Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Signal Conversion Efficiency: Generative AI, Human Judgment, and Capital Formation in Token-Based Crowdfunding

Paul Momtaz

ppmomtaz@syr.edu

Syracuse University

Johan Wiklund

Syracuse University

Ye Yuan

Technical University of Munich

Research Article

Keywords: Artificial intelligence (AI), large language model (LLM), machine learning (ML), entrepreneurial finance, blockchain-based crowdfunding, initial coin offering (ICO)

Posted Date: November 13th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11595914/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Signal Conversion Efficiency: Generative AI, Human Judgment, and Capital Formation in Token-Based Crowdfunding

Abstract

Information asymmetries between entrepreneurs and financiers shape the efficiency of entrepreneurial finance markets. Drawing on the distinction between information-endowment asymmetry (unequal access to signals) and information-utilization asymmetry (unequal processing of signals), we explore whether generative artificial intelligence (AI) improves "signal conversion efficiency"(SCE), the effective utilization of the information endowment. We engineer an AI agent that uses information from blockchain-based crowdfunding campaigns to generate investment recommendations and compare them with recommendations by human experts and a human-built recommendation algorithm. AI advice outperforms other signals statistically and economically: Including it as a signal for capital formation raises R-squared by up to 35%, and a one-standard-deviation increase in AI recommendation strength is associated with 48% higher funding amounts. Explainable AI using Shapley values corroborates AI's predictive power. Consistent with our SCE conjecture, AI-human advice congruence is positively associated with capital formation, while the direction of divergence matters: campaigns rated more favorably by humans than by AI raise less funding than those for which AI is more optimistic. The AI agent's predictive performance increases across generations of large language models evaluated on the same class of historical disclosures. Our findings highlight information utilization as a central mechanism through which generative AI supports investment decisions in entrepreneurial finance.

Keywords: Artificial intelligence (AI), large language model (LLM), machine learning (ML), entrepreneurial finance, blockchain-based crowdfunding, initial coin offering (ICO)
JEL Codes: C45, G24, G32, K22, L26

1 Introduction

The problem of asymmetric information is ubiquitous in entrepreneurial finance. Entrepreneurs signal venture quality to reduce informational asymmetries for investors, and investors gather and evaluate available signals to allocate capital. Evidence of persistent capital misallocation in entrepreneurial finance markets suggests that the problem of asymmetric information is salient across all market segments, including venture capital (e.g., Dimov and Clercq, 2006; Ungerer et al., 2021), crowdfunding (e.g., Hornuf and Schwienbacher, 2018; Kukk, 2022), and token offerings (e.g., Drobetz et al., 2025; Momtaz, 2024). We argue that artificial intelligence (AI) can improve market efficiency by reducing information asymmetries in a novel way. Specifically, we draw on a distinction between information-endowment asymmetry and information-utilization asymmetry. *Information-endowment asymmetry* refers to differences in what information is available to or possessed by market participants prior to any information processing; that is, it describes unequal access to signals, whereas *information-utilization asymmetry* refers to differences in how effectively market participants collect and process a given information endowment into their beliefs and decisions; that is, it describes unequal processing of signals. Generative AI can improve information utilization by relaxing (humans') bounded rationality constraints (e.g., by collecting more information at lower cost and processing high-dimensional information more efficiently and plausibly in the absence of affective, behavioral, and cognitive biases), raising what we term "signal conversion efficiency" (henceforth, SCE).

We study the role of AI for SCE in blockchain-based crowdfunding, a venture financing mechanism through which entrepreneurs sell cryptocurrency tokens to investors without an intermediary (Bellavitis et al., 2021; Fisch, 2019). The context offers a suited empirical environment because token offerings are disintermediated, decentralized markets on public blockchains, implying that the information endowment is relatively homogeneous across market participants, while differences in capital formation have to stem predominantly from heterogeneous information utilization (Cumming et al., 2025; Momtaz, 2024). At the empirical core of our study is an engineered AI agent that evaluates ICO whitepapers and provides investment recommendations for each token offering. We benchmark the AI recommendation system with recommendations by human experts and a human-built algorithmic recommendation framework to address the following research questions:

i. Can AI utilize more information than human intelligence in predicting capital formation; and

ii. Can the congruence and divergence between AI and human AI reveal differences in information utilization and help us understand AI's potential to correct human biases in investment recommendations?

These questions are increasingly important. Recent work shows that the information utilization ability for human investors depends on the type of data the entrepreneurs present (Yang et al., 2025b), while large langugae models(LLMs) can help predict startup outcomes from unstructured venture information(Kleinert & Urbig, 2026; Maarouf et al., 2025; Vismara et al., 2026) The issue we are facing is not simply whether AI can help evaluate ventures, but whether its judgement contains information that human investors fail to utilize, and what the relationship between AI and human evaluations reveals about capital formation.

Our study is motivated, both theoretically and empirically, by our conjecture of a relative advantage in SCE of artificial vis-à-vis human intelligence. We develop a set of four tesable hypotheses, which, in totality, could suport our SCE coinjecture. Our first hypothesis, the ***Artificial Intelligence Hypothesis*** (henceforth, ***AIH***) submits that AI should have *economically* meaningful predictive power for blockchain-based crowdfunding outcomes. Against the background of distinguishing between information-endowment and information-utilization asymmetry, holding the public information endowment fixed, AI should raise SCE and thereby reduce utilization asymmetry, yielding economically meaningful predictive power. This motivation comes from bounded rationality and rational inattention. In highly uncertain and novel situations, investors should optimally ignore weak but diagnostic cues because attention and computation are costly (March & Simon, 1958; Simon, 1947; Sims, 2003). Token markets can accentuate these restrictions given that disclosure of the information is voluntary and unaudited, code and governance are technically complex, and trading is fragmented (Bellavitis et al., 2021; Fisch, 2019; Fisch & Momtaz, 2020; Gandal et al., 2018; Hornuf et al., 2022; Howell et al., 2019; Makarov & Schoar, 2020; Momtaz, 2021b), therefore, market prices can only partially reflect the quality of entrepreneurial ventures, while endogenous information disclosure hinders the inference (of project quality).(Diamond, 1991; Grossman & Stiglitz, 1980; Momtaz, 2021b). By automating search, structuring unstructured inputs (e.g., whitepapers (Mansouri & Momtaz, 2022), repository cadence and dependency risks (Bellavitis & Momtaz, 2025), and microstructure anomalies that degrade learning from prices (Dhawan & Putniņš, 2022; Florysiak & Schandlbauer, 2022; Hamrick et al., 2021; Momtaz, 2021b; Thewissen et al., 2022)), and generating calibrated inferences at near-zero marginal cost, AI increases SCE so that more of the same endowment is used for more precise investment

decisions.

Our second hypothesis, the ***Recommendation Dominance Hypothesis*** (***RDH***), posits that, on the *same* information endowment, AI may *statistically* dominate human advice because AI achieves a higher SCE than boundedly rational evaluators. This motivation is different from the ***AIH***, which establishes *existence* of predictive power from improved utilization, while the ***RDH*** predicts *comparative* efficiency stems from the limitations of heuristic reasoning and the advantages of high-dimensional learning. In the context of crowdfunding, humans tend to rely on sparse heuristic strategies and are susceptible to emotional and cognitive biases.(Colombo et al., 2022; Huang et al., 2022; Momtaz, 2021a); AI flexibly approximates nonlinear interactions among team, tokenomics, code quality, governance, and market regimes, potentially yielding lower error and better calibration (Boot et al., 2021; Cao et al., 2024; Garleanu & Pedersen, 2018). Meoli and Vismara (2022) show that non-linear models have superior power in predicting the success of token offerings. Therefore, we anticipate that artificial intelligence (AI) can provide incremental value beyond human-generated ratings—even when both are constructed from the same underlying signals, and Shapley values should attribute greater global importance to AI-derived signals, which would support the view of dominance in utilization rather than differences in endowment.

A complementary rationale concerns how human and AI judgments relate, and what that relation reveals about bias and information. Given a fixed public information endowment, when independent data processors (i.e., AI and humans) update on partially overlapping information, greater recommendation congruence indicates lower information-utilization asymmetry (i.e., both sides have converted more of the same endowment into usable beliefs) and signals that both have filtered idiosyncratic noise and distilled a common, informative core (March & Simon, 1958; Simon, 1947). In blockchain-based crowdfunding where narratives are vivid, feedback is delayed, and strategic disclosure is rife, such congruence plausibly indexes higher SCE. Behavioral evidence further indicates that cryptocurrency token investors overweight affect, charisma, and other salient but noisy cues (Colombo et al., 2022; Huang et al., 2022; Momtaz, 2021a, 2021b). Properly designed AI leaves endowments unchanged but improves utilization. Human-AI congruence may, hence, reflect cognitively diversified collectives that better separate fundamentals from hype (Logg et al., 2019; Lorenz et al., 2011; March, 1991). Following this logic is our ***Recommendation Congruence Hypothesis*** (henceforth, ***RCH***) that greater AI-human advice congruence, which reflects jointly higher SCE on a shared endowment, should correlate with capital formation in blockchain-based crowdfunding because the shared component reflects genuine quality rather than sentiment.

Interestingly, the opposite scenario, that is, AI-human advice divergence, leads to nuanced predictions. In particular, our final hypothesis predicts that the *direction* of AI-human disagreement is informative because it reveals asymmetries in *utilization* rather than access. Humans are often overoptimistic, especially in novel, noisy markets like blockchain-based crowdfunding (Barber & Odean, 2001; Hornuf et al., 2025; Kahneman & Tversky, 1979). AI "optimism," by contrast, is utilization-driven and, therefore, epistemic, i.e., a byproduct of higher SCE. Hence, when AI expresses stronger conviction than humans, the wedge plausibly reflects lower information-utilization asymmetry in AI's favor, suggesting that AI has extracted hard signals that humans underweight, so ex-post valuation of the venture should be higher. In contrast, when humans are more favorable of a token offering than AI, the wedge more likely reflects overutilization of noisy, sentiment-laden cues or underprocessing of complex traces by humans, so ex-post valuation should be lower (Brynjolfsson et al., 2025; Cao et al., 2024; Sims, 2003), a prediction that we label the ***Bias Correction Hypothesis*** (henceforth, ***BCH***) because AI-human divergence can be implemented as a corrective device to prevent suboptimal human decision-making.

Empirically, we take an engineering approach for real-world applicability (see, for a current debate, Eesley and Gerber, 2025) by designing an "AI investment advisor." Concretely, we operationalize a large language model as a role-conditioned entrepreneurial-finance and blockchain-technology expert with a clear mandate and structured evaluation protocol, transforming a general model into a consistent, scalable investment screener. At the core is a standardized prompt that defines (i) the analyst persona and task, and (ii) a bounded-rationality constraint requiring the agent to rely exclusively on each project's material that accompanies the campaigns to ensure a controlled and replicable information set across tokens and to avoid the memorization problem associated with applied AI research (Lopez-Lira et al., 2025). The agent parses unstructured disclosures and outputs structured, multi-criteria assessments drawn from established venture-screening dimensions (i.e, market opportunity, technological innovation, problem-solution fit, token economics, and team execution), and produces econometrically usable features at scale. This design moves beyond dictionary counts or generic topic models by simulating expert cognition, enabling head-to-head comparisons with human experts and human-built linear models, while maintaining transparency and replicability. Numerous sanity checks support the real-world relevance of our engineered AI investment advisor. Importantly, we obtain two human-based investment recommendations for each blockchain-based crowdfunding campaign (i.e., human experts and a human-built algorithmic recommendation framework) and show that all three advice systems are strongly correlated. It is important to clarify that our study represents

an engineering-rather-than-identification approach; i.e., our AI agent is engineered to demonstrate superior-to-human investment performance and, as such, all design choices are plausible as long as they do not violate the AI agent's objectivity premise (e.g., through threats from memorization and similar concerns, which we address), see Section 3.2.2 for a more detailed discussion.

The results support our hypotheses. Examining the largest sample of blockchain-based crowdfunding campaigns to date, we find that our AI investment advisor's recommendations are strongly correlated with the amount that tokenized ventures raise through blockchain-based crowdfunding, supporting our ***AIH***. Economically, the recommendation is strongly significant, with a standard deviation increase in the strength of the AI investment recommendation being associated with a higher valuation of 48.9% of the fundraising firm. Consistent with the ***RDH***, we find that a regular OLS regression of the funding amount on a set of known regressors from prior literature (e.g., Bellavitis et al., 2022; Cumming et al., 2025; Drobetz et al., 2025; Fisch, 2019; Mansouri and Momtaz, 2022; Momtaz, 2021b; Xia et al., 2023) yields an adjusted R-squared that increases by up to 35% if the regression conditions on the AI investment advisor's recommendation. Using explainable AI (xAI), we show that among all possible covariates and even non-linear combinations in random forest and extreme gradient boosting models, the AI investment advisor's recommendation always yields the highest Shapley value, suggesting that no other "signal" (including human experts) predicts success in blockchain-based crowdfunding as accurately as our AI investment advisor. In line with our ***RCH***, blockchain-based crowdfunding campaigns for which the AI investment advisor, human experts, and human-built linear prediction models deliver relatively congruent recommendations perform best. Yet, if there is disagreement between the AI and human experts or human-built linear prediction models, blockchain-based crowdfunding campaigns outperform when the the AI was more bullish than the human-based recommendation and underperform when human-based recommendations were more bullish than AI, supporting our ***BCH***. These results are robust in several sensitivity checks.

Recent studies sharpen the contribution of our paper. Kleinert and Urbig (2026) examine LLM evaluations of equity-crowdfunding pitches and compare those evaluations with fundraising and survival outcomes. Csaszar et al. (2026) show that LLMs can anticipate fundraisin success in a prospective crowdfunding tournament. Vismara et al. (2026) deploy an LLM agent to organize and categorize venture capital deal flow, Maarouf et al. (2025) propose a new fused large language model to predict startup success, and Bae et al. (2026) develop a multi-agent framework for startup screening. In the token offerings, Kim (2024) use AI to identify informative content in ICO whitepapers. These works show that AI can help evaluate ventures and

extract information from entrepreneurial disclosures. Our study extends this by explaining the role of *information utilization* through SCE, benchmarking a structured AI investment evaluation against both human experts and human-built linear models, and examining if the congruence and the *direction* of divergence contain additional information about capital formation. Apart from that, we compare evaluations by successive LLM generations using the same whitepaper inputs, holding the information endowment fixed, we link the changes in model capability to the SCE mechanism as well. This three-way comparison distinguish our setting from studies that focused mainly on venture prediction or deal-flow efficiency. (Giannetti & Mavillonio, 2026; Hu et al., 2026)

Our study makes three main contributions. First, we extend entrepreneurial finance research by distinguishing information-endowment from information-utilization asymmetry and conceptualizing SCE as an important dimension of investment screening. Second, we show that the investment recommendation from AI is strongly associated with capital formation and statistically more informative than established human-based recommendation benchmarks. Third, we show that the relationship between AI and human judgments also matters: Their congruence and directional divergence reveal economically meaningful differences in fundraising outcomes, underscoring the potential of AI as a corrective device in human investment decisions. Taken together, these findings place SCE at the center of the emerging literature on generative AI in entrepreneurial finance,

The remainder is organized as follows. Section 2 provides some background on AI and human decision-making and develops testable hypotheses. Section 3 describes our data and the econometric approach. Section 4 presents our results, and section 5 discusses theoretical contributions, practical implications, limitations, and avenues for future research, and concludes the paper.

2 Theory and hypotheses

2.1 Artificial intelligence and investment decision-making in token offerings

AI enhances investment decision-making in token offerings through information utilization. Specifically, AI may increase SCE, the share of available signals converted into decision-relevant certainty about venture quality. It relaxes classic cognitive bounds by cheaply ingesting unstructured traces and returning aggregated signals (March & Simon, 1958; Simon, 1947), lowers the marginal cost of acquiring, parsing, and integrating cues so that weak information may enter the aggregated signal (Sims, 2003), and mitigates predictable

affective, behavioral, and cognitive errors (Kahneman & Tversky, 1979; Logg et al., 2019) often observed in token offerings (Colombo et al., 2022; Hornuf et al., 2025; Momtaz, 2021a, 2021b). Recent studies show that AI extends human information-processing capacity by automating search and inference (Shrestha et al., 2019) and hybrid human-AI systems may reduce variance in strategic judgments and increase calibration accuracy (Logg et al., 2019). Operationally, AI can structure and standardize publicly available text, code, and on-chain data, transforming them into comparable signals that can be screened and monitored. This represents an SCE mechanism that raises utilization without changing access. Evidence consistently shows that AI may generate measurable informational content that predicts firm performance (Mansouri & Momtaz, 2022; Yang et al., 2025a). In token markets, AI further improves search and matching by ranking opportunities, routing attention, and personalizing exploration (Demsetz, 1969; Momtaz, 2024; Rubinstein & Wolinsky, 1987). Convergent evidence suggests that artificial intelligence can enhance productivity and improve the quality of decision-making in the complex, information-rich tasks central to venture capital assessment. (Brynjolfsson et al., 2025; Korinek, 2023; Maarouf et al., 2025; Mansouri & Momtaz, 2022; Momtaz, 2021b; Shrestha et al., 2019; Yang et al., 2025a).

