AI-Powered (Finance) Scholarship*

Robert Novy-Marx† Mihail Velikov‡

March 17, 2025

Abstract

This paper describes a process for automatically generating academic finance papers using large language models (LLMs). It demonstrates the process' efficacy by producing hundreds of complete papers on stock return predictability, a topic particularly well-suited for our illustration. We first mine over 30,000 potential stock return predictor signals from accounting data, and apply the Novy-Marx and Velikov (2024) "Assaying Anomalies" protocol to generate standardized "template reports" for 96 signals that pass the protocol's rigorous criteria. Each report details a signal's performance predicting stock returns using a wide array of tests and benchmarks it to more than 200 other known anomalies. Finally, we use state-of-the-art LLMs to generate three distinct complete versions of academic papers for each signal. The different versions include creative names for the signals, contain custom introductions providing different theoretical justifications for the observed predictability patterns, and incorporate citations to existing (and, on occasion, imagined) literature supporting their respective claims. This experiment illustrates AI's potential for enhancing financial research efficiency, but also serves as a cautionary tale, illustrating how it can be abused to industrialize HARKing (Hypothesizing After Results are Known).

JEL classification: A12, C12, C18, C45, G11, G12
Keywords: Artificial Intelligence, Large Language Models, Hypothesis Testing, Anomalies

*Novy-Marx provides consulting services to Dimensional Fund Advisors, an investment firm headquartered in Austin, Texas with strong ties to the academic community. The thoughts and opinions expressed in this article are those of the authors alone, and no other person or institution has any control over its content. For helpful comments and suggestions, we thank Delina Agnosteva, Andrew Chen, Matt Gustafson, Bryan Kelly, Milena Novy-Marx, along with conference and seminar participants at Syracuse University and Utah Winter Finance Conference.

†Simon Business School and NBER, University of Rochester, 500 Joseph C. Wilson Blvd., Box 270100, Rochester, NY 14627. Email: robert.novy-marx@simon.rochester.edu.

‡Smeal College of Business, Pennsylvania State University, 366 Business Building, University Park, PA 16802. Email: velikov@psu.edu.

1 Introduction

Consider this scenario: a junior professor submits a paper documenting a novel return predictor, which includes precisely formulated hypotheses and robust empirical evidence. The paper is well written, the analysis appears correct, and the hypotheses accurately predict the patterns observed in the data. Should it matter if an AI system generated these hypotheses after seeing the results? This question cuts to the heart of how we understand scientific discovery and hypothesis formation, and how our views are being tested by the introduction of Large Language Models (LLMs).

In modern academia we face an inherent tension in our treatment of hypothesis formation. We often view post-hoc theorizing with suspicion, labeling it as "HARKing" (Hypothesizing After Results are Known) (Kerr, 1998). The prevailing academic standard insists that researchers should first develop their theories and predictions and then test them against data. Few significant scientific discoveries in history have, however, adhered to this standard. Legend has it that in 1666 Isaac Newton observed an apple falling from a tree. This observation led him to hypothesize about universal gravitation, forming testable predictions that explained not just falling apples, but planetary motion, tides, and countless other phenomena. Newton developed his hypotheses after observing the phenomena they would later explain. Post-observation hypothesis generation has always been, and will always be, a crucial part of scientific discovery.

This tension between an idealized scientific method and practical discovery is particularly evident in empirical asset pricing. Recent work by Chen, Lopez-Lira, and Zimmermann (2024) shows a striking parallel between data mining and traditional peer review, finding that both methods achieve similar predictability rates with approximately only 50% of predictive power persisting out-of-sample. Their finding that peer review systematically mislabels mispricing or luck as risk is consistent with a view that the anomaly literature often develops ex-post theoretical explanations

to fit observed empirical patterns rather than testing pre-specified economic mechanisms. In fact, the peer-review process strongly encourages this. Reviewers and editors often require papers documenting interesting new aspects of the data to include an "economic story" even when other authors could be far better suited to explain the economics behind the results documented in a paper.

