustworthiness, human-centered generative AI design, and RAG evaluation, which emphasizes factual grounding, transparency, robustness, privacy, accountability, and user control [21], [29], [30], [31], [32].

Acknowledgements

The authors thank the participating founders and ecosystem support actors who

contributed to the needs analysis, pilot evaluation, and post-use feedback.

Conflicts of Interest

The authors declare no conflict of interest.

References

[1] Kővári, A. AI for decision support: Balancing accuracy, transparency, and trust across sectors. Information, 2024, 15(11), 725. DOI: 10.3390/info15110725.

[2] Found.ation, EIT Digital and Hellenic Development Bank of Investments. Startups in Greece 2024 to 2025: Venture Financing Report. Found.ation, 2024.

[3] OECD and European Commission. Inclusive Entrepreneurship Policy Assessment Country Notes: Greece. OECD Publishing, 2025. DOI: 10.1787/fb58fe49-en.

[4] González, M.A., Terzidis, O., Lütz, P. and Heblich, B. Critical decisions at the early stage of start-ups: A systematic literature review. Journal of Innovation and Entrepreneurship, 2024, 13, 83. DOI: 10.1186/s13731-024-00438-9.

[5] Ahmed, A.E., Ucbasaran, D., Cacciotti, G. and Williams, T.A. Integrating psychological resilience, stress, and coping in entrepreneurship: A critical review and research agenda. Entrepreneurship Theory and Practice, 2022, 46(3), 497-538. DOI: 10.1177/10422587211046542.

[6] Le Moal, M., Thurik, R., Torrès, O. and Soenen, G. Mental health of entrepreneurs and daily recovery experiences. Small Business Economics, 2025. DOI: 10.1007/s11187-025-01087-2.

[7] Ali, R., Hussain, A., Nazir, S., Khan, S. and Khan, H.U. Intelligent decision support systems: An analysis of machine learning and multicriteria decision-making methods. Applied Sciences, 2023, 13(22), 12426. DOI: 10.3390/app132212426.

[8] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: AI RMF 1.0. NIST AI 100-1, 2023.

[9] World Health Organization. Ethics and Governance of Artificial Intelligence for Health: WHO Guidance. World Health Organization, 2021.

[10] Hevner, A.R., March, S.T., Park, J. and Ram, S. Design science in information systems research. MIS Quarterly, 2004, 28(1), 75-105. DOI: 10.2307/25148625.

[11] Peffers, K., Tuunanen, T., Rothenberger, M.A. and Chatterjee, S. A design science research methodology for information systems research. Journal of Management Information Systems, 2007, 24(3), 45-77. DOI: 10.2753/MIS0742-1222240302.

[12] Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Wang, M. and Wang, H. Retrieval-augmented generation for large language models: A survey. arXiv preprint arXiv:2312.10997, 2023.

[13] World Health Organization. Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models. World Health Organization, 2025.

[14] Guo, Z., Lai, A., Thygesen, J.H., Farrington, J., Keen, T. and Li, K. Large language models for mental health applications: Systematic review. JMIR Mental Health, 2024, 11, e57400. DOI: 10.2196/57400.

[15] High-Level Expert Group on Artificial Intelligence. Ethics Guidelines for Trustworthy AI. European Commission,

2019.
[16] Venable, J., Pries-Heje, J. and Baskerville, R. FEDS: A framework for evaluation in design science research. European Journal of Information Systems, 2016, 25(1), 77-89. DOI: 10.1057/ejis.2014.36.
[17] Eldridge, S.M., Chan, C.L., Campbell, M.J., Bond, C.M., Hopewell, S., Thabane, L. and Lancaster, G.A. CONSORT 2010 statement: Extension to randomised pilot and feasibility trials. BMJ, 2016, 355, i5239. DOI: 10.1136/bmj.i5239.
[18] Davis, F.D. Perceived usefulness, perceived ease of use, and user acceptance of information technology. MIS Quarterly, 1989, 13(3), 319-340. DOI: 10.2307/249008.
[19] Choung, H., David, P. and Ross, A. Trust in AI and its role in the acceptance of AI technologies. International Journal of Human-Computer Interaction, 2023, 39(9), 1727-1739. DOI: 10.1080/10447318.2022.2050543.
[20] European Union. General Data Protection Regulation, Article 5: Principles relating to processing of personal data. 2016.
[21] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1, 2024.
[22] OWASP Foundation. LLM01: Prompt Injection. OWASP Top 10 for Large Language Model Applications, 2025.
[23] OWASP Foundation. LLM02: Sensitive Information Disclosure. OWASP Top 10 for Large Language Model Applications, 2025.
[24] Braun, V. and Clarke, V. Using thematic analysis in psychology. Qualitative Research in Psychology, 2006, 3(2), 77-101. DOI: 10.1191/1478088706qp063oa.
[25] Brooke, J. SUS: A quick and dirty usability scale. In Jordan, P.W., Thomas, B., Weerdmeester, B.A. and McClelland, I.L., Eds., Usability Evaluation in Industry. Taylor and Francis, London, 1996, 189-194.
[26] Jian, J.Y., Bisantz, A.M. and Drury, C.G. Foundations for an empirically determined scale of trust in automated systems. International Journal of Cognitive Ergonomics, 2000, 4(1), 53-71. DOI: 10.1207/S15327566IJCE0401_04.
[27] Malhotra, N.K., Kim, S.S. and Agarwal, J. Internet users' information privacy concerns: The construct, the scale, and a causal model. Information Systems Research, 2004, 15(4), 336-355. DOI: 10.1287/isre.1040.0032.
[28] Tavakol, M. and Dennick, R. Making sense of Cronbach's alpha. International Journal of Medical Education, 2011, 2, 53-55. DOI: 10.5116/ijme.4dfb.8dfd.
[29] Zhou, Y., Liu, Y., Li, X., Jin, J., Qian, H., Liu, Z., Li, C., Dou, Z., Ho, T. and Yu, P.S. Trustworthiness in retrieval-augmented generation systems: A survey. arXiv preprint arXiv:2409.10102, 2024.
[30] Weisz, J.D., He, J., Muller, M., Hoefer, G., Miles, R. and Geyer, W. Design principles for generative AI applications. arXiv preprint arXiv:2401.14484, 2024.
[31] European Commission. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. 2024.
[32] Upadhyay, S., Thakur, N., Pradeep, R., Craswell, N., Campos, D. and Lin, J. Overview of the TREC 2025 Retrieval Augmented Generation (RAG) Track. arXiv preprint arXiv:2603.09891, 2026.