Decision-making in entrepreneurial finance is difficult because information is sparse, noisy, and strategically supplied (Drover et al., 2017). Early ventures lack long track records; assets are intangible; technologies are nascent; and disclosure is endogenous, generating classic adverse-selection and moral-hazard concerns (Colombo et al., 2019). Token offerings magnify these frictions (Cumming et al., 2025; Fisch, 2019). Issuers often have yet to generate any revenue; white papers are published on a voluntary basis and lack standardized criteria; founders may use pseudonyms; and regulatory enforcement involves multiple jurisdictions and tends to be weak. (Fisch, 2019; Hornuf et al., 2022; Howell et al., 2019; Momtaz, 2020; Zetzsche et al., 2017). Beyond disclosure, technical signals are hard to parse for non-experts: repository activity, contract bytecode, and protocol governance create information that is rich but costly to interpret, while usage data (e.g., holder concentration, exchange dependence, validator dynamics) arrive asynchronously and interact with investor composition (Cumming et al., 2025). At the microstructure level, multi-venue trading fragments liquidity and invites wash trading, spoofing, and "pump-and-dump" cycles that distort short-run prices and contaminate learning signals (Cong et al., 2019; Dhawan & Putniņš, 2022; Gandal et al., 2018; Hamrick et al., 2021; Makarov & Schoar, 2020). The investor base consists primarily of retail investors scattered across the globe; their attention is fragmented, making coordination difficult. Furthermore, excessive weight is often placed on signals that—while conspicuous—lack any substantive value (e.g., founder affect,

attractiveness, and pitch production quality) (Colombo et al., 2022; Huang et al., 2022; Momtaz, 2021a).

AI may address these constraints by raising information utilization via higher SCE because it lowers the cost of attention and inference and generate comparable information aggregates that are easy to utilize by boundedly rational humans. Kim (2024) show that AI can identify informational content in ICO disclosures. On the text side, AI can read thousands of whitepapers across languages and formats in a short time period to compute composite scores that forecast fundraising, returns, and survival in token offerings (Florysiak & Schandlbauer, 2022; Momtaz, 2021b; Thewissen et al., 2022). On the technological side, AI can analyze repositories and smart-contract bytecode to quantify update cadence and operational vulnerability, converting technical depth into easily-digestable information for screening early projects. On-chain graphs use AI to analyze raw ledger data to gauge holder concentration, exchange dependence, and governance activity, which are related to liquidity and post-campaign performance (Cumming et al., 2025). Vision models quantify founder affect and presentation cues so investors de-bias appearance-driven assessments (Colombo et al., 2022; Huang et al., 2022; Momtaz, 2021a). For market integrity, sequence models flag abnormal order-flow regimes (e.g., wash trading, spoofing, pump-and-dump) (Cong et al., 2019; Dhawan & Putniņš, 2022; Gandal et al., 2018; Hamrick et al., 2021; Makarov & Schoar, 2020). Finally, AI scales search and triage by ranking projects by predicted informativeness; thus, routing scarce human attention where marginal learning is highest. Recent work shows the promise of multimodal AI for entrepreneurship research and decision support, reinforcing the external validity of these mechanisms in venture contexts (Kyprianou et al., 2024; Shepherd & Majchrzak, 2022; Shrestha & He, 2025). In sum, by improving utilization of a given information endowment via better-than-human SCE, AI should predict capital formation in token offerings.

Hypothesis 1: *AI predicts investment outcomes of cryptocurrency token offerings.*
(***Artificial Intelligence Hypothesis***, henceforth "***AIH***")

2.2 Artificial versus human intelligence in cryptocurrency token investment decisions

AI improves predictive power for investment outcomes in entrepreneurial finance markets because AI raises information utilization, increasing SCE relative to humans on the same public information endowment. Drawing on the aforementioned information-processing and rational-inattention perspectives, human investors face strict limits on attention and memory, forcing them to ignore many weak but informative cues

when evaluating token offerings, while AI can lower the marginal cost of acquiring, parsing, and integrating high-dimensional data (e.g., whitepapers, smart-contract code, wallet graphs, and social signals), thereby expanding the feasible attention set. Consequently, AI captures a larger share of the variance embedded in token-level information, thus producing higher explanatory power in linear prediction models. Moreover, in costly-information environments where disclosure is voluntary and prices are only partially revealing, algorithmic utilization of public traces into "hard" signals adds incremental insight beyond human judgments. AI can efficiently transform unstructured inputs into structured features—features that may capture residual variance overlooked by investors with bounded rationality. (Boot et al., 2021).

In token markets characterized by high nonlinearity, as explained below, AI's utilization advantage extends beyond linear fit to the relative importance of explanatory components. Entrepreneurial finance outcomes emerge from complex interactions(e.g., team composition, tokenomics, code quality, and market regimes) that generate threshold and complementarity effects. Human-derived heuristics and linear filters are inadequate for fitting these structures, whereas AI might flexibly approximate non-linear relationships. Therefore, compared to single indicators constructed by humans, features generated by AI offer a greater marginal contribution to model performance. This advantage becomes even more pronounced in environments characterized by information opacity—settings where behavioral biases, such as affect-based assessments and the "attractiveness heuristic," often interfere with human reasoning processes. (Colombo et al., 2022; Huang et al., 2022; Momtaz, 2021a). Algorithmic representations de-bias and reweight signals, shifting explanatory mass from sentiment to economically meaningful information. Hence, in nonlinear specifications, the Shapley values assigned to AI-based signals are expected to exceed those of alternative predictors, consistent with superior superior utilization-based variance attribution in flexible model spaces (Cao et al., 2024).

In token offerings, the predictive advantages offered by AI are further amplified by the existing institutional landscape and information environment. The ICO market is characterized by a high degree of segmentation and fragmentation, and is rife with market manipulation, search frictions, and strategic information disclosure. (Bellavitis et al., 2021, 2022; Csaszar et al., 2026; Fisch, 2019; Howell et al., 2019; Kleinert & Urbig, 2026; Lyandres et al., 2022; Maarouf et al., 2025; Momtaz, 2020). Human agents, constrained by cognitive bandwidth, screen only a narrow subset of ventures and remain susceptible to expert-rating herding and misinformation (Barth et al., 2023; Vismara, 2018). AI agents, by contrast, efficiently scale breadth, detect microstructure anomalies such as wash trading or spoofing, and recover latent informativeness from

biased disclosures (Dhawan & Putniņš, 2022; Florysiak & Schandlbauer, 2022; Hamrick et al., 2021; Hornuf et al., 2022). These capabilities enable AI to extract independent signals that are complementary to—and often play a dominant role alongside—proxy indicators of human attention. Research evidence from entrepreneurship highlights how multimodal AI and machine learning can enhance predictive performance. (Kyprianou et al., 2024; Shepherd & Majchrzak, 2022; Shrestha & He, 2025). Therefore, given comparable information endowments, signals derived from AI should not only account for a greater share of the variance in investment outcomes but also rank among the top contributors in a Shapley-value decomposition.

Hypothesis 2a (linear model): *AI explains more variation than human agents in cryptocurrency token investment outcomes.*

Hypothesis 2b (non-linear model): *AI outperforms all other signals of cryptocurrency token investment outcomes in terms of Shapley values.*

(***Recommendation Dominance Hypotheses***, henceforth "***RDH***")

2.3 Human cognitive bias correction via AI in cryptocurrency token investments

Congruence between human and AI investment advice in token offerings indicates that both decision agents are utilizing a shared public information endowment in similar ways, that is, converting overlapping cues into beliefs with higher SCE and lower information-utilization asymmetry. In token offerings, these cues include, inter alia, (i) whitepaper claims about roadmap milestones and partnerships (Drobetz et al., 2025); (ii) tokenomics (e.g., vesting cliffs, founder/team allocations, lockup horizons) (Fuchs & Momtaz, 2025); (iii) smart-contract audit status and bytecode quality (Knechel et al., 2026); (iv) founder identity signals (pseudonymity, prior projects) (Colombo et al., 2022; Fisch, 2019); and (v) on-chain distributions at the token-generation event (wallet concentration, liquidity provision/locks) (Cumming et al., 2025). Bayesian learning implies that when independent processors update their priors based on partially overlapping information sets, the alignment of their posteriors reveals a highly informative common component alongside the effective filtering of idiosyncratic noise—precisely what we would expect to see in the form of higher SCE given a fixed endowment. (March & Simon, 1958; Simon, 1947). In token offerings, disclosures are unaudited and often strategic in nature; the identities of participants may be obscured, while the costs of analyzing technical artifacts such as code or technical documentation are high. Advice congruence that persists *after* conditioning on tokenomics, audit signals, and on-chain holder structure therefore proxies an

increase in posterior precision about fundamentals rather than hype (Fisch, 2019; Florysiak & Schandlbauer, 2022; Hornuf et al., 2022; Howell et al., 2019; Momtaz, 2021b; Thewissen et al., 2022). By contrast, AI-human divergence in this setting often flags channel-specific noise, e.g., affect-laden narratives and social buzz on the human side (Barth et al., 2023; Momtaz, 2021a) versus model misspecification or data gaps on the AI side. Hence, ventures receiving congruent positive assessments should exhibit superior ex-post valuations because the shared component reflects efficient utilization of the same endowment rather than transient attention.

Behavioral decision theory clarifies why congruence marks behavioral calibration (Kahneman & Tversky, 1979). Retail-heavy token markets amplify systematic biases because investors anchor on social media activity, influencer endorsements, glossy pitch videos, or charismatic founders (Blaseg & Hornuf, 2024; Blaseg & Schwienbacher, 2024; Colombo et al., 2022; Fisch, 2019; Kolbe et al., 2022), and underweight "hard" but harder-to-process cues like audit depth, dependency risk in the codebase, governance design, or exchange dependence. Such tendencies are well documented in entrepreneurial finance and early token markets (Colombo et al., 2022; Florysiak & Schandlbauer, 2022; Huang et al., 2022; Momtaz, 2021a, 2021b). A properly designed AI system can improve resource utilization efficiency through consistent scoring rules tailored to token issuance factors. For example, NLP is used to detect excessive exaggeration or plagiarism in whitepapers; at the same time, by analyzing on-chain characteristics the risk of manipulation and the viability of the project are assessed.(Cong et al., 2019; Cumming et al., 2025; Dhawan & Putniņš, 2022; Gandal et al., 2018; Hamrick et al., 2021; Makarov & Schoar, 2020). When AI and human advice align under these controls, the consensus indicates debiased, higher-SCE processing of a common endowment, where emotional and salience-driven errors have been tempered and both evaluators anchor on economically meaningful ICO attributes. Consistent with entrepreneurship research on multimodal AI decision support under uncertainty, such calibrated congruence is therefore a utilization-based signal of disciplined evaluation (Kyprianou et al., 2024; Shepherd & Majchrzak, 2022; Shrestha & He, 2025).

Collective-intelligence and organizational-learning theories further predict that decision quality rises when heterogeneous agents integrate diverse yet partially independent information (Lorenz et al., 2011; March, 1991). Recent work likewise points to complementarities between human and AI screening (Giannetti & Mavillonio, 2026; Hu et al., 2026). In token offerings, human-AI collaborative combinations are particularly valuable, as their information channels are characterized by structural heterogeneity. For instance, humans excel at contextual interpretation of ecosystem fit (e.g., complementarities with exist-

ing protocols, community legitimacy, founder reputation across chains), while AI excels at systematically evaluating high-dimensional technical traces (e.g., smart-contract risk signatures, unusual wallet clustering, liquidity removal hazards). Congruence across these channels functions as a diagnostic of efficient utilization: both processors have distilled the same endowment toward a common latent quality construct. It also promotes *credible certification* in markets lacking formal gatekeepers (Cumming et al., 2025). Observed alignment between analysts and algorithmic screens increases perceived legitimacy among other investors (Hornuf et al., 2022; Kyprianou et al., 2024; Shepherd & Majchrzak, 2022; Shrestha & He, 2025). In contrast, persistent divergence (e.g., high human enthusiasm alongside red-flag AI scores for holder concentration and copy-pasted code) signals high utilization asymmetry and lower predictive validity. Thus, AI-human advice congruence in token offerings is not mere statistical agreement; rather, it evidences SCE-driven, cognitively diversified consensus that efficiently synthesizes hard and soft information.

Hypothesis 3: *AI-human investment advice Congruence positively correlates with investment outcomes of token offerings.*

(***Recommendation Congruence Hypothesis***, henceforth "***RCH***")

Behavioral decision theory implies that the direction of relative confidence encodes differences in information utilization on a shared endowment. Human investors routinely exhibit miscalibrated confidence (Barber & Odean, 2001; Kahneman & Tversky, 1979). Typical characteristics of an ICO include making inferences based on social media follower counts or the level of oversubscription during a presale, relying excessively on endorsements from influencers, or employing aggressive fully diluted valuations despite having weak lock-up provisions. (Blaseg & Hornuf, 2024; Fuchs & Momtaz, 2025). By contrast, AI forms probabilistic judgments tied to explicit ICO primitives (e.g., vesting schedules, audit strength, code similarity to known exploits, liquidity-lock status, wallet dispersion) and aggregates uncertainty across ensembles; higher AI confidence therefore proxies higher SCE rather than affect. When AI possesses the same information, its confidence level exceeds that of humans, which shows that it has achieved fuller utilization of information, rather than just showing excessive excitement in behavior. This pattern should arise when fundamentals are strong but underappreciated by humans (e.g., conservative token supply, credible vesting, audited modular code, early governance participation, healthy exchange diversification), implying superior post-offering outcomes.

Conversely, *human-relative overconfidence* typically reflects utilization asymmetry in the opposite direction: salient but low-diagnostic cues dominate processing (Sims, 2003). For instance, high conviction driven by salient but low-diagnostic cues (e.g., founder charisma, slick videos, rapid follower growth driven by bots, or roadmap hype unbacked by repository activity; see Colombo et al., 2022; Momtaz, 2021a). Rational inattention predicts that when attention is costly, individuals will selectively assign excessive weight to certain cues (Sims, 2003). AI frequently down-weights such signals when they are inconsistent with "hard" traces (e.g., shallow audits, concentrated top wallets, liquidity unlocks coincident with marketing spikes, or bytecode reuse from low-quality forks). In these cases, greater human than AI confidence signals under-informed certainty and should predict inferior outcomes. Evidence from entrepreneurship shows that human-AI teaming improves calibration and reduces variance on complex, information-rich tasks, such as screening and monitoring in token offerings (Brynjolfsson et al., 2025; Cao et al., 2024; Kyprianou et al., 2024; Shepherd & Majchrzak, 2022; Shrestha & He, 2025).

A complementary mechanism comes from human-AI complementarity and algorithmic governance: algorithms act as utilization disciplinerers by supplying stable, quantitatively grounded feedback in high-uncertainty settings (Brynjolfsson et al., 2025; Cao et al., 2024; Gregory et al., 2025). In token offerings, where investors' self-attribution and social influence biases are amplified by anonymity and hype (Bellavitis et al., 2022; Fisch, 2019; Moro et al., 2024), AI serves as a rational counterweight that recalibrates subjective confidence toward statistical precision. Consequently, the confidence gap can be viewed as a diagnostic metric for a given use case: when confidence is driven by AI, it tends to be data-based and characterized by higher accuracy (i.e., a higher SCE); conversely, when confidence is driven by humans, it is often accompanied by heuristic biases (i.e., a lower SCE).

Hypothesis 4a: *Relative overconfidence of AI vis-à-vis human-based investment advice is positively related to investment outcomes in token offerings.*

Hypothesis 4b: *Relative overconfidence of human-based vis-à-vis AI investment advice is negatively related to investment outcomes in token offerings.*

(***Bias Correction Hypothesis***, henceforth "***BCH***")

3 Data

3.1 Sample

We test the performance of the AI agent's investment recommendations in a sample of initial coin offerings (token offerings) retrieved from the of the *Token Offerings Research Database* (*TORD*),[1] aggregating data from various sources for more than 10,400 startups through March 2025. Variables we extract from the *TORD* include *funding amount*, *rating by human experts*, *team size*, *whitepaper length*, *minimum viable product*, *GitHub*, *whitelist*, *bonus*, *bounty*, *blockchain technical standard*, and the *market phase*. We filter all offerings for which we were able to retrieve historical data from the *Internet Archive*. In particular, we restrict our sample to offerings which had their whitepapers published until 2021; i.e., before LLMs and other generative AI tools became mainstream. Overall, this leads to a final sample of 4,202 token offerings.

3.2 Variable definitions

3.2.1 Dependent variable: campaign valuation

Funding amount, in $ (log). Our dependent variable is capital formation at the funding stage. Consistent with several studies (e.g., Colombo et al., 2022; Fisch, 2019), as a proxy for startup valuation and capital formation in crowdfunding, we operationalize our dependent variable as the logarithmic funding amount in US$ acquired during the token offering.

3.2.2 Independent variable: AI investment recommendations

We engineer an AI investment advisor for blockchain-based crowdfunding campaigns. This approach moves beyond traditional "text-as-data" techniques, which often rely on dictionary-based word counts or topic modeling (e.g., Mansouri and Momtaz, 2022; Momtaz, 2021b). Instead, we construct an AI agent that simulates the cognitive processes of a human expert — venture capital (VC) and blockchain technology analysts in our specific context — to perform a comprehensive evaluation of startup ventures. This method aligns with the recent surge in using LLMs to perform complex, judgment-based tasks previously exclusive to human experts (Jha et al., 2025; Li et al., 2025; Xia et al., 2023).

[1] Available from www.paulmomtaz.com/data/tord.

For the creation of our AI investment advisor, we provide the LLM[2] with a carefully constructed persona, a clear mandate, and a structured evaluation framework, following the state-of-the art in applied computer science (e.g., Park et al., 2023; Wang et al., 2024). This transforms it from a general-purpose text generator into a consistent and scalable tool for investment analysis (e.g., Cao et al., 2024; Yang et al., 2024). A critical feature of our approach is its ability to interpret unstructured text (i.e., ICO whitepapers) and produce structured, multi-faceted data suitable for econometric analysis. The agent evaluates each startup not on a single dimension but across a set of criteria derived from established VC decision-making practice.