The emergence of powerful LLMs has transformed this tension from a matter of scientific practice to one of technological capability. van Inwegen et al. (2023) shows that algorithmic writing assistance can improve outcomes without compromising signal value, while Horton (2023) explores how LLMs can function as *homo silicus* - computational analogues of *homo economicus*. These advances suggest that AI systems can meaningfully engage with economic reasoning and prediction. Building on this, Manning et al. (2024) present methods for automatically generating and testing scientific hypotheses *in silico*, though noting that LLMs may struggle with precise magnitude estimates. Si et al. (2024) show that LLMs can generate novel research ideas while maintaining high standards of feasibility and scientific merit. Bail (2024), Korinek (2023), and Liang et al. (2024) document LLMs' expanding capabilities across research domains. Most notably, Lu et al. (2024) develop an "AI Scientist" system that can independently generate research ideas, conduct experiments, and produce papers exceeding typical acceptance thresholds.

Drawing on these advancing capabilities, in this paper we demonstrate a complete pipeline for automated academic research production in finance, from hypothesis generation through full paper creation. Using stock return predictability as a testing ground, we first mine accounting data to identify over 30,000 potential predictors. We subsequently apply the Novy-Marx and Velikov (2024) "Assaying Anomalies" protocol to identify 96 signals that pass rigorous statistical criteria. Then, using GPT-3.5-turbo, we systematically generate descriptive names for these empirically discovered "return predictors," ensuring consistent and meaningful terminology across

papers. Using Claude 3.5-Sonnet and "template reports" generated by the "Assaying Anomalies" protocol, we machine-generate complete academic papers for each predictor. For each signal, we create three distinct versions of full papers, including abstract, introduction, data, results, and conclusion sections. For each of the signals, the three different versions of the papers contain different hypotheses and economic "explanations" while maintaining consistency with the empirical findings.

The 288 fully programmatically-generated papers contain introductions that follow standard academic conventions, developing theoretical arguments that connect the documented return patterns to established economic mechanisms, incorporating citations to existing (and, at least for now, on occasion hallucinated) literature. Each paper includes comprehensive descriptions of the data and methodology, detailed discussion of results, and contextualized conclusions. All of these papers are available at https://github.com/velikov-mihail/AI-Powered-Scholarship. While the papers and their theoretical frameworks are automatically generated, it's important to note that all empirical analyses and statistical validations are conducted using rigorous methods developed in the academic literature, ensuring the reliability (if not the interpretation) of the underlying findings.

This scale and automation of AI-powered research generation raises fundamental concerns about the integrity of knowledge production in the academic finance community. The profession has institutional safeguards in place against potential abuse of data mining and post-hoc theorizing. Perhaps most importantly, the profession rewards (at least in the long-run) scholarly reputation built through sustained contribution of influential work rather than mere quantity of publications. Researchers who consistently publish low-quality papers, even in substantial numbers, rarely achieve the field's highest honors or most prestigious appointments. The peer review process provides additional screening, with referees and editors scrutinizing not just statistical significance but theoretical foundations, methodological rigor, and broader

contribution to the literature (though, as noted earlier, for good or for bad this process in practice probably encourages more HARKing). The practice of presenting work at research seminars and conferences before publication creates opportunities for detailed questioning about theoretical mechanisms and research design choices. The recent increasing emphasis on replicability, including requirements for publicly sharing data and code, has added another layer of quality control. In a world powered by AI, this requirement is especially important because the same tools that make testing easier, and thus raise overfitting concerns, also lower the bar on replication and independently verifying robustness through additional tests.

These various institutional safeguards have generally served the profession well in maintaining research standards. The emergence of sophisticated AI systems capable of generating (multiple) plausible theoretical frameworks at scale, however, poses novel challenges to these traditional mechanisms. Chen and Dim (2024) demonstrate how rigorous data mining can produce predictive signals comparable to top finance journals. When AI systems can rapidly produce hundreds of seemingly coherent theoretical explanations for mined empirical results, how do we maintain meaningful quality control through our existing institutions and avoid overwhelming our traditional peer review process (a process that was already stressed by the profession's growth outpacing the growth in the number of quality outlets for the profession's output)?