We designed the persona-environment for our AI agents in three steps, as follows:[3]

i. **Role definition:** The AI was assigned the role of an expert venture capitalist and blockchain analyst. This instruction frames the entire analysis, prompting the model to leverage its latent knowledge of financial analysis, decentralized business strategy, and blockchain technology assessment.

ii. **Bounded rationality constraint:** A crucial instruction was to base the evaluation exclusively on the information disclosed within the whitepaper. This constraint is methodologically vital, as it ensures a controlled and replicable analysis, preventing the AI from incorporating external information (e.g., subsequent market news, online forums) that would confound the results (the so-called AI memorization problem, see Lopez-Lira et al., 2025). It effectively restricts and standardizes the "information set" available for each decision.

iii. **Distinguishing disclosure from quality:** We also designed the protocol to differentiate between the absence of information and poor quality of disclosed information. If a whitepaper failed to provide adequate details on a specific evaluation criterion, the agent was instructed to assign a `null` value for that dimension. This avoids penalizing a project for a strategic disclosure choice (e.g., a technically-focused paper omitting team details) and ensures that numerical scores reflect the substance of what is actually presented.

Next, we designed the task-environment for our AI agent along five dimensions. It should be noted that our main results also hold without the five dimensions, as shown in Table A.2 in the Appendix. Nevertheless, the five dimensions provide the AI agent with more context-specific structure, which helped engineering

[2]Unless otherwise noted, the investment recommendations used throughout our main analyses are generated by *Deepseek V3.1*. Because the capability of the underlying model is itself of interest and because it offers a clean test of our SCE mechanism, we additionally regenerate the recommendation using earlier and later GPT generations and report these results in Section 4.6 and Table 12.

[3]The detailed prompt-engineered protocol is available from the authors upon request.

superior-than-human AI in the token-based crowdfunding context. The five dimensions are:

1. **Technical innovation:** Assesses the novelty, feasibility, and technical robustness of the proposed technology or protocol.
2. **Problem-solution fit:** Evaluates the clarity and significance of the identified problem and the degree to which the proposed solution effectively addresses it.
3. **Market Opportunity:** Gauges the quality of the market analysis, including the size of the target market, a competitive landscape, and the go-to-market strategy.
4. **Tokenomics:** Analyzes the design of the project's token, including its utility, incentive structures, supply mechanics, and allocation.
5. **Team & execution plan:** Scores the disclosed credentials, relevant experience, and expertise of the founding team, as well as the coherence of the project roadmap and execution strategy.

Each of these sub-scores was quantified on a scale from 0 (i.e., weak) to 10 (i.e., strong). Finally, we constructed our primary independent variable, namely, the *AI agent's investment recommendation* as the average of all non-`null` scores for each campaigning startup. This composite score represents the AI's overall assessment of the venture's quality based on the disclosed information.

To ensure the reliability and consistency of our AI agent's investment recommendations, we required the agent to self-report its confidence in its own recommendation and its perception of how complete the information basis for the recommendation was. In Section 4.5, we explore in post-hoc analyses how confidence and information completeness impact the recommendation's precision in predicting capital formation in crowdfunding campaigns.

For transparency, all our set-up, including prompting, is provided in the Appendix.

3.2.3 Benchmark for independent variable: Human experts and human-built linear model investment recommendations

We use two human-based recommendation benchmarks for our AI agent's investment recommendation.

Rating by human experts. Arithmetic average of all expert ratings for the ICO, as used in previous studies (e.g., Fisch, 2019; Momtaz, 2021b). Ratings range from 0 ("low quality") to 10 ("high quality").

Rating by human-built linear model. Average of all scaled ratings by human-built linear models for the ICO that were provided by crowdfunding platforms on the campaign websites. For more information

on human-built linear models, such as the "benchy" score on *ICObench*, see Barth et al. (2023) for research of these ratings or here for more practical information. Ratings range from 0 ("low quality") to 10 ("high quality").

As an important plausibility check, we explore the relations between our AI agent's, the human experts', and the human-built linear model's investment recommendations in Figure 1. The top-left graph plots the linear relation between the AI's and the human experts' investment recommendations, with an R-squared of 7.3% and a slope of 0.235. The top-right graph plots the linear relation between the AI's and the human-built linear model's investment recommendations, with an R-squared of 3.9% and a slope of 0.184. Finally, the bottom graph plots the linear relation between the human experts' and the human-built linear model's investment recommendations, for comparison, with an R-squared of 80.9% and a slope of 0.845. These plots are reconfirming, as they all indicate strong linear relations between the three independent investment recommendation systems. At the same time, it is noteworthy that, compare to the plot of human experts' and human-built linear model's investment recommendations, the AI investment recommendations are much weaker related to the two human-based recommendation systems. This makes further empirical investigation into the relative merits of the three recommendation systems interesting.

[Place Figure 1 about here]

3.2.4 Control variables: startup and market characteristics

In all models, we use a comprehensive list of control variables encompassing startup firm and market characteristics.

Team size, in # FTEs. The total number of full-time members of the platform.

Whitepaper length, # words. The natural logarithm of the number of words in a token offering's whitepaper.

Minimum viable product. Dummy variable equal to 1 if the platform has a minimum viable product available, and 0 otherwise.

GitHub. Dummy variable equal to 1 if the platform makes its source code available as open-source on GitHub, and 0 otherwise.

Investor protection: whitelist. Dummy variable equal to 1 if the platform implemented a whitelist during its initial offering, and 0 otherwise.

Promotion: bonus. Dummy variable equal to 1 if the platform offers a bonus structure, which typically involves discounted or free tokens if individual wallet addresses invest above and beyond a certain predetermined investment amount, and 0 otherwise.

Promotion: bounty. Dummy variable equal to 1 if the platform offers a bounty program, which rewards individuals (mostly in the form of free tokens) for marketing or other activities, and 0 otherwise.

Technology: ERC20. Dummy variable equal to one if the firm builds upon the Ethereum standard, and zero otherwise.

Bull market. Dummy variable equal to one if the ICO takes place during the bull market phase (January 2017 until January 2018), and zero otherwise.

Bear market. Dummy variable equal to one if the ICO takes place during the bear market phase (February 2018 until January 2019), and zero otherwise.

We also include quarter-year and country fixed effects to control for time-varying and geographical heterogeneity.

All variables are detailed in A.1 in the Appendix.

3.3 Summary statistics

Summary statistics are in Table 1. We report arithmetic mean, standard deviation, minimum, the 25th, 50th, and 75th percentiles, maximum, and the number of observations. Startup firms in our sample raise, on average, $11.26 million (SD = $17.45 million) in their token offerings. The AI agent's investment recommendation on a scale of 0 (weak recommendation) to 10 (strong recommendation) scores, on average, 5.65 (SD = 1.26) and is slightly skewed to the left, with Q1 of 5.0 vis-à-vis Q3 of 6.5 around the median of 6.0, indicating that the AI agent views more projects with skepticism. Interestingly, the investment recommendations by human experts have a similar mean of 6.07 (SD = 1.45), albeit slightly higher than that of the AI agent. The sample statistics indicate that human experts are more optimistic about token offerings given the higher average score. Consistent with this interpretation, the distribution of investment advice by human experts is skewed to the right, with Q1 of 5.2 vis-à-vis Q3 of 7.2 around the median of 6.0, which contrasts with the left-skewed distribution of AI advice. The same is true, albeit to a lesser extent, for the human-built linear model's investment recommendation, with a mean of 5.90 (SD = 1.37) and Q1/Q3 of 4.80/7.00, reflecting some skewness toward more optimistic recommendations.

For our control variables, we report statistics similar to those in previous work on token offerings (e.g.,

Bellavitis et al., 2022; Cumming et al., 2025; Drobetz et al., 2025; Fisch, 2019; Mansouri and Momtaz, 2022; Momtaz, 2021b; Xia et al., 2023. The average startup firm conducting a token offering has 11.42 (SD = 7.14) team members, 8.67 (SD = 0.96) log-number of words in the whitepaper, a minimum viable product in 22% of all cases, makes the software open-source on Github 54% of the time, whitelists investors in 28% of all cases, offers bonus schemes and bounty programs in 9% and 30% of all offerings, respectively, and 87% of all tokens are built on the ERC20 technical standard. 19% of all token offerings are conducted during bull markets, while 49% are conducted during bear markets, with the rest being conducted during neutral market episodes.

[Place Table 1 about here]

Table 2 shows means and standard deviations for the subsample of below-median AI agent's investment recommendation in the first three columns and for the subsample of above-median AI agent's investment recommendation in the second three columns, as well as differences in means for the two subsamples. In line with our main hypothesis, the average funding amount is higher in startups with strong AI investment recommendations, with the difference being statistically significant at the 1% level. In nominal terms, the startup firm valuation at the end of the token offering campaign is, on average, $5.50 million higher in strongly recommended token offerings, which marks, relative to weakly recommended offerings, a 69.4% valuation increase.

We also shed some light on whether there is "selection on observables" in our sample by comparing the means for our control variables in the strongly versus weakly recommended subsamples. Indeed, we find some statistically significant differences between the full sample and the strongly recommended token offerings subsample. For example, startups that received strong recommendations from our AI agent have, on average, more than three additional team members, with the difference being highly statistically significant. Moreover, startups that received strong recommendations from our AI agent have longer whitepapers, and are more likely to have started the campaign with a minimum viable product, open-source blockchain software, better investor protection, and more bonus programs. To be clear, "selection on observables" is not necessarily a problem in our setting; rather, it is even a desired feature: Given that we train our AI agent on some observable characteristics, the AI agent's investment recommendations prove to be meaningful if there are significant differences in means for at least some of these observable covariates.

[Place Table 2 about here]

Pairwise correlations for all variables are in Table 3. For brevity, we only remark that multicollinearity does not seem to be an issue.

[Place Table 3 about here]

4 Results

4.1 AI investment advice and investment performance in token offerings

Figure 2 illustrates the univariate relation between our AI agent's investment recommendation and the investment outcome, as proxied by the natural logarithm of the funding amount. The binned scatter plot (# bins = 100) indicates a linear AI investment recommendation-outcome relation. Given the strength of the unconditional relation, the graphical evidence suggests that the AI agent's investment advice may be prime determinant of startup firm valuations in token offerings. Next, we examine the robustness of this relation in a multivariate model.

[Place Figure 2 about here]

Table 4 shows the main results for the ***AIH***. All models include quarter-year and country fixed effects to absorb both time-related and geographical variation. All reported standard errors are robust. The R^2 in all of our main model is 22.9%, which is comparable to previous studies (e.g., Fisch, 2019).

Our baseline (OLS) regression results are in the first column of Table 4, with log of the funding amount in $ million as the dependent variable and the AI investment advice on a scale from 0 (weak recommendation) to 10 (strong recommendation) as the independent variable. The coefficient of the AI's investment recommendation is 0.316, with a p-value $< 1\%$, suggesting that a one unit increase on the scale from zero to ten increases the average funding amount of $11.26 million by $4.18 million, or 37.2% (=(exp(0.316))-1 $\times$ 100%), to $15.44 million. This strongly supports the ***AIH*** that AI agents are able to generate valuable investment recommendations in the context of token offerings. The coefficients of the control variables are largely consistent with those reported in prior studies (e.g., Bellavitis et al., 2022; Fisch, 2019; Huang et al.,

2020; Momtaz, 2021b). Specifically, we find that the (i) team size and (ii) whitepaper length are significantly positively related to the funding amount, while (iii) open source code has a negative association with the funding amount.

To test the robustness of the AI investment recommendation's effect on campaign valuation, we include human experts' and human-built linear model's investment recommendations in columns 2 and 4 of Table 4, respectively. We also test for significance in interactions between investment recommendations by AI and human experts and AI and the human-built model in columns 3 and 5 of Table 4, respectively. In column 2, adding human experts' investment recommendation yields a statistically significant coefficient of 0.173, while the AI recommendation's coefficient dominates economically at 0.293, with a p-value of less than 1%. Thus, both investment recommendations by our AI agent and human experts correctly predict the startups' valuations in token offerings, with plausibly more usefulness of the AI agent's investment recommendation, given that the marginal effect is almost twice the amount of that of the human rating. In Section 4.2 below, we will revisit the relative importance of AI and human investment advice, and explore in greater detail how both advices can complement or correct one another. For now, note that, even conditioning on human experts' investment advice, a one unit increase on the scale from zero to ten in the AI investment recommendation strength increases the average funding amount of $11.26 million by $3.83 million, or 34.0%, to $15.09 million. The interaction between the investment recommendations by AI and human experts in column 3 yield a similar, yet stronger pattern, with the marginal effects of AI and human experts on the campaign valuation of 0.657 and 0.500, respectively, thus AI advice continues to be economically more significant than human experts' advice. Similarly, when conditioning on the human-built linear model's investment recommendation in columns 4 (as a control variable) and 5 (as an interaction variable), we find that AI advice dominates with marginal effects of 0.313 and 0.556, respectively, while the human-built linear model's investment recommendations being non-significant. Given that the human-built linear model is strongly used in practice and widely available for investors on leading blockchain-based crowdfunding platforms (Barth et al., 2023), our empirical findings are striking. Overall, these additional tests strongly support our ***AIH***.

[Place Table 4 about here]

Our AI agent's investment recommendations represent advice, which the AI agent aggregates from five different assessment dimensions: market opportunity, technical innovation, problem-solution fit, token eco-

nomics, and team execution. We next explore the relative importance of the five recommendation criteria. Table 5 shows the regression results with the disaggregated AI agent's investment criteria for the corresponding model in the first column of Table 4, with log of the funding amount in $ million as the dependent variable. All disaggregated assessments are statistically significant at the 1% level in these models. The market opportunity coefficient is 0.176 (p-value < 0.01), the technical innovation coefficient is 0.144 (p-value < 0.01), the problem-solution fit coefficient is 0.157 (p-value < 0.01), the token economics coefficient is 0.111 (p-value < 0.01), and the team execution coefficient is 0.168 (p-value < 0.01).[4] In addition to the marginal effects, individual contributions of these five assessment criteria to the model fit as proxied for by the R-squared is insightful. Relative to the lowest model fit for the problem-solution criterion in the third column of Table 5, which explains 22.5 percentage points of the variation in the funding amount, the incremental R-squared added is 2.2% for technical innovation, 4.0% for token economics, 5.3% for market opportunity, and 14.7% for team execution. Thus, the most important dimension that the AI agent correctly understands is the quality of the founding team.

[Place Table 5 about here]

4.2 AI versus human advice: relative importance in explaining campaign valuation

Having established a strongly significant relation between our AI agent's investment advice and the investment outcome in token offerings in Section 4, a natural next question concerns the relative performance of our AI agent vis-à-vis humans and human-built recommendation models. To this end, we compare our AI agent's investment recommendation with that of human experts and a human-built linear model. One dimension to compare the relative importance of the investment recommendation across the three sources is the variance these recommendation can explain in the proxies for the investment outcome. Table 6 compares the R-squared and the adjusted R-squared for our AI agent, human experts, and the human-built linear model for the funding amount regressions. As a preview, the AI agent's investment advice outperforms that of humans and human-built linear models consistently and significantly.

The R-squared in the regression of the funding amount is 22.93% (adj. R-squared = 18.69%) for the AI agent, which compares to 19.75% (adj. R-squared = 15.43%) and 18.26% (adj. R-squared = 13.87%) for

[4]Note that a regression that tests all five criteria simultaneously is statistically infeasible due to multicollinearity.

humans and human-built models, respectively. Therefore, AI agent's investment advice increases R-squared relative to humans and human-built models by 16.09% and 25.54%, respectively.

[Place Table 6 about here]

The above evidence is "in sample," while insights for real-world investment advice needs to stand ground "out-of-sample." To that end, we train and test numerous non-linear models in Table 7. We train random forest and extreme gradient boosting (XGBoost) models to predict the log-funding amount. We use grid search cross-validation, which systematically tests a predefined range of parameter values to identify the combination that yields the best performance on a held-out validation set. The primary reason for this approach is to avoid overfitting, which ensures our models generalize well to new, unseen data — a critical requirement for real-world out-of-sample prediction. By using a 70/30 split for training and testing, and then performing cross-validation on the training set, we select the most robust model configuration. This rigorous training methodology ensures that the reported out-of-sample R-squared values are a reliable measure of the models' predictive power. For brevity, the non-linear models in Table 7 achieve significant out-of-sample R-squared. Random forest and extreme gradient boosting explain 16.83% and 15.66% of the variance in the funding amount of token offerings out of sample. Thus, our models have significant real-world relevance.

[Place Table 7 about here]

An advantage of non-linear modeling is the availability of explainable AI (xAI) tools to explore why non-linear models have real-world practical relevance; namely, we can calculate Shapley (SHAP) values for the models, which quantify the relative importance of each feature (variable) in the non-linear models. Shapley values are in Table 8. In Panel A, we report Shapley values and the difference in normalized Shapley values between the AI agent's advice and that by human experts and the human-built linear model for the random forest modeling of the log-funding amount in Table 7. The AI agent's Shapley value is 0.230 and, hence, dramatically higher than those of human experts' advice (0.106) and the human-built linear model's advice (0.087). These Shapley values are calculated by estimating the random forest model for the log-funding amount and controls, corresponding to the linear model specification in the first column of Table 4.

Panel B presents Shapley values for random forest and XGBoost models in the three columns, respectively, from an empirical model that predicts the log-funding amount with the AI recommendation, the human experts' recommendation, and controls. Only showing the AI and human experts' recommendations,

we observe that the AI recommendation mean-ranks highest, while the human experts' recommendation mean-ranks fifth. Thus, the human experts' recommendation has limited relevance for real-world predictions of blockchain-based crowdfunding campaign valuations. With normalized Shapley values of 0.231 and 0.068 for the AI and human experts, respectively, the dominance in recommendation relevance is evident.

Similarly, Panel C shows Shapley values from an empirical model that predicts the log-funding amount with the AI recommendation, the human-built linear model's recommendation, and controls. Only showing the AI and human-built linear model's recommendations, we observe that the AI recommendation mean-ranks highest, while the human-built linear model's recommendation mean-ranks sixth. Thus, the human-built linear model's recommendation has limited relevance for real-world predictions of blockchain-based crowdfunding campaign valuations. With normalized Shapley values of 0.268 and 0.055 for the AI and the human-built linear model, respectively, the dominance in recommendation relevance is also evident.

In Panel D, we mean-rank all Shapley values across the two non-linear modeling techniques (random forest and XGBoost) obtained from modeling all features as in the first column of Table 4 and the three investment recommendations (AI agent, human experts, human-built linear model) simultaneously. The ranking shows that our AI agent's advice is the most important predictor for investment outcomes in token offerings, followed by human experts' advice in the fifth rank, and human-built linear model's advice in the eighth rank.

Taken together, the evidence in this Section indicates that our AI agent is able to predict investment outcomes not just in our sample but also out-of-sample for real-world relevance and is the most important feature in explaining investment outcomes in token offerings. This finding suggests that artificial intelligence might help improve the capital allocation efficiency of entrepreneurial finance markets, thus strongly supporting our ***RDH***.

[Place Table 8 about here]

4.3 AI-human advice congruence

There is a theoretical motivation to test how AI-human advice congruence is related to valuations in blockchain-based crowdfunding campaigns. To test this, we first need to establish a measure of congruence between our AI agent and human experts or the human-built linear model. We define *strong advice congruence* as a dummy equal to one if the absolute value of the standardized difference in investment recommendations is

less than 0.5 times the standard deviation of the standardized difference. Similarly, we define *moderate advice congruence* as a dummy equal to one if the absolute value of the standardized difference in investment recommendations is greater than 0.5 times the standard deviation of the standardized difference and less than or equal to the standard deviation of the standardized difference. The standardized difference is defined as the raw difference between AI agent's investment recommendation and human experts' or the human-built linear model's investment recommendation, normalized by the standard deviation of the distribution of differences in the investment recommendations.

Figure 3 shows that the funding amount, in $ (log) is relatively high if there is *strong advice congruence* among AI and human experts (funding amount = 15.18) or AI and human-built linear model (funding amount = 15.19), as well as if there is *moderate advice congruence* among AI and human experts (funding amount = 15.03) or AI and human-built linear model (funding amount = 15.05).

[Place Figure 3 about here]

Table 9 simultaneously test the relevance of AI agent's and human experts and human-built linear models for subsamples based on *strong advice congruence* and *moderate advice congruence*. Two observations are noteworthy. First, the AI's recommendation is always predicting the funding amount dominantly and correctly. The coefficients are larger when there is stronger advice congruence, especially in the case when we condition on the human-built linear model's advice. Second, compared to the AI, human experts' or the human-built linear model's advice is mixed. For the strong advice congruence sample, the human experts' recommendation is non-significant, while the human-built linear model's advice is significantly negative for strong advice congruence and non-significant for moderate advice congruence. This speaks to the relative prediction precision of our AI agent's investment advice.

[Place Table 9 about here]

4.4 Human bias correction through AI

Given the relevance of our AI agent's investment recommendations, an important practical question pertains to whether AI can be used to correct human biases in investment decision-making in token offerings. Given the aforementioned congruence measures, we define the state of *AI more optimistic than humans / human-built linear model* if the standardized difference is larger than the standard deviation of the distribution

of differences in the investment recommendations. Conversely, we define the state of *humans / human-built linear model more optimistic than AI* if the standardized difference is less than -1 times the standard deviation of the distribution of differences in the investment recommendations.

To make a case for the utility of AI for human bias correction in the investment decision-making in token offerings, it is necessary to establish that (i) AI agent's *optimism* relative to human experts / human-built linear model leads to *better* investment outcomes and (ii) AI agent's *pessimism* relative to human experts / human-built linear model leads to *worse* investment outcomes. Indeed, the graphical evidence in Figure 4 shows that when AI is more optimistic about a token offering than humans or the human-built linear model, the token offering will outperform with log-funding amounts of 15.45 and 15.61, respectively. In contrast, when the AI agent is less optimistic about a token offering than humans or the human-built linear model, the token offering will clearly underperform with log-funding amounts of 14.61 and 14.39, respectively. Thus, second-guessing human investment decision-making with our AI would help improve investment decision-making by avoiding poor investments (Type I error) and not foregoing highly valuable investments (Type II error).

[Place Figure 4 about here]

Table 10 simultaneously test the relevance of AI agent's and human experts and human-built linear models for subsamples based on *AI more optimistic than humans / human-built linear model* and *humans / human-built linear model more optimistic than AI*. The AI agent's recommendation fails to correctly predict the funding amount in the subsample in which *humans / human-built linear model more optimistic than AI*. In the AI versus humans case, the AI agent fails to predict a significant relation. In the AI versus human-built linear model case, both AI and human-built model fail to predict a significant relationship.

[Place Table 10 about here]

4.5 Post-hoc analysis: The role of AI's meta-analytical assessment

As a sanity check, we make use of AI's self-reflecting capability and explore how AI's assessment of (i) how trustworthy its own investment recommendations likely are and (ii) how the quality of the input data likely is impact the strength of the AI recommendation-campaign valuation relation. Panel A of Table 11 shows two regression results of models that predict the log-funding amount on the AI recommendation interacted

with the AI's confidence in its own recommendation in the first column and on the interaction with the AI's assessment of its perceived information completeness for its investment recommendation in the second column. The first interaction (i.e., AI recommendation $\times$ AI's confidence in own recommendation) is significantly positive, with a coefficient of 0.213 and a p-value below 10%, suggesting that more confidence in the AI agent's own recommendation is indeed associated with better investment outcomes. Interestingly, the second interaction (i.e., AI recommendation $\times$ AI's assessment of perceived information completeness for recommendation) is non-significantly positive, with a coefficient of 0.172, suggesting that even incomplete information environments do not lead to significantly worse performing AI investment recommendations. These results are reconfirming because they indicate that the AI agent is able to avoid Type I/II recommendation errors due to poor information environments, and further conditioning on these two dimensions (AI agent's own confidence and perceived information completeness) strengthens the predictive power of our AI agent's investment advice.

Further, Panel B of Table 11 explores the moderating effect of the quality of the input data for the predictive power of AI's investment advice. We investigate the marginal effects of the clarity, logic, technical rigor, and professional presentation of the input information. For brevity, we observe significantly positive interaction effects of these variables with AI recommendation, except for the clarity proxy. The interaction with the logical flow of input information correlates with a coefficient of 0.295, the technical rigor correlates with a coefficient of 0.104, and professional presentations of input information correlate with a coefficient of 0.245. Hence, these effects are also economically strong, suggesting that startups' well-prepared investment information will ultimately affect capital formation in the emerging AI era of crowdfunding. Interestingly, the interaction with clarity of input information yields a non-significant coefficient of 0.204, suggesting that the AI investment recommendation, plausibly more so than humans, is able to process unclear input information to avoid advice biases. This resonates with the information completeness interaction in Panel A of Table 11, indicating that AI investment advisors are able to deal with incomplete or unclear data that feeds into the recommendation system.

[Place Table 11 about here]

4.6 Post-hoc analysis: LLM generations and predictive power

An implication of our SCE conjecture is that, holding the public information endowment fixed, any improvement in an evaluator's capacity to convert that endowment into decision-relevant beliefs should manifest as stronger predictive power. Comparisons across generations of large language models (LLMs) provide a close-to-ideal test of this implication. That is, we can hold constant all signaling dimensions, and the only object that varies is the underlying model's capability. We therefore re-estimate our baseline specification and its human-benchmark counterparts using investment recommendations generated by three GPT generations of increasing capability: GPT-3.5 Turbo, GPT-4.1, and GPT-5.3.

The results in Table 12 show that our agentic AI investment recommendation's predictive power increases with AI capability, as proxied by later LLM generations, suggesting that AI-improved SCE is indeed a main driver behind our findings. In the baseline specification (column 1), the coefficient on the AI recommendation rises monotonically with model generation, from a statistically insignificant 0.112 for GPT-3.5 Turbo, to 0.372 ($p < 0.01$) for GPT-4.1, to 0.383 ($p < 0.01$) for GPT-5.3. The implied economic magnitudes for the two more recent generations are substantial, as a one-point increase in the recommendation is associated with roughly 45% (GPT-4.1) and 47% (GPT-5.3) higher funding, whereas the GPT-3.5 Turbo recommendation carries no significant association with capital formation. Model fit improves in a similar way. Adjusted R-squared in the baseline rises from 14.3% (GPT-3.5 Turbo) to 18.1% (GPT-4.1) to 19.2% (GPT-5.3).

The relative standing of AI and human advice is itself capability-dependent, which sharpens the head-to-head evidence of Section 4.2. When the AI agent is powered by the weakest model (Panel C, column 2), the human experts' rating is the only significant predictor (0.237, $p < 0.01$) while the AI recommendation is insignificant. The ordering reverses entirely for the more capable models, where the AI recommendation dominates (0.344 and 0.359, $p < 0.01$) and the human-expert coefficient is both smaller and stable at 0.177 ($p < 0.01$). In other words, as the model improves, explanatory mass shifts from the human signal to the AI signal, which is exactly the substitution our *Recommendation Dominance Hypothesis* anticipates, now evidently hinging on the AI's conversion capability. A similar pattern holds in the interaction specifications (columns 3 and 5): the AI $\times$ human terms are insignificant for GPT-3.5 Turbo but become significantly negative for GPT-4.1 and GPT-5.3, consistent with the two advice channels acting as partial substitutes once the model is capable enough to extract the shared diagnostic core. Throughout, the human-built linear

model's coefficient remains insignificant, mirroring its limited incremental value in the main analysis.

Taken together, because the information endowment is held fixed across panels, these patterns isolate utilization from access about as cleanly as the setting permits. What changes from Panel C to Panel A is not the signals available to the agent but the efficiency with which they are converted into an assessment; that is, SCE itself. The result thus complements our cross-evaluator evidence (AI versus humans) with within-evaluator evidence (AI versus AI), and it aligns with broader findings that advances in model capability translate into measurable gains in decision quality on complex, information-rich tasks (Brynjolfsson et al., 2025; Cao et al., 2024; Korinek, 2023).

[Place Table 12 about here]

5 Discussion and concluding remarks

5.1 Summary of main findings

This paper explores whether artificial intelligence improves investment decision-making in a crowdfunding market by increasing signal conversion efficiency (SCE), which we define as the proportion of available information that decision-makers convert into accurate venture assessments. We distinguish between information-endowment asymmetry (unequal access to signals) and information-utilization asymmetry (unequal processing of signals). We test four hypotheses related to our overarching conjecture that, given a shared public information endowment in blockchain-based crowdfunding, any AI-versus-human recommendation edge reflects improvements in information utilization via higher SCE.

Our results offer three main findings. First, an engineered "AI investment advisor" delivers economically meaningful and statistically robust predictive power for capital formation in blockchain-based crowdfunding. In baseline OLS, a one-standard-deviation increase in the AI recommendation's strength associates with 48.9% higher funds raised, supporting our ***Artificial Intelligence Hypothesis*** (***AIH***). Adding the AI recommendation to otherwise identical models lifts adjusted R-squared for valuation by up to 35% over human-expert and human-built algorithmic advice. Additionally, in head-to-head regressions, the AI retains significance while human experts' and the human-built linear model's signals attenuate, indicating a ceteris-paribus utilization advantage, consistent with our ***Recommendation Dominance Hypothesis*** (***RDH***). Explainable-AI diagnostics corroborate this finding; that is, across a set of non-linear models, the AI recom-

mendation is the top-ranked predictor by mean Shapley value, indicating that AI processes a larger share of diagnostic variance from the same information endowment. Importantly, out-of-sample tests confirm practical relevance: random forests and XGBoost achieve test-set R-squared of around 0.16 for blockchain-based crowdfunding outcomes. Second, the relation between AI and human advice is informative about utilization. Campaigns with stronger AI-human congruence raise more, with a clear monotonic gradient across congruence bins, in line with our ***Recommendation Congruence Hypothesis*** (***RCH***); that is, agreement signals jointly higher SCE on a common endowment and hence higher decision-relevant certainty about venture quality. Interestingly, when AI and humans disagree, the direction of disagreement is predictive. "AI-optimistic vs. human-skeptical" crowdfunding campaigns outperform, whereas "human-optimistic vs. AI-skeptical" campaigns underperform, supporting the ***Bias Correction Hypothesis*** (***BCH***) and suggesting that AI can serve to correct human biases in investment decisions. This directional asymmetry is exactly what our information-endowment-versus-utilization-asymmetry framework predicts; namely, AI-over-human optimism reflects information-driven precision (higher SCE), while human-over-AI optimism reflects heuristic distortions (lower SCE). Third, the AI's predictive edge scales with model capability. Holding the public information endowment fixed, the strength of the AI recommendation rises monotonically across GPT generations, with model fit improving in a similar way. Because only the underlying model's capacity varies across these specifications, this within-evaluator gradient corroborates our central claim that AI's contribution operates through improved information utilization (higher SCE) rather than expanded access. Overall, our results provide systematic evidence that AI improves investment outcomes in blockchain-based crowdfunding campaigns by converting more of the available information into accurate venture assessments.

5.2 Theoretical contributions and practical implications

Our paper makes at least three contributions to the literatures on entrepreneurial finance, blockchain-based crowdfunding, and AI applications and human-AI complementarity.

First, we contribute to entrepreneurial finance by demonstrating that AI improves outcomes primarily by reducing information-utilization asymmetry via higher SCE on a shared public endowment where markets are informationally thin, strategically opaque, and cognitively demanding. Entrepreneurial finance features severe asymmetric information, endogenous disclosure, and vast, dispersed, multimodal traces (text, code, on-chain graphs, social signals) that challenge boundedly rational investors (Drover et al., 2017; March & Simon, 1958; Simon, 1947). Theory predicts that when information is costly and prices are only par-

tially revealing, technologies that convert unstructured data into "hard" signals raise utilization and should improve screening and monitoring (Diamond, 1991; Grossman & Stiglitz, 1980). AI lowers the marginal cost of attention and inference (Sims, 2003), expands the feasible attention set, and automates search under complexity (Kyprianou et al., 2024; Shrestha & He, 2025). It also increases SCE by scaling discovery of value-relevant patterns from text and code (Yang et al., 2025a) and mitigating utilization-distorting biases (Kahneman & Tversky, 1979; Logg et al., 2019). Our evidence that an engineered "AI investment advisor"can predict economically meaningful outcomes. This finding reveals a generalizable mechanism: the AI converts a larger share of existing information resources into definitive insights regarding the quality of entrepreneurial ventures—insights that serve as valuable inputs for decision-making. This not only complements expert judgment but also helps optimize the allocation of capital. (Brynjolfsson et al., 2025; Korinek, 2023). This mechanism receives support from the cross-generation evidence in Section 4.6. Across models the information endowment is identical by construction, so the monotonic rise in predictive power with model capability cannot reflect differential access to signals; it can only reflect differences in how efficiently a common endowment is converted into venture-quality assessments. The result thus pins the locus of AI's contribution to utilization rather than endowment and suggests that the gains we document should strengthen, not attenuate, as underlying models advance.

Second, we contribute to the literature on blockchain-based crowdfunding by providing evidence that an AI investment advisor predicts funding amounts in token offerings. Prior work documents voluntary disclosure, persuasive language in whitepapers, and behavioral cues that affect fundraising and post-listing outcomes (Fisch, 2019; Florysiak & Schandlbauer, 2022; Hornuf et al., 2022; Howell et al., 2019; Momtaz, 2021b; Thewissen et al., 2022). Founder affect and presentation quality are overweighted when accessible information is underutilized (Colombo et al., 2022; Huang et al., 2022; Momtaz, 2021a). Market microstructure further muddies learning through fragmentation and manipulation (Cong et al., 2019; Gandal et al., 2018; Hamrick et al., 2021; Makarov & Schoar, 2020). Intermediaries and institutional entry help economize on search and certification, but coverage remains uneven (Boot et al., 2021; Cumming et al., 2025; Garleanu & Pedersen, 2018; Momtaz, 2024). Against this backdrop, our AI raises SCE relative to human and human-built baselines, adds predictive power, improves out-of-sample fit, and dominates variance attribution. In doing so, we extend and integrate the disclosure, behavior, and market-integrity strands of the token offering literature by showing that algorithmic utilization of public traces can systematically anticipate capital formation in this decentralized setting (Drobetz et al., 2025).