Practical challenges amplify these concerns to academic integrity. Each of our AI-generated papers naturally includes citations to the literature to support its hypothesis development. When scaled to hundreds or thousands of papers, this automated citation generation could artificially inflate the citation counts of existing works, including our own.[1] The ease with which AI can generate convincing theoret-

[1]If Google's algorithms pick up the 96 paper titles we've generated and posted on our github page, that would result in at least $96 \times 3 = 288$ (Novy-Marx) and $96 \times 4 = 384$ (Velikov) additional google scholar citations for our own work (each paper cites Novy-Marx at least three times and Velikov at least four times; it is entirely possible that our instructions to the AI that generate cites

ical frameworks that reference prior literature may inadvertently create a new form of academic arbitrage – where researchers can boost their citation counts through automated paper generation. It is actually easy to imagine a scenario in which entire fictitious sub-fields of a literature emerge in which all of the citations are from AI-generated papers to other reciprocally citing AI-generated papers. It is actually hard to imagine a task more suited to LLM's current capability than perpetrating a large-scale version of the "Sokal hoax."[2]

Our paper makes several contributions to this emerging literature. First, we provide a concrete demonstration of how LLMs can be used to automate the generation of academic finance papers at scale. Second, we highlight the potential for systematic manipulation of traditional academic metrics through AI-powered paper generation. Finally, we argue for the development of new standards in research evaluation that are robust to these technological capabilities.

The field of empirical asset pricing provides an ideal laboratory for demonstrating these issues, as we identify robust return predictors through comprehensive data mining of accounting ratios and use AI to generate testable hypotheses explaining the findings. The results challenge our understanding of how scientific hypotheses are (or should be) produced and validated in the age of artificial intelligence. While we demonstrate this capability in finance, the implications extend far beyond: any field where researchers develop theoretical frameworks to explain empirical patterns—from biology to physics to social sciences—could be transformed by similar AI-powered approaches to hypothesis generation. As Lu et al. (2024) demonstrate in their development of the "AI Scientist" system, these capabilities are rapidly gener-

to related literature yield additional citations to our own work). If Google's algorithms counted the three distinct versions of each paper identified with a common title as distinct papers (and we could effortlessly generate distinct titles), then these minimum numbers would be 864 and 1,152, respectively!

[2]Physics professor Alan Sokal's (1996) unauthorized human subject test of the hypothesis that "a leading North American journal of cultural studies" would publish any "article liberally salted with nonsense if (a) it sounded good and (b) it flattered the editors' ideological preconceptions." His erudite sounding gobbledygook was accepted and published in short order.

alizing across scientific domains, suggesting our findings may foreshadow a broader transformation in how theoretical frameworks are developed across the sciences.

2 Methodology

This section outlines the procedures used to identify robust cross-sectional return predictors and subsequently generate manuscript-ready research outputs at scale through AI-driven methods.

2.1 Data-Driven Signal Construction and Filtering

Our methodology for identifying potential return predictors follows the broad framework established in recent literature on factor discovery and cross-sectional asset pricing (Yan and Zheng, 2017; Hou et al., 2020; Chen et al., 2024). We begin by assembling a comprehensive candidate set of firm-level signals from COMPUSTAT. This initial dataset comprises 31,460 potential predictors, each formed by combining accounting variables and their temporal differences. These signals are constructed to span a wide array of firm characteristics, ensuring a rich and diverse search space.

We then implement a series of data-quality and sufficiency filters to refine the candidate set. Table 1, Panel A reports the results. First, we eliminate redundant measures, where multiple combinations of underlying accounting items produce essentially identical metrics, reducing the universe to 29,315 unique signals. Second, we impose a minimum breadth requirement, retaining only those signals that have at least 30 stocks represented in each cross-section, ensuring that the resulting portfolios are sufficiently diversified. This criterion reduces the set to 25,852 signals.

Next, we confine the sample period to signals available through December 2023 and require at least 360 months of historical data to enable robust statistical inference, leaving 19,834 signals after the temporal restriction and 17,074 after enforcing

the longevity requirement. These 17,074 candidate signals serve as the initial input into our systematic validation stage.

2.2 Statistical Validation and Robustness Tests

We subject the remaining signals to a series of increasingly stringent validation tests designed to identify those that yield economically meaningful and statistically reliable patterns in the cross-section of stock returns. Table 1, Panel B, summarizes these tests.

First, we evaluate the predictive capacity of each candidate signal by sorting stocks into equal-weighted decile portfolios. Among the 17,074 candidate signals, 7,102 (approximately 41.6%) generate statistically significant return spreads at the 5% level. Narrowing further, we consider equal-weighted quintile portfolios; only 1,249 signals (7.3%) remain significant under this additional more restrictive sorting method.