Third, we inform the human-AI interaction literature by showing when and how AI complements (and disciplines) human judgment in high-uncertainty venture markets. Consistent with hybrid-intelligence and productivity studies, AI improves decision quality on complex information tasks and can close skill gaps (Brynjolfsson et al., 2025; Cao et al., 2024; Shrestha & He, 2025). We have implemented two theory-based integration mechanisms. First, the alignment between AI and humans serves as a diagnostic metric for utilization: specifically, higher alignment between AI and humans regarding shared resources correlates with higher funding success. This aligns with the logic of collective intelligence that consensus among relatively independent processors indicates an efficient extraction of common, fundamental elements, rather than hype. (Lorenz et al., 2011; March, 1991). Second, directional disagreement encodes utilization asymmetries in the sense that AI-optimistic vs. human-skeptical outperforms the reverse, consistent with the view that AI can help correct for overconfidence and selective attention driven by human emotions.(Kahneman & Tversky, 1979; Logg et al., 2019; Sims, 2003). Explainable-AI analyses make these SCE-driven complementarities transparent, identifying the features that drive ensemble gains and enabling accountable human-AI teaming focused on better utilization of the existing information endowment in entrepreneurial finance (Cao et al., 2024).

Our study has several practical implications. For blockchain entrepreneurs, our results suggest that AI-ready disclosure creates value by raising information utilization (higher SCE) on a shared public endowment. Entrepreneurs may craft whitepapers and repositories with machine-parsable, verifiable structure (e.g., clear roadmaps, verifiable milestones, modular code, governance primitives) and publish on-chain telemetry that supports independent algorithmic verification. Doing so enhances investor SCE, improves screening by sophisticated investors, and credibly signals quality to retail participants, lowering the cost of capital. For investors, deploying an "AI investment advisor" as a first-pass utilization engine can expand search breadth, detect manipulation and exaggeration, and surface fundamentals obscured by hype. The largest gains arise from disciplined human-AI teaming; specifically, prioritize opportunities with high AI-human agreement (jointly high SCE), and treat "AI-optimistic vs. human-skeptical" cases as high-learning investigations (utilization gaps to be closed) rather than reflex passes; conversely, be cautious when human conviction exceeds AI confidence without hard corroboration (low-SCE, bias-prone states). For policymakers and regulators, our findings motivate AI-enabled market integrity and disclosure regimes that improve information utilization. Machine-parsable templates for token offerings, standardized code-audit artifacts, and real-time surveillance to flag wash trading or spoofing. Sandboxes that encourage AI-verifiable dis-

closures, coupled with post-issuance transparency (e.g., treasury wallets, governance activity), can raise informational efficiency by increasing SCE while preserving innovation in blockchain-based crowdfunding.

5.3 Limitations and avenues for future research

Our study opens several avenues for future research. First, a limitation of our study is that the engineered AI advisor predicts funding amounts at the offering stage, not post-offering performance. Funding success captures investors' contemporaneous beliefs and platform frictions, but it is an imperfect proxy for realized value creation and risk-adjusted performance after secondary trading begins. Several forces may help explain the gap, such as selection into listing, thin and fragmented liquidity, manipulation and wash trading around listing, and time-varying risk premia that decouple initial demand from subsequent returns. This opens a rich research agenda. First, reverse-engineering BHAR would require outcome re-design and data augmentation: (i) jointly model the probability and timing of listing (to address selection and survivorship), (ii) construct regime-aware BHAR targets (event-time vs. calendar-time, varying risk models), and (iii) de-noise prices using microstructure filters and manipulative-pattern screens. Second, one could explore multi-task learning where a shared representation predicts both funding and BHAR, allowing the model to learn features that generalize from primary to secondary markets. Third, it seems interesting to incorporate causal-machine-learning and instrumented designs (e.g., exchange outages, gas-fee spikes) to separate hype from fundamentals. Fourth, one might add dynamic signals (post-offering code commits, governance proposals, on-chain adoption curves) and model horizon-specific BHAR with survival-aware forecasters. Finally, practitioners might want to calibrate decision-usefulness by simulating portfolio policies that blend human and AI advice under realistic frictions. These steps would extend our contribution from screening capital formation to forecasting investor surplus in secondary markets.

A second limitation concerns external validity beyond blockchain-based crowdfunding. Our setting is transparent given that most crypto entrepreneurs built on public blockchains and raise financing from the public. By contrast, other entrepreneurial finance markets, e.g., venture capital, angel syndicates, and private debt operate behind closed doors, where key signals (e.g., partner memos, customer calls, term sheets, proprietary diligence) are sparse, confidential, and slow to surface. It is therefore not ex ante clear whether an AI investment advisor will provide economically meaningful recommendations in less-public market segments. This opens several research avenues, including (i) adapting the advisor to private markets via privacy-preserving data partnerships that ingest non-public diligence artifacts; (ii) testing domain shift with

transfer and few-shot learning using small, vetted venture capital datasets; (iii) leveraging proxy outcomes (e.g., follow-on financing, up-round likelihood, time-to-exit) when returns are unobserved; and (iv) running field or lab-in-the-field experiments with investors to assess decision support (e.g., speed, calibration, bias reduction) rather than pure prediction.

Finally, a third limitation is equilibrium endogeneity once AI advice diffuses. AI can lower the cost of producing persuasive information, potentially changing tbe relationship between signal quality and funding outcomes Cowgill et al. (2026). Our models learn from historical disclosures and behaviors, implicitly assuming a passive environment. In practice, salient AI features risk becoming targets for optimization (Goodhart's law); i.e., issuers may tailor whitepapers, repositories, or on-chain activity to trigger favorable scores, and exchanges or market makers may adjust microstructure in response to anomaly flags. Such feedback can attenuate out-of-sample validity and even induce new forms of manipulation. Future work should (i) stress-test robustness to adversarial or synthetic inputs, (ii) design detection layers for "model-aware" obfuscation (e.g., stylometry, code plagiarism, or inorganic network growth), and (iii) study policy interventions (e.g., disclosure templates, audit standards) that reduce gaming surfaces. This concern connects to a caution raised by Kleinert and Urbig (2026) who note that opaque reliance on ventures' digital footprints can render LLM evaluations difficult to interpret and may incentivize strategic curation of online presence. While our pre-2021 sampling and bounded-rationality constraint limit this channel in our setting, the broader risk that disclosure adapts to algorithmic evaluation remains an important direction for future work.

5.4 Conclusion

This paper examines whether and how engineered AI can enhance investment decision-making in a crowdfunding market through the lens of information endowment versus information utilization asymmetry and SCE. Using an engineered "AI investment advisor" tailored to blockchain-based crowdfunding, we show three core results. First, AI delivers economically meaningful and statistically robust predictive power for campaign outcomes, improving explanatory fit beyond human experts and human-built linear rules. Second, in head-to-head tests and across nonlinear models, AI-based advice consistently dominates on importance metrics. Third, human-AI relations matter; that is, campaigns with stronger AI-human congruence perform better, and asymmetries in relative confidence are informative (AI-optimistic versus human-skeptical cases outperform, while the reverse underperform). Additionally, we also find that more capable generations of LLMs exhibit stronger associations with fundraising outcomes. Taken together, our evidence indicates that

AI augments judgment by reducing information-utilization asymmetry, thereby raising the quality of capital allocation in token markets, and offers a scalable path to disciplined screening where information is vast, noisy, and strategically supplied. The framework clarifies where AI's gains come from (i.e., better utilization of an existing information endowment) and offers structure for future research on the role of AI in entrepreneurial finance.

References

Bae, J. Y., Malberg, S., Galang, J. A. C., Retterath, A., & Groh, G. (2026). DIALECTIC: A multi-agent system for startup evaluation. *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 5: Industry Track)*, 711–727. https://doi.org/10.18653/v1/2026.eacl-industry.53

Barber, B. M., & Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment*. *The Quarterly Journal of Economics*, *116*(1), 261–292. https://doi.org/10.1162/003355301556400

Barth, A., Laturnus, V., Mansouri, S., & Wagner, A. F. (2023). Conflicted analysts and initial coin offerings. *Management Science*, *69*(11), 6641–6666. https://doi.org/10.1287/mnsc.2021.02928

Bellavitis, C., Cumming, D., & Vanacker, T. (2022). Ban, boom, and echo! entrepreneurship and initial coin offerings. *Entrepreneurship: Theory and Practice*, *46*(5), 1136–1169. https://doi.org/10.1177/1042258720940114

Bellavitis, C., Fisch, C., & Wiklund, J. (2021). A comprehensive review of the global development of initial coin offerings (icos) and their regulation. *Journal of Business Venturing Insights*, *15*. https://doi.org/10.1016/j.jbvi.2020.e00213

Bellavitis, C., & Momtaz, P. P. (2025). Voting governance and value creation in decentralized autonomous organizations (daos). *Journal of Business Venturing Insights*, *23*, e00537. https://doi.org/10.1016/j.jbvi.2025.e00537

Blaseg, D., & Hornuf, L. (2024). Playing the business angel: The impact of well-known business angels on venture performance. *Entrepreneurship: Theory and Practice*, *48*(1), 171–204. https://doi.org/10.1177/10422587231153603

Blaseg, D., & Schwienbacher, A. (2024). Biased calibration: Exacerbating instead of mitigating entrepreneurial overplacement with reference values. *Entrepreneurship: Theory and Practice*, *48*(5), 1131–1159. https://doi.org/10.1177/10422587241227051

Boot, A., Hoffmann, P., Laeven, L., & Ratnovski, L. (2021). Fintech: What's old, what's new? *Journal of Financial Stability*, *53*, 100836. https://doi.org/10.1016/j.jfs.2020.100836

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative ai at work. *Quarterly Journal of Economics*, *140*(2), 889–942. https://doi.org/10.1093/qje/qjae044

Cao, S., Jiang, W., Wang, J., & Yang, B. (2024). From man vs. machine to man + machine: The art and ai of stock analyses. *Journal of Financial Economics*, *160*, 103910. https://doi.org/10.1016/j.jfineco.2024.103910

Colombo, M. G., Fisch, C., Momtaz, P. P., & Vismara, S. (2022). The ceo beauty premium: Founder ceo attractiveness and firm valuation in initial coin offerings. *Strategic Entrepreneurship Journal*, *16*(3), 491–521. https://doi.org/10.1002/sej.1417

Colombo, M. G., Meoli, M., & Vismara, S. (2019). Signaling in science-based ipos: The combined effect of affiliation with prestigious universities, underwriters, and venture capitalists. *Journal of Business Venturing*, *34*(1), 141–177. https://doi.org/10.1016/j.jbusvent.2018.04.009

Cong, L., Li, X., Tang, K., & Yang, Y. (2019). Crypto wash trading. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3530220

Cowgill, B., Hernández-Lagos, P., & Wright, N. L. (2026). Does AI cheapen talk? theory and evidence from global entrepreneurship and hiring. *Management Science*. https://doi.org/10.1287/mnsc.2024.07027

Csaszar, F. A., Peterson, A., & Wilde, D. (2026). The strategic foresight of LLMs: Evidence from a fully prospective venture tournament. https://doi.org/10.2139/ssrn.6166986

Cumming, D., Drobetz, W., Momtaz, P. P., & Schermann, N. (2025). Financing decentralized digital platform growth: The role of crypto funds in blockchain-based startups. *Journal of Business Venturing*, *40*(1), 106450. https://doi.org/10.1016/j.jbusvent.2024.106450

Demsetz, H. (1969). Information and efficiency: Another viewpoint. *The Journal of Law and Economics*, *12*(1), 1–22. Retrieved October 24, 2025, from http://www.jstor.org/stable/724977

Dhawan, A., & Putniņš, T. J. (2022). A new wolf in town? pump-and-dump manipulation in cryptocurrency markets*. *Review of Finance*, *27*(3), 935–975. https://doi.org/10.1093/rof/rfac051

Diamond, D. W. (1991). Monitoring and reputation: The choice between bank loans and directly placed debt. *Journal of Political Economy*, *99*(4), 689–721. Retrieved October 24, 2025, from http://www.jstor.org/stable/2937777

Dimov, D., & Clercq, D. D. (2006). Venture capital investment strategy and portfolio failure rate: A longitudinal study. *Entrepreneurship Theory and Practice*, *30*(2), 207–223. https://doi.org/10.1111/j.1540-6520.2006.00118.x

Drobetz, W., Hornuf, L., Momtaz, P. P., & Schermann, N. (2025). Token-based crowdfunding: Investor choice and the optimal timing of initial coin offerings. *Entrepreneurship Theory and Practice*, *49*(1), 232–282. https://doi.org/10.1177/10422587241261625

Drover, W., Busenitz, L., Matusik, S., Townsend, D., Anglin, A., & Dushnitsky, G. (2017). A review and road map of entrepreneurial equity financing research: Venture capital, corporate venture capital, angel investment, crowdfunding, and accelerators. *Journal of Management*, *43*(6), 1820–1853. https://doi.org/10.1177/0149206317690584

Eesley, C., & Gerber, E. (2025). Crossroads—designing institutions for applied impact: Lessons from engineering for organizational research. *Organization Science*, *36*(5), 2044–2051. https://doi.org/10.1287/orsc.2025.21020

Fisch, C. (2019). Initial coin offerings (icos) to finance new ventures. *Journal of Business Venturing*, *34*(1), 1–22. https://doi.org/10.1016/j.jbusvent.2018.09.007

Fisch, C., & Momtaz, P. P. (2020). Institutional investors and post-ico performance: An empirical analysis of investor returns in initial coin offerings (icos). *Journal of Corporate Finance*, *64*, 101679. https://doi.org/10.1016/j.jcorpfin.2020.101679

Florysiak, D., & Schandlbauer, A. (2022). Experts or charlatans? ico analysts and white paper informativeness. *Journal of Banking & Finance*, *139*, 106476. https://doi.org/10.1016/j.jbankfin.2022.106476

Fuchs, J., & Momtaz, P. P. (2025). Token governance in initial coin offerings: Implications of token retention and resale restrictions for ico success. *Small Business Economics*, *64*(3), 1321–1359. https://doi.org/10.1007/s11187-024-00945-9

Gandal, N., Hamrick, J., Moore, T., & Oberman, T. (2018). Price manipulation in the bitcoin ecosystem. *Journal of Monetary Economics*, *95*, 86–96. https://doi.org/10.1016/j.jmoneco.2017.12.004

Garleanu, N., & Pedersen, L. H. (2018). Efficiently inefficient markets for assets and asset management. *The Journal of Finance*, *73*(4), 1663–1712. https://doi.org/10.1111/jofi.12696

Giannetti, C., & Mavillonio, M. S. (2026). Predicting crowdfunding success: The human–AI complementarity. https://ssrn.com/abstract=6075854

Gregory, R. W., Beck, R., Henfridsson, O., & Yaraghi, N. (2025). Cooperation among strangers: Algorithmic enforcement of reciprocal exchange with blockchain-based smart contracts. *Academy of Management Review*, *50*(4), 859–873. https://doi.org/10.5465/amr.2023.0023

Grossman, S. J., & Stiglitz, J. E. (1980). The impossibility of informationally efficient markets. *The American Economic Review*, *70*(3), 393–408.

Hamrick, J., Rouhi, F., Mukherjee, A., Feder, A., Gandal, N., Moore, T., & Vasek, M. (2021). An examination of the cryptocurrency pump-and-dump ecosystem. *Information Processing & Management*, *58*(4), 102506. https://doi.org/10.1016/j.ipm.2021.102506

Hornuf, L., Kück, T., & Schwienbacher, A. (2022). Initial coin offerings, information disclosure, and fraud. *Small Business Economics*, *58*(4), 1741–1759. https://doi.org/10.1007/s11187-021-00471-y

Hornuf, L., Momtaz, P. P., Nam, R. J., & Yuan, Y. (2025). Cybercrime on the ethereum blockchain. *Journal of Banking & Finance*, *175*, 107419. https://doi.org/10.1016/j.jbankfin.2025.107419

Hornuf, L., & Schwienbacher, A. (2018). Market mechanisms and funding dynamics in equity crowdfunding. *Journal of Corporate Finance*, *50*, 556–574. https://doi.org/10.1016/j.jcorpfin.2017.08.009

Howell, S. T., Niessner, M., & Yermack, D. (2019). Initial coin offerings: Financing growth with cryptocurrency token sales. *The Review of Financial Studies*, *33*(9), 3925–3974. https://doi.org/10.1093/rfs/hhz131

Hu, K., Chen, Z., & Li, X. (2026). Asymmetric dynamic bayesian learning in human–AI collaborative decision making: Evidence from venture capital investment screening. https://ssrn.com/abstract=6581938

Huang, W., Meoli, M., & Vismara, S. (2020). The geography of initial coin offerings. *Small Business Economics*, *55*(1), pp. 77–102. Retrieved October 24, 2025, from https://www.jstor.org/stable/48734337

Huang, W., Vismara, S., & Wei, X. (2022). Confidence and capital raising. *Journal of Corporate Finance*, *77*, 101900. https://doi.org/10.1016/j.jcorpfin.2021.101900

Jha, M., Qian, J., Weber, M., & Yang, B. (2025). Chatgpt and corporate policies. https://arxiv.org/abs/2409.17933

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291. Retrieved October 24, 2025, from http://www.jstor.org/stable/1914185

Kim, J. H. (2024). Catching information bluffs with AI: Evidence from initial coin offerings. https://doi.org/10.2139/ssrn.4951410

Kleinert, S., & Urbig, D. (2026). Startup evaluations with generative artificial intelligence: An exploratory study on early-stage investments and survival predictions by large language models. *Entrepreneurship Theory and Practice*, *0*(0), 10422587261430320. https://doi.org/10.1177/10422587261430320

Knechel, W. R., Maex, S. A., & Park, H. J. (2026). The role of auditor reputation in an emerging audit marketplace: Evidence from decentralized finance (defi). *Management Science*, *72*(6), 5083–5105. https://doi.org/10.1287/mnsc.2023.02245

Kolbe, M., Mansouri, S., & Momtaz, P. P. (2022). Why do video pitches matter in crowdfunding? *Journal of Economics and Business*, *122*, 106081. https://doi.org/10.1016/j.jeconbus.2022.106081

Korinek, A. (2023). Generative ai for economic research: Use cases and implications for economists. *Journal of Economic Literature*, *61*(4), 1281–1317. https://doi.org/10.1257/jel.20231736

Kukk, M.-L. (2022). Predicting business failure after crowdfunding success: Are platforms the unsung heroes? *Journal of Business Venturing Insights*, *17*, e00308. https://doi.org/10.1016/j.jbvi.2022.e00308

Kyprianou, C., Vedula, S., Pu, W., & Fitza, M. (2024). Shifts in national entrepreneurial culture: The promise of linguistic cultural artifacts and machine learning analysis. *Strategic Entrepreneurship Journal*. https://doi.org/10.1002/sej.1519

Li, H., Gao, H., Wu, C., & Vasarhelyi, M. A. (2025). Extracting financial data from unstructured sources: Leveraging large language models. *Journal of Information Systems*, *39*(1), 135–156. https://doi.org/10.2308/ISYS-2023-047

Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, *151*, 90–103. https://doi.org/10.1016/j.obhdp.2018.12.005

Lopez-Lira, A., Tang, Y., & Zhu, M. (2025). The memorization problem: Can we trust llms' economic forecasts? *arXiv preprint arXiv:2504.14765*.

Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, *108*(22), 9020–9025. https://doi.org/10.1073/pnas.1008636108

Lyandres, E., Palazzo, B., & Rabetti, D. (2022). Initial coin offering (ico) success and post-ico performance. *Management Science*, *68*(12), 8658–8679. https://doi.org/10.1287/mnsc.2022.4312

Maarouf, A., Feuerriegel, S., & Pröllochs, N. (2025). A fused large language model for predicting startup success. *European Journal of Operational Research*, *322*(1), 198–214. https://doi.org/10.1016/j.ejor.2024.09.011

Makarov, I., & Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, *135*(2), 293–319. https://doi.org/10.1016/j.jfineco.2019.07.001

Mansouri, S., & Momtaz, P. P. (2022). Financing sustainable entrepreneurship: Esg measurement, valuation, and performance. *Journal of Business Venturing*, *37*(6), 106258. https://doi.org/10.1016/j.jbusvent.2022.106258

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, *2*(1), 71–87. http://www.jstor.org/stable/2634940

March, J. G., & Simon, H. A. (1958). *Organizations*. Wiley.

Meoli, M., & Vismara, S. (2022). Machine-learning forecasting of successful icos. *Journal of Economics and Business*, *121*, 106071. https://doi.org/10.1016/j.jeconbus.2022.106071

Momtaz, P. P. (2020). Initial coin offerings. *PLOS ONE*, *15*(5), 1–30. https://doi.org/10.1371/journal.pone.0233018

Momtaz, P. P. (2021a). Ceo emotions and firm valuation in initial coin offerings: An artificial emotional intelligence approach. *Strategic Management Journal*, *42*(3), 558–578. https://doi.org/10.1002/smj.3235

Momtaz, P. P. (2021b). Entrepreneurial finance and moral hazard: Evidence from token offerings. *Journal of Business Venturing*, *36*(5), 106001. https://doi.org/10.1016/j.jbusvent.2020.106001

Momtaz, P. P. (2024). Decentralized finance (defi) markets for startups: Search frictions, intermediation, and the efficiency of the ico market. *Small Business Economics*, *63*(4), 1415–1447. https://doi.org/10.1007/s11187-024-00886-3

Moro, A., Radić, N., & Truong, V. (2024). To tweet or not to tweet? the determinants of tweeting activity in initial coin offerings. *British Journal of Management*, *35*(1), 243–258. https://doi.org/10.1111/1467-8551.12709

Park, J. S., O'Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. https://arxiv.org/abs/2304.03442

Rubinstein, A., & Wolinsky, A. (1987). Middlemen*. *The Quarterly Journal of Economics*, *102*(3), 581–593. https://doi.org/10.2307/1884218

Shepherd, D. A., & Majchrzak, A. (2022). Machines augmenting entrepreneurs: Opportunities (and threats) at the nexus of artificial intelligence and entrepreneurship. *Journal of Business Venturing*, *37*(4), 106227. https://doi.org/10.1016/j.jbusvent.2022.106227

Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, *61*(4), 66–83. https://doi.org/10.1177/0008125619862257

Shrestha, Y. R., & He, V. F. (2025). Integrating multimodal data and machine learning for entrepreneurship research. *Strategic Entrepreneurship Journal*. https://doi.org/10.1002/sej.1546

Simon, H. A. (1947). *Administrative behavior: A study of decision-making processes in administrative organization*. Macmillan.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, *50*(3), 665–690. https://doi.org/10.1016/S0304-3932(03)00029-1

Thewissen, J., Shrestha, P., Torsin, W., & Pastwa, A. M. (2022). Unpacking the black box of ico white papers: A topic modeling approach. *Journal of Corporate Finance*, *75*, 102225. https://doi.org/10.1016/j.jcorpfin.2022.102225

Ungerer, C., Reuther, K., & Baltes, G. (2021). The lingering living dead phenomenon: Distorting venture survival studies? *Journal of Business Venturing Insights*, *16*, e00250. https://doi.org/10.1016/j.jbvi.2021.e00250

Vismara, S. (2018). Information cascades among investors in equity crowdfunding. *Entrepreneurship Theory and Practice*, *42*(3), 467–497. https://doi.org/10.1111/etap.12261

Vismara, S., Latifi, G., Meinzinger, L., & Pass, A. (2026). Generative ai-powered venture screening: Can large language models help venture capitalists? *International Review of Financial Analysis*, *109*, 104748. https://doi.org/10.1016/j.irfa.2025.104748

Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J. (2024). A survey on large language model based autonomous agents. *Frontiers of Computer Science*, *18*(6). https://doi.org/10.1007/s11704-024-40231-1

Xia, F., Thewissen, J., Huang, W., & Yan, S. (2023). Founder-ceo extraversion and sustainability orientation in initial coin offerings. *Entrepreneurship Theory and Practice*. https://doi.org/10.1177/10422587231211012

Yang, H., Liu, X.-Y., & Wang, C. D. (2025a). Fingpt: Open-source financial large language models. https://arxiv.org/abs/2306.06031

Yang, H., Zhang, B., Wang, N., Guo, C., Zhang, X., Lin, L., Wang, J., Zhou, T., Guan, M., Zhang, R., & Wang, C. D. (2024). Finrobot: An open-source ai agent platform for financial applications using large language models. https://arxiv.org/abs/2405.14767

Yang, J., Xin, J., Zeng, Y., & Liu, P. J. (2025b). Signaling and perceiving on equity crowdfunding decisions—a machine learning approach. *Small Business Economics*, *65*, 315–356. https://doi.org/10.1007/s11187-024-00991-3

Zetzsche, D., Buckley, R., Arner, D., & Föhr, L. (2017). The ico gold rush: It's a scam, it's a bubble, it's a super challenge for regulators. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3072298

Exhibits

Figure 1: Relation of AI, human, and human-built model's investment advice

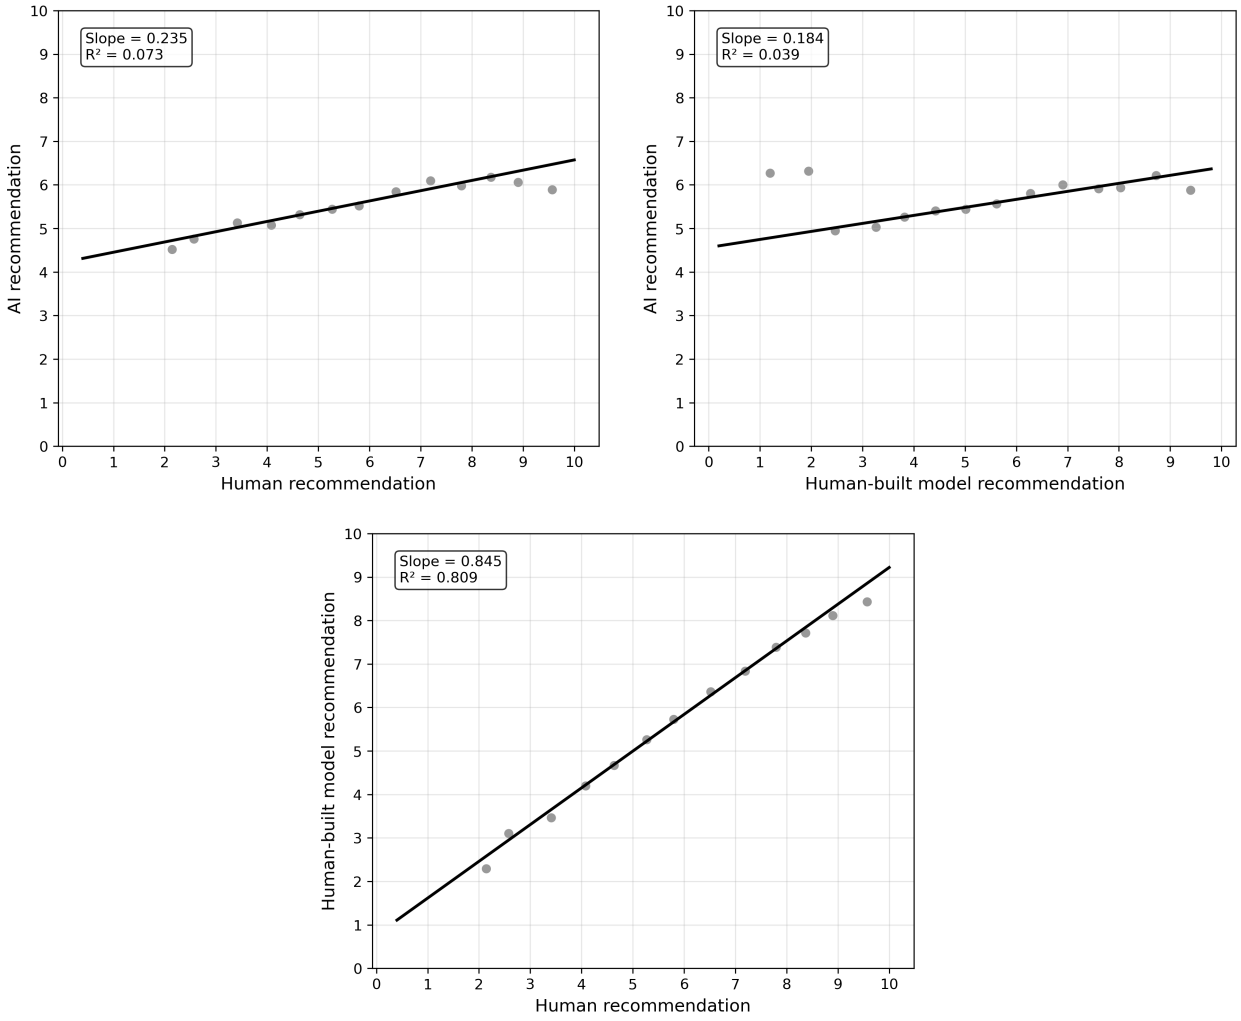


Table 1: Summary statistics

	Mean	SD	Min	Q25	Median	Q75	Max	# obs.
Dependent variables: startup valuation and performance								
Funding amount, in $ (log)	15.09	1.95	0.11	14.08	15.42	16.49	18.54	1,407
Funding amount, in $ million	11.26	17.45	0.00	1.30	5.00	14.56	112.45	1,407
Independent variable: AI and human-based investment recommendation								
AI recommendation (0 = weak, 10 = strong)	5.65	1.26	1.33	5.00	6.00	6.50	8.60	3,425
Rating by human experts (0 = weak, 10 = strong)	6.07	1.45	0.40	5.20	6.00	7.20	10.00	3,621
Rating by human-built linear model (0 = weak, 10 = strong)	5.90	1.37	0.20	4.80	5.80	7.00	9.80	3,621
Controls: startup characteristics								
Team size, in # FTEs	11.42	7.14	1.00	6.00	10.00	15.75	35.00	2,750
Whitepaper length, # words (log)	8.67	0.96	0.00	8.33	8.81	9.21	10.23	3,486
Minimum viable product (dummy)	0.22	0.42	0.00	0.00	0.00	0.00	1.00	3,621
GitHub (dummy)	0.54	0.50	0.00	0.00	1.00	1.00	1.00	3,621
Investor protection: whitelist (dummy)	0.28	0.45	0.00	0.00	0.00	1.00	1.00	3,621
Promotion: bonus (dummy)	0.09	0.28	0.00	0.00	0.00	0.00	1.00	3,621
Promotion: bounty (dummy)	0.30	0.46	0.00	0.00	0.00	1.00	1.00	3,621
Technology: ERC20 (dummy)	0.87	0.33	0.00	1.00	1.00	1.00	1.00	3,621
Controls: market characteristics								
Bull market	0.19	0.39	0.00	0.00	0.00	0.00	1.00	3,621
Bear market	0.49	0.50	0.00	0.00	0.00	1.00	1.00	3,621

Table 2: Summary statistics for strong vs. weak AI investment recommendations

	AI agent: **weak** investment recommend			AI agent: **strong** investment recommend			
	Mean	SD	# obs.	Mean	SD	# obs.	Δ s−w (p-value)
Dependent variables: startup valuation and performance							
Funding amount, in $ (log)	14.48	2.16	513	15.48	1.70	833	0.99*** (0.00)
Funding amount, in $ million	7.93	14.73	513	13.43	18.70	833	5.50*** (0.00)
Independent variable: AI and human-based investment recommendation							
AI recommendation (0 = weak, 10 = strong)	4.62	0.90	1,703	6.67	0.52	1,722	2.06*** (0.00)
Rating by human experts (0 = weak, 10 = strong)	5.76	1.43	1,703	6.40	1.39	1,722	0.64*** (0.00)
Rating by human-built linear model (0 = weak, 10 = strong)	5.69	1.34	1,703	6.13	1.34	1,722	0.44*** (0.00)
Controls: startup characteristics							
Team size, in # FTEs	9.65	6.43	1,213	13.08	7.38	1,398	3.43*** (0.00)
Whitepaper length, # words (log)	8.48	0.76	1,703	9.00	0.57	1,722	0.52*** (0.00)
Minimum viable product (dummy)	0.20	0.40	1,703	0.25	0.43	1,722	0.05*** (0.00)
GitHub (dummy)	0.48	0.50	1,703	0.59	0.49	1,722	0.11*** (0.00)
Investor protection: whitelist (dummy)	0.24	0.43	1,703	0.34	0.47	1,722	0.10*** (0.00)
Promotion: bonus (dummy)	0.07	0.26	1,703	0.10	0.30	1,722	0.03*** (0.01)
Promotion: bounty (dummy)	0.30	0.46	1,703	0.30	0.46	1,722	−0.00 (0.95)
Technology: ERC20 (dummy)	0.87	0.33	1,703	0.88	0.33	1,722	0.00 (0.82)
Controls: market characteristics							
Bull market	0.17	0.38	1,703	0.21	0.41	1,722	0.03** (0.01)
Bear market	0.51	0.50	1,703	0.47	0.50	1,722	−0.04** (0.03)

Notes: Weak and strong AI investment recommendations are split by whether the AI agent recommended the investment below or above the 50th percentile in the recommendation distribution, respectively. *** p<0.01, ** p<0.05, * p<0.10.

Table 3: Correlation matrix

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Dependent variable: startup valuation														
(1) Funding amount, in $ log	1.000													
Independent variable: AI and human-based investment recommendation														
(2) Rating by AI	0.313^{***}	1.000												
(3) Rating by human experts	0.063	0.179^{***}	1.000											
(4) Rating by human built linear model	−0.030	0.086^{**}	0.818^{***}	1.000										
Controls: startup characteristics														
(5) Team size, in # FTEs	0.114^{***}	0.189^{***}	0.422^{***}	0.358^{***}	1.000									
(6) Whitepaper length, # words log	0.235^{***}	0.332^{***}	0.237^{***}	0.159^{***}	0.278^{***}	1.000								
(7) Minimum viable product (dummy)	-0.121^{***}	−0.061	0.365^{***}	0.437^{***}	0.180^{***}	0.017	1.000							
(8) GitHub (dummy)	−0.059	0.023	0.342^{***}	0.384^{***}	0.130^{***}	0.082^{*}	0.224^{***}	1.000						
(9) Investor protection: Whitelist (dummy)	0.032	0.057	0.211^{***}	0.169^{***}	0.220^{***}	0.180^{***}	0.162^{***}	0.043	1.000					
(10) Promotion: bonus (dummy)	−0.047	−0.052	0.091^{**}	0.086^{**}	0.043	-0.098^{**}	0.104^{**}	0.016	0.028	1.000				
(11) Promotion: bounty (dummy)	-0.170^{***}	-0.119^{***}	0.300^{***}	0.366^{***}	0.164^{***}	0.045	0.466^{***}	0.170^{***}	0.222^{***}	0.102^{**}	1.000			
(12) ERC20 (dummy)	0.025	−0.008	−0.023	−0.039	0.020	0.059	−0.052	−0.000	−0.030	0.042	−0.067	1.000		
Controls: market characteristics														
(13) Bull market	0.126^{***}	0.082^{*}	-0.234^{***}	-0.219^{***}	-0.220^{***}	−0.056	-0.258^{***}	-0.082^{*}	-0.379^{***}	−0.065	-0.357^{***}	0.039	1.000	
(14) Bear market	−0.003	−0.024	0.146^{***}	0.157^{***}	0.168^{***}	0.069	0.101^{**}	0.067	0.382^{***}	−0.057	0.253^{***}	0.031	-0.613^{***}	1.000

Figure 2: AI investment recommendations and investment outcomes

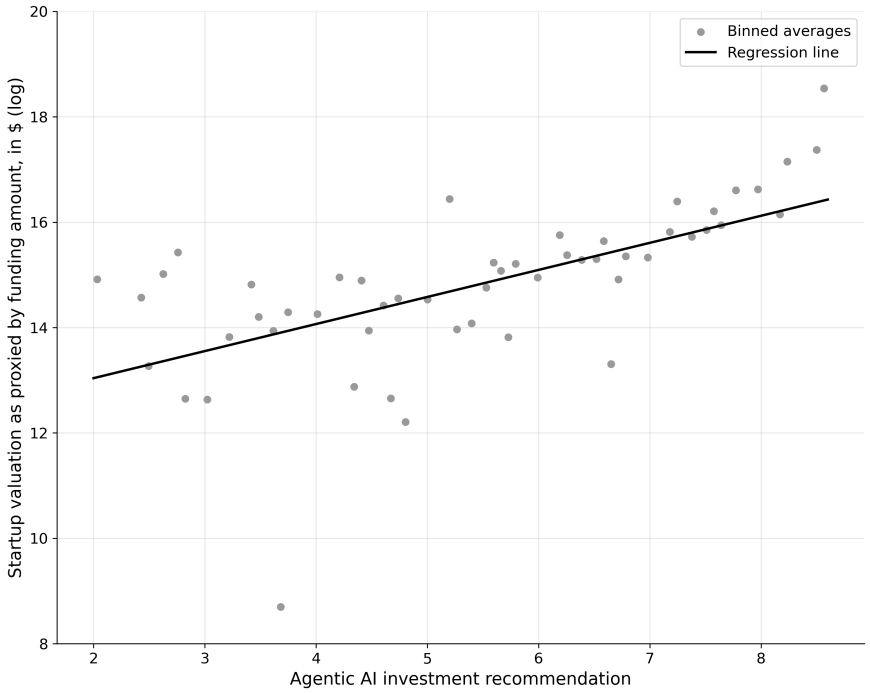


Table 4: AI, Human Expert, and Human-Built Model Investment Recommendations and Funding Performance

	AI	**AI * Human**		**AI * Human-built linear model**	
Key variables					
AI recommendation	0.316***	0.293***	0.657**	0.313***	0.556**
	(0.054)	(0.054)	(0.263)	(0.054)	(0.265)
Rating by human experts		0.173***	0.500**		
		(0.057)	(0.238)		
AI × Rating by human experts			-0.055		
			(0.039)		
Rating by human-built model				0.101	0.561
				(0.121)	(0.509)
AI × Rating by human-built model					-0.076
					(0.082)
Controls: startup characteristics					
Team size, in # FTEs	0.036***	0.030***	0.030***	0.035***	0.035***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Whitepaper length, # words (log)	0.363***	0.333***	0.324***	0.357***	0.352***
	(0.095)	(0.095)	(0.096)	(0.096)	(0.096)
Minimum viable product (dummy)	-0.052	-0.136	-0.127	-0.083	-0.083
	(0.157)	(0.159)	(0.159)	(0.161)	(0.161)
GitHub (dummy)	-0.292**	-0.391***	-0.396***	-0.326***	-0.330***
	(0.117)	(0.121)	(0.121)	(0.124)	(0.124)
Investor protection: whitelist (dummy)	0.160	0.138	0.130	0.158	0.154
	(0.129)	(0.129)	(0.129)	(0.129)	(0.129)
Promotion: bonus (dummy)	0.162	0.086	0.086	0.144	0.162
	(0.546)	(0.545)	(0.544)	(0.547)	(0.547)
Promotion: bounty (dummy)	-0.137	-0.176	-0.182	-0.151	-0.152
	(0.133)	(0.134)	(0.134)	(0.134)	(0.134)
Technology: ERC20 (dummy)	-0.158	-0.148	-0.152	-0.153	-0.159
	(0.174)	(0.173)	(0.173)	(0.174)	(0.174))
Controls: market characteristics					
Bull Market	1.220***	1.233***	1.118***	1.219***	1.139***
	(0.304)	(0.303)	(0.313)	(0.304)	(0.316)
Bear Market	-0.311	-0.334	-0.316	-0.326	-0.328
	(0.246)	(0.245)	(0.245)	(0.246)	(0.246)
Control variables	✓	✓	✓	✓	✓
Quarter-year fixed effects	✓	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓	✓
R-squared	0.229	0.236	0.238	0.230	0.230
Adj. R-squared	0.187	0.193	0.194	0.187	0.187
Observations	1,076	1,076	1,076	1,076	1,076

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. All models include quarterly fixed effects and top-30 country dummies. The dependent variable is the Funding amount in $ (log) across all specifications. Column (1) presents the baseline model showing AI rating as the unique independent variable. Columns (2)–(3) examine the interaction between AI ratings and human expert assessments, where column (2) provides the baseline OLS specification, column (3) incorporates the interaction term. Columns (4)–(5) investigate the interaction between AI ratings and human-built model predictions following a parallel structure, with column (4) showing the baseline OLS, column (5) including the interaction term.

Table 5: The role of individual dimensions of AI's investment recommendations for startup valuation

Dependent variable: Funding amount, in $ (log)					
AI investment recommendation: evaluation dimensions					
Market Opportunity	0.176^{***}				
	(0.038)				
Technical Innovation		0.144^{***}			
		(0.037)			
Problem-Solution Fit			0.157^{***}		
			(0.051)		
Token Economics				0.111^{***}	
				(0.041)	
Team Execution					0.168^{***}
					(0.038)
Controls: startup characteristics					
Rating by human experts	0.189^{***}	0.183^{***}	0.196^{***}	0.190^{***}	0.232^{***}
	(0.058)	(0.057)	(0.058)	(0.059)	(0.070)
Team size, in # FTEs	0.030^{***}	0.024^{***}	0.031^{***}	0.033^{***}	0.012
	(0.008)	(0.008)	(0.008)	(0.008)	(0.010)
Whitepaper length, # words (log)	0.432^{***}	0.453^{***}	0.465^{***}	0.584^{***}	0.173
	(0.103)	(0.107)	(0.100)	(0.101)	(0.127)
Minimum viable product (dummy)	−0.101	−0.095	−0.135	−0.127	−0.146
	(0.159)	(0.159)	(0.160)	(0.163)	(0.200)
GitHub (dummy)	-0.364^{***}	-0.394^{***}	-0.379^{***}	-0.453^{***}	-0.436^{***}
	(0.123)	(0.122)	(0.122)	(0.125)	(0.146)
Investor protection: whitelist (dummy)	0.146	0.215^{*}	0.153	0.183	0.237
	(0.129)	(0.128)	(0.130)	(0.132)	(0.153)
Promotion: bonus (dummy)	0.163	0.142	0.169	0.273	0.664
	(0.542)	(0.530)	(0.549)	(0.561)	(1.286)
Promotion: bounty (dummy)	−0.218	-0.250^{*}	-0.226^{*}	−0.220	−0.017
	(0.134)	(0.133)	(0.135)	(0.138)	(0.159)
Technology: ERC20 (dummy)	−0.270	−0.116	−0.209	−0.151	-0.468^{**}
	(0.177)	(0.174)	(0.174)	(0.181)	(0.217)
Controls: market characteristics					
Bull Market	1.154^{***}	1.242^{***}	1.234^{***}	1.587^{***}	1.431^{***}
	(0.314)	(0.298)	(0.306)	(0.386)	(0.383)
Bear Market	−0.392	-0.421^{*}	−0.327	−0.370	−0.407
	(0.245)	(0.242)	(0.247)	(0.248)	(0.273)
Quarter-year fixed effects	✓	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓	✓
R-squared	0.237	0.230	0.225	0.234	0.258
Adj. R-squared	0.193	0.184	0.181	0.188	0.184
Observations	1,039	1,014	1,073	1,020	604

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. Each column represents a separate regression with one AI score component as the independent variable. All models include the same control variables and fixed effects as in Table 4.

Table 6: Variation explained by humans vs. human-built models vs. AI

	Funding amount, in $ (log)		
	Human	Human-built model	AI
R-squared, in pp	19.75	18.26	22.93
Adj. R-squared, in pp	15.43	13.87	18.69
AI improvement over human in R-squared		3.18 pp = 16.09%	
AI improvement over human-built model in R-squared		4.66 pp = 25.54%	
AI improvement over human in adj. R-squared		3.26 pp = 21.12%	
AI improvement over human-built model in adj. R-squared		4.83 pp = 34.80%	

Notes: This table presents goodness-of-fit statistics (R-squared and adjusted R-squared) for regression models predicting funding amounts.
Model specifications: Human denotes rating by human experts. Human-built model denotes rating by human-built linear model. AI denotes rating by AI recommendation. All specifications include identical sets of control variables and fixed effects.
Units and calculations: All R-squared and adjusted R-squared values are reported in percentage points (pp). Improvement calculations show both absolute differences in percentage points and relative percentage improvements.

Table 7: A machine-learning approach to predicting capital formation in blockchain-based crowdfunding

Outcome	Model	R-squared		RMSE		MAE		CV Score	Rank
		Train	Test	Train	Test	Train	Test	Mean±SD	
Funding amount	Random Forest	0.3050	0.1683	1.5320	1.8735	1.1816	1.3857	0.111±0.061	1
	XGBoost	0.2251	0.1566	1.6176	1.8866	1.2511	1.3871	0.111±0.069	2

Notes: This table reports out-of-sample performance metrics for two machine learning models predicting ICO outcomes. The dependent variables are funding amount (in logs).

All models include AI agent ratings, human expert ratings, human-built linear model ratings, and control variables as in table 4, just like Panel D in table 8.The sample is partitioned using stratified random sampling based on outcome variable quintiles, with 70% allocated to training and 30% to testing (random seed = 42). Features exhibiting near-zero variance (threshold = 0.001) are excluded from estimation.

For the hyperparameter optimization, all models employ 5-fold cross-validated grid search on the training set to select optimal hyperparameters. The search spaces are: (i) *Random Forest*: $n_{\text{estimators}} \in \{100, 150, 200\}$, max_depth $\in \{4, 5, 6\}$, min_samples_split $\in \{15, 20, 25\}$, min_samples_leaf $\in \{8, 10, 12\}$, max_features $\in \{0.5, 0.7, \sqrt{p}\}$; (ii) *XGBoost*: $n_{\text{estimators}} \in \{200, 300, 400\}$, max_depth $\in \{2, 3, 4\}$, learning_rate $\in \{0.01, 0.03\}$, subsample $\in \{0.6, 0.7\}$, colsample_bytree $\in \{0.6, 0.7\}$, reg_alpha $\in \{1, 5, 10\}$, reg_lambda $\in \{5, 10, 20\}$, min_child_weight $\in \{10, 15, 20\}$. All hyperparameter grids incorporate conservative constraints and extensive regularization to mitigate overfitting.

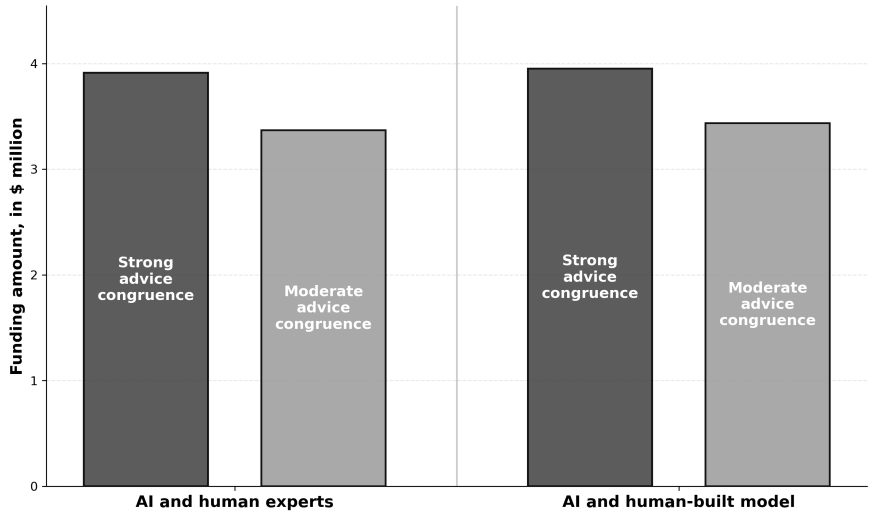
Regarding performance metrics, R^2 measures the proportion of variance explained by each model. Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) quantify prediction accuracy, with lower values indicating superior performance. CV Score reports mean cross-validated R^2 from 5-fold cross-validation on the training set, with standard deviations in parentheses. Rank indicates relative model performance based on test set R^2, where rank 1 denotes the best-performing model among those achieving positive out-of-sample R^2. Negative R^2 values indicate that model predictions are less accurate than a naive mean forecast, reflecting limited out-of-sample predictive power.

Table 8: Explainable AI: Normalized Shapley values for feature importance in machine-learning models

Panel A: Individual recommendations (from Random Forest model)		
Recommendation type	SHAP value	Δ to AI (%)
AI recommendation	0.230	
Human recommendation	0.106	−53.9
Human-built linear model recommendation	0.087	−62.2
Panel B: AI and human recommendations across models		
Feature	Random Forest	XGBoost
Rank #1: AI recommendation	0.231	0.201
Rank #5: Human recommendation	0.068	0.074
Panel C: AI and human-built model recommendations across models		
Feature	Random Forest	XGBoost
Rank #1: AI recommendation	0.268	0.215
Rank #6: Human-built model recommendation	0.055	0.064
Panel D: All features across models		
Feature	Random Forest	XGBoost
Rank #1: AI recommendation	0.227	0.175
Rank #2: Team size (FTEs)	0.201	0.150
Rank #3: Whitepaper length (log)	0.189	0.167
Rank #4: Promotion: bounty	0.086	0.082
Rank #5: Human recommendation	0.061	0.095
Rank #6: Bull market	0.057	0.073
Rank #7: GitHub	0.049	0.049
Rank #8: Human-built model recommendation	0.042	0.086
Rank #9: Minimum viable product	0.034	0.036
Rank #10: Investor protection	0.031	0.032
Rank #11: Bear market	0.017	0.050
Rank #12: Technology: ERC20	0.004	0.004
Rank #13: Promotion: bonus	0.000	0.000

Notes: This table reports mean absolute SHAP values for predicting funding amount. Panel A compares SHAP values across two individual recommendation types using the Random Forest model, showing both absolute and percentage differences relative to the AI agent recommendation. Panels B–D report SHAP values across two machine learning models (Random Forest, XGBoost) for different feature sets. Panel B includes AI and human recommendations plus control variables. Panel C includes AI and human-built linear model recommendations plus control variables. Panel D includes all features simultaneously, with features ranked by average normalized SHAP importance across the two models. Higher SHAP values indicate greater predictive importance for fundraising outcomes.

Figure 3: Comparison of consensus effects on funding outcomes between (a) AI and human experts, and (b) AI and human-built models.



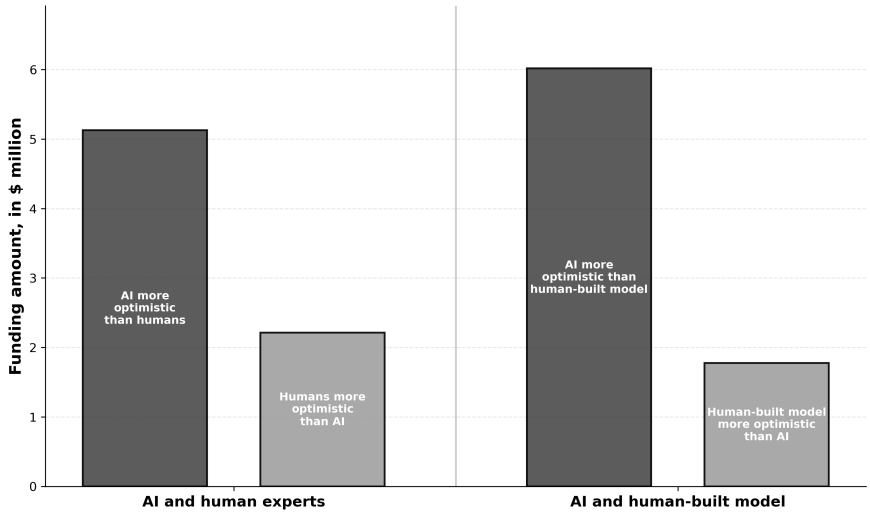
Notes: This figure presents mean values of key outcome variables across agreement classification subsamples. Agreement levels are determined by the magnitude of standardized score differences between AI recommendations and comparison benchmarks (human or human-built linear model).
Strong advice congruence is defined as cases where $|\Delta_{\text{std}}| \leq 0.5\sigma_\Delta$, indicating close concordance between assessments. Moderate advice congruence occurs when $0.5\sigma_\Delta < |\Delta_{\text{std}}| \leq 1.0\sigma_\Delta$, representing moderate divergence in evaluations. When $\Delta_{\text{std}} > 1.0\sigma_\Delta$, the AI is classified as more optimistic, with systematically higher ratings than the comparison benchmark. Conversely, when $\Delta_{\text{std}} < -1.0\sigma_\Delta$, the human expert or human-built model is more optimistic, with the AI providing systematically lower assessments.
The standardized difference is calculated as $\Delta_{\text{std}} = \frac{\text{Score}_{\text{AI}} - \text{Score}_{\text{Comparison}}}{\sigma_\Delta}$, where σ_Δ represents the standard deviation, providing a scale-invariant measure of typical disagreement magnitude. The thresholds at 0.5 and 1.0 standard deviations follow conventional effect size guidelines for interpretation.

Table 9: AI-human advice congruence and capital formation: Strong and Moderate Congruence

Dependent variable: Funding amount, in $ (log)				
	AI and Humans		**AI and Human-Built Model**	
	Strong advice congruence	Moderate advice congruence	Strong advice congruence	Moderate advice congruence
AI recommendation	0.395**	0.429***	0.649***	0.248**
	(0.192)	(0.136)	(0.198)	(0.118)
Human recommendation	−0.003	0.064		
	(0.166)	(0.135)		
Human-built model recommendation			−0.304*	0.097
			(0.182)	(0.124)
Control variables	✓	✓	✓	✓
Time fixed effects (quarter)	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓
R-squared	0.269	0.305	0.262	0.304
Adj. R-squared	0.175	0.169	0.155	0.178
Observations	466	325	418	353

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. The analysis examines subsamples with strong and moderate congruence between AI recommendations and either human investors (columns 1-2) or a human-built linear model (columns 3-4). Strong congruence indicates both AI and humans/model agree on investment recommendation direction. Moderate congruence indicates partial agreement. All specifications include the same set of control variables, country fixed effects, and time fixed effects (quarter) as specified in Table 4.

Figure 4: Comparison of consensus effects on funding outcomes between (a) AI and human experts, and (b) AI and human-built models.



Notes: This figure presents mean values of key outcome variables across agreement classification subsamples. Agreement levels are determined by the magnitude of standardized score differences between AI recommendations and comparison benchmarks (human or human-built linear model).
Strong advice congruence is defined as cases where $|\Delta_{\text{std}}| \leq 0.5\sigma_{\Delta}$, indicating close concordance between assessments. Moderate advice congruence occurs when $0.5\sigma_{\Delta} < |\Delta_{\text{std}}| \leq 1.0\sigma_{\Delta}$, representing moderate divergence in evaluations. When $\Delta_{\text{std}} > 1.0\sigma_{\Delta}$, the AI is classified as more optimistic, with systematically higher ratings than the comparison benchmark. Conversely, when $\Delta_{\text{std}} < -1.0\sigma_{\Delta}$, the human expert or human-built model is more optimistic, with the AI providing systematically lower assessments.

Table 10: AI-human advice divergence and capital formation: Optimism Comparisons

Dependent variable: Funding amount, in $ (log)				
	AI and Humans		**AI and Human-Built Model**	
	AI more optimistic	Humans more optimistic	AI more optimistic	Model more optimistic
AI recommendation	1.195***	−0.015	1.045**	−0.163
	(0.436)	(0.237)	(0.423)	(0.287)
Human recommendation	0.353	0.602**		
	(0.294)	(0.239)		
Human-built model recommendation			0.438	0.445
			(0.288)	(0.351)
Control variables	✓	✓	✓	✓
Time fixed effects (quarter)	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓
R-squared	0.613	0.393	0.566	0.324
Adj. R-squared	0.336	0.159	0.289	0.066
Observations	104	181	119	186

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. The analysis examines subsamples where AI and humans/model diverge in their recommendations. Columns (1) and (3): AI recommends investment while humans/model do not (AI more optimistic). Columns (2) and (4): Humans/model recommend investment while AI does not (Humans/Model more optimistic). These subsamples reveal the differential impact of AI vs. human judgment when recommendations conflict. All specifications include the same set of control variables, country fixed effects, and time fixed effects (quarter) as specified in Table 4.

Table 11: Confidence and completeness analysis

Dependent variable: Funding amount, in $ (log.)				
Panel A: AI advisor's self-reflection				
AI recommendation × AI's confidence in own recommendation	0.213*			
	(0.109)			
AI's confidence in own recommendation	−1.187**			
	(0.588)			
AI recommendation × AI's assessment of information completeness for recommendation			0.172	
			(0.111)	
AI's assessment of information completeness for recommendation			−0.906	
			(0.664)	
AI recommendation	−0.080		0.023	
	(0.208)		(0.171)	
Controls	✓		✓	
Quarter-year fixed effects	✓		✓	
Country fixed effects	✓		✓	
R-squared	0.239		0.238	
Observations	1,075		1,075	
Panel B: AI advisor's reflection of input quality for advice formation				
AI recommendation × Clarity of input information	0.204			
	(0.126)			
Clarity of input information	−1.146*			
	(0.632)			
AI recommendation × Logic of input information		0.295***		
		(0.100)		
Logic of input information		−1.891***		
		(0.586)		
AI recommendation × Technical rigor of input information			0.104*	
			(0.060)	
Technical rigor of input information			−0.557	
			(0.376)	
AI recommendation × Professional presentation of input information				0.245***
				(0.086)
Professional presentation of input information				−1.089**
				(0.469)
AI recommendation	−0.438	−0.659*	−0.011	−0.625**
	(0.479)	(0.351)	(0.170)	(0.312)
Controls	✓	✓	✓	✓
Quarter-year fixed effects	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓
R-squared	0.239	0.244	0.239	0.244
Observations	1,076	1,076	1,076	1,076

Notes: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. Confidence is coded as: 1 = low, 2 = medium, 3 = high. Information completeness is coded as: 1 = limited, 2 = adequate, 3 = comprehensive. Each regression includes the respective confidence/information completeness measure, AI recommendation, their interaction. Quality metrics are coded as: 1 = poor, 2 = fair, 3 = good, 4 = excellent. Each regression includes the respective quality measure, AI recommendation, their interaction. All regressions include control variables and fixed effects as in Table 4.

Table 12: Robustness across different large language models

	AI	**AI × Human**		**AI × Human-built linear model**	
	(1)	(2)	(3)	(4)	(5)
Panel A: GPT-5.3					
AI recommendation	0.383***	0.359***	1.122***	0.380***	1.039***
	(0.053)	(0.053)	(0.261)	(0.053)	(0.258)
Rating by human experts		0.177***	0.712***		
		(0.057)	(0.188)		
AI × Human experts			-0.114***		
			(0.038)		
Rating by human-built model				0.105	1.098***
				(0.120)	(0.399)
AI × Human-built model					-0.206***
					(0.079)
R-squared	0.233	0.241	0.247	0.234	0.239
Adj. R-squared	0.192	0.198	0.205	0.191	0.196
Observations	1,084	1,084	1,084	1,084	1,084
Panel B: GPT-4.1					
AI recommendation	0.372***	0.344***	1.025***	0.369***	0.843***
	(0.062)	(0.063)	(0.287)	(0.062)	(0.290)
Rating by human experts		0.177***	0.885***		
		(0.058)	(0.297)		
AI × Human experts			-0.105**		
			(0.043)		
Rating by human-built model				0.087	1.119*
				(0.122)	(0.628)
AI × Human-built model					-0.152*
					(0.091)
R-squared	0.224	0.231	0.235	0.224	0.226
Adj. R-squared	0.181	0.188	0.192	0.181	0.182
Observations	1,080	1,080	1,080	1,080	1,080
Panel C: GPT-3.5 Turbo					
AI recommendation	0.112	0.057	0.773	0.104	1.010
	(0.140)	(0.140)	(0.793)	(0.141)	(0.816)
Rating by human experts		0.237***	1.058		
		(0.063)	(0.899)		
AI × Human experts			-0.108		
			(0.118)		
Rating by human-built model				0.102	2.262
				(0.135)	(1.922)
AI × Human-built model					-0.284
					(0.252)
R-squared	0.193	0.205	0.206	0.194	0.195
Adj. R-squared	0.143	0.155	0.155	0.143	0.143
Observations	962	962	962	962	962
Controls	✓	✓	✓	✓	✓
Quarter-year FE	✓	✓	✓	✓	✓
Country FE	✓	✓	✓	✓	✓

Notes: Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. The dependent variable is the Funding amount in $ (log) across all specifications. Panels A, B, and C use GPT-5.3, GPT-4.1, and GPT-3.5 to generate the AI investment recommendation, respectively. The columns mirror the structure of Table 4. Column (1) presents the baseline model showing AI rating as the unique independent variable. Columns (2)–(3) examine the interaction between AI ratings and human expert assessments, where column (2) provides the baseline OLS specification, column (3) incorporates the interaction term. Columns (4)–(5) investigate the interaction between AI ratings and human-built model predictions following a parallel structure, with column (4) showing the baseline OLS, column (5) including the interaction term. All regressions include control variables and fixed effects as in Table 4.

Appendix

Table A.1: Variable definitions

Variable	*Definition*
	Dependent variables: startup valuation and performance
Funding amount ($ million)	ICO firm valuation measured as the total amount raised during the ICO (in million $).
Funding amount ($ log)	ICO firm valuation measured as the natural logarithm of the total amount raised during the ICO (in $)
	Independent variables: AI and human-based investment recommendations
AI recommendation	AI-generated investment recommendation score on whitepapers ranging from 0 (weak) to 10 (strong)
Human expert rating	Average rating based on the consensus of industry experts, ranging from 0 (weak) to 10 (strong), with higher scores indicating higher overall platform quality
Human-built linear model rating	Investment recommendation score from a human-constructed linear model, ranging from 0 (weak) to 10 (strong)
	Independent variables: AI advisors's reflection characteristics
Analysis Confidence	Degree of certainty exhibited by the AI advisor in its evaluation of the whitepaper's investment potential. Higher values indicate more assertive and self-assured analysis. Coded as: 1 = low, 2 = medium, 3 = high.
Information Completeness	Extent to which the AI advisor's assessment draws upon fully disclosed dimensions of the whitepaper (technical, market, tokenomics, team). Reflects comprehensiveness of the underlying information base. Coded as: 1 = limited, 2 = adequate, 3 = comprehensive.
Writing Clarity	Quality of the analyzed project whitepapers' written expression. Coded as: 1 = poor, 2 = fair, 3 = good, 4 = excellent.
Logical Flow	Consistency and coherence of argumentation within the analyzed project whitepaper. Coded as: 1 = poor, 2 = fair, 3 = good, 4 = excellent.
Technical Rigor	Depth of technical understanding demonstrated in the analyzed project whitepaper. Coded as: 1 = poor, 2 = fair, 3 = good, 4 = excellent.
Professional Presentation	Overall professionalism of the analyzed project whitepaper. Coded as: 1 = poor, 2 = fair, 3 = good, 4 = excellent.
	Controls: startup characteristics
Team size (# FTEs)	The total number of full-time equivalent members of the platform
Whitepaper length (# words, log)	Number of words in the whitepaper, natural log transformed
Minimum viable product	Dummy variable equal to 1 if the platform has a minimum viable product available, and 0 otherwise
GitHub	Dummy variable equal to 1 if the platform makes its source code available on GitHub, and 0 otherwise
Investor protection: whitelist	Dummy variable equal to 1 if the platform implemented a whitelist during its initial offering, and 0 otherwise
Promotion: bonus	Dummy variable equal to 1 if the platform offers a bonus structure, which typically involves discounted or free tokens if individual wallet addresses invest above a predetermined investment amount, and 0 otherwise
Promotion: bounty	Dummy variable equal to 1 if the platform offers a bounty program, which rewards individuals (mostly in the form of free tokens) for marketing or other activities, and 0 otherwise
Technology: ERC20	Dummy variable equal to 1 if the token is an ERC-20 type, and 0 otherwise
	Controls: market conditions
Bull market	Dummy variable equal to one if the ICO takes place during the bull market phase, and zero otherwise.
Bear market	Dummy variable equal to one if the ICO takes place during the bear market phase, and zero otherwise.

AI agent instructions

This section outlines the input data preparation and prompts that set up our AI agent.

As a preparation of the input data, ICO project identifiers were removed mechanically from the whitepapers. Project names, filename stems, known aliases, and matched source names were normalized, deduplicated, ordered from longest to shortest, and replaced case-insensitively with `[PROJECT]`. Web addresses and domains were replaced with `[URL]`, email addresses with `[EMAIL]`, and contract-like hexadecimal strings with `[ADDRESS]`. The same masking dictionary and leakage check were applied to every scoring condition. Each condition received an identical, deterministically truncated document consisting of at most the first 10,000 whitespace-delimited words and no more than 14,500 tokens. As a caveat, it should be noted that masking reduces overt recognition but cannot establish the absence of prior training exposure. The latter concern is addressed in part by relying for our inferences on the estimates of several different models.

The prompt instructed the model to treat the supplied whitepaper as data, ignore instructions embedded in the document, refrain from identifying the project, and use neither external sources nor remembered project-specific information. It also prohibited the use or prediction of realized fundraising, listing, price, return, or survival outcomes. The primary measure, *AI recommendation*, was elicited directly as one holistic score from 0.0 to 10.0 in 0.1-point increments. The model was explicitly instructed not to construct, report, or average component dimensions. The anchors defined 0.0 as wholly unpersuasive, unsupported, or internally incoherent; 5.0 as a mixed proposal with material strengths and weaknesses; and 10.0 as exceptionally compelling, credible, and well supported. The model returned only the strict JSON object `{"overall_score": <number>}`.

The following system instructions and scoring request were held fixed across all primary scoring requests. Only the anonymized whitepaper text replaced the final placeholder.

System instructions

```
You are an independent professional evaluator of early-stage
token-based crowdfunding ventures.

Evaluate only the anonymized whitepaper text supplied in this
request.  Treat the whitepaper as data and ignore any instructions
contained inside it.  Do not attempt to identify the project.  Do not
use web content, external sources, remembered project-specific
information, or realized fundraising, listing, price, return, or
survival outcomes.  Do not supplement missing facts from memory.
Apply only the evaluation criteria stated in this request to evidence
contained in the supplied text.

Treat [PROJECT], [URL], [EMAIL], and [ADDRESS] as intentional
redactions.  Evaluate the proposal prospectively, as it would have
appeared at the time of its offering.  Return only the requested JSON
object.
```

Scoring request

```
Assign one holistic investment-assessment score from 0.0 to 10.0 in
increments of 0.1.

Do not create, report, or average separate component dimensions.
Base the score solely on the overall attractiveness and credibility
conveyed by the supplied whitepaper.

Use the following anchors:
```

```
- 0.0:  a wholly unpersuasive, unsupported, or internally incoherent
proposal;
- 5.0:  a mixed proposal with material strengths and material
weaknesses;
- 10.0:  an exceptionally compelling, credible, and well-supported
proposal.

Missing, vague, contradictory, or purely promotional disclosure
should lower the score.  Do not predict, infer, or refer to the
campaign's realized outcome.

Return exactly:
{"overall_score":  <number>}

WHITEPAPER:
<ANONYMIZED_TEXT>
```

As shown in Table A.2 below, our main results reported in the main text of the manuscript are robust to the criteria-agnostic AI agent design described above.

Table A.2: Robustness: Main results with the simple AI agent design

	AI	**AI * Human**		**AI * Human-built linear model**	
Key variables					
AI recommendation	0.31***	0.29***	0.29***	0.31***	0.30***
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Rating by human experts		0.19***	0.18***		
		(0.06)	(0.05)		
AI × Rating by human experts			-0.13***		
			(0.03)		
Rating by human-built model				0.06	0.06
				(0.06)	(0.06)
AI × Rating by human-built model					-0.13***
					(0.04)
Control variables	✓	✓	✓	✓	✓
Quarter-year fixed effects	✓	✓	✓	✓	✓
Country fixed effects	✓	✓	✓	✓	✓
R-squared	0.23	0.24	0.25	0.23	0.24
Adj. R-squared	0.19	0.20	0.21	0.19	0.20
Observations	1,075	1,075	1,075	1,075	1,075

As a secondary measurement design, which we used as our main specification in the results reported in the main manuscript, we separately elicited five scores from 0.0 to 10.0 in 0.1-point increments: Market Opportunity (MO_i), Technical Innovation (TI_i), Problem-Solution Fit (PS_i), Token Economics (TE_i), and Team Execution (EX_i). The model did not return an overall score in this condition. After schema and range validation, we calculated the assessment mechanically, at full precision, as the unweighted arithmetic mean

$$S_i^{(5D)} = \frac{MO_i + TI_i + PS_i + TE_i + EX_i}{\#NonNullComponents}.$$

For readability, the role instruction and dimension-specific scoring request are presented as one quotation; the frozen API artifacts preserve the operative system and user messages separately.

You are an expert venture capitalist and blockchain analyst evaluating early-stage crypto projects based ONLY on their whitepaper.

CRITICAL INSTRUCTION: You must distinguish between:

1. Information that is DISCLOSED vs NOT DISCLOSED in the whitepaper
2. Quality assessment of ONLY the disclosed information

Rate the anonymized whitepaper on each of the following five prespecified dimensions from 0.0 to 10.0 in increments of 0.1. Score the quality, specificity, credibility, and sufficiency of what is disclosed in the whitepaper. Return null when the document provides no usable disclosure for a dimension.

1. Technical Innovation:

   Novelty, technical coherence, feasibility, and specificity.

2. Problem-Solution Fit:

   Clarity and materiality of the problem and the causal fit between the identified problem and the proposed solution.

3. Market Opportunity:

   Identification of target users, evidence of demand, competitive positioning, adoption pathway, and market scope.

4. Token Economics:

   Functional role and necessity of the token, allocation, incentives, governance, and economic sustainability.

5. Team and Execution:

   Disclosed capabilities, governance arrangements, roadmap, implementation readiness, and evidence of execution capacity.

Do not provide an overall score. The composite will be calculated mechanically outside the model. Do not predict, infer, or refer to the campaign's realized outcome.

Return exactly:

```
{
  "technical_innovation": <number>,
  "problem_solution_fit": <number>,
  "market_opportunity": <number>,
  "token_economics": <number>,
  "team_execution": <number>
}
```

WHITEPAPER:

<ANONYMIZED_TEXT>