We then implement more stringent portfolio construction criteria by using NYSE breakpoints for quintile formation, which helps mitigate potential biases associated with firm size and improves cross-sectional comparability. Under these conditions, 808 signals (4.7%) produce significant return spreads using equal-weighted portfolios, and 640 signals (3.7%) retain significance with value-weighted portfolios, which further reduces susceptibility to small-firm effects.

To account for known systematic risks, we next adjust each signal's returns using the Fama and French (2018) six-factor model. After controlling for these well-established risk factors, only 183 signals (1.1%) remain statistically significant, suggesting genuine incremental predictive power beyond standard factor benchmarks.

Finally, we subject the remaining 183 signals to the Novy-Marx and Velikov (2024) "Assaying Anomalies" protocol. This state-of-the-art methodology rigorously benchmarks each candidate predictor against the expansive "zoo" of published

anomalies from Chen and Zimmermann (2022), providing a transparent and standardized basis for gauging its relative performance. The protocol automatically produces detailed PDF "template reports" for each signal, including comprehensive statistical assessments and robustness checks. After reviewing these results and filtering out signals that fail to demonstrate robust performance relative to closely related anomalies, we are left with only 96 signals (0.6%) that survive all layers of validation. These high-quality signals and the associated pdf outputs from the "Assaying Anomalies" protocol form the foundation for generating full-length academic papers at scale.

Figure 2 displays a Sankey plot showing the mapping of accounting variables used to construct stock return predictor signals. It shows the complex network of relationships between accounting variables in our universe of signals. The diagram maps the flow from numerator variables on the left side to denominator variables on the right, with line thickness proportional to the frequency of each variable pairing. Orange connections represent difference-type signals, constructed as the difference between a COMPUSTAT variable and its lag, scaled by a lagged COMPUSTAT variable, while blue connections indicate ratio-type signals formed by dividing one COMPUSTAT variable by another. As we can observe in the figure, certain accounting items feature prominently in our signal construction process. Common stock (CSTK) and debt issuance (DLTIS) serve as primary numerator variables, creating numerous potential signals when paired with various denominator variables.

2.3 AI-Driven Paper Generation Pipeline

The final stage of the methodology leverages state-of-the-art large language models (LLMs) and automated text-processing scripts to produce fully formed academic manuscripts for each of the 96 validated return predictors starting with the PDF "template reports" from the "Assaying Anomalies" protocol.

2.3.1 Systematic Signal Naming and Classification

We first apply an AI-powered naming system to assign each validated signal a descriptive, academically credible name. Using GPT-3.5-turbo, we produce signal identifiers that integrate COMPUSTAT variable names and acronyms into an informative, non-generic label.[3] This step ensures that each signal is readily interpretable and easily distinguishable within the literature, even as we scale the production of papers across dozens of return predictors.

2.3.2 Content Generation and Structuring

We next employ a more advanced LLM (Claude 3.5-Sonnet) to generate the core textual content of each paper. The introduction, composed of roughly 1,100 words, is subdivided into four sections to ensure a balanced, academically coherent narrative:[4]

1. Motivation (200 words): Frames the research question within the broader asset pricing literature, discussing market efficiency, cross-sectional predictability, and recent developments in factor research.

2. Hypothesis Development (300 words): Proposes economic mechanisms justifying the signal's predictive power, citing relevant theoretical and empirical studies to maintain a scholarly tone and contextualize the new factor.

3. Results Summary (300 words): Presents key empirical findings, highlighting statistical significance, robustness checks, and comparisons to established anomalies.

4. Contribution (300 words): Places the proposed signal in relation to 3–4 closely related studies, articulating how the new evidence enhances our understanding

[3] Appendix A.1 contains our specific prompt.
[4] Appendix A.2 contains our specific prompt.

of systematic return drivers and contributes to ongoing debates in the literature.

All generated text adheres to a formal academic writing style and utilizes active voice. It carefully distinguishes correlation from causation, avoids unwarranted claims, and ensures appropriate application of tense to reflect established knowledge versus new findings. Citations are embedded using LaTeX-formatted references, and all writing conventions align with norms in leading finance journals.

The other added sections of each manuscript, including Data and Conclusion, are generated following similarly structured prompts.[5]

2.3.3 Document Assembly, Formatting, and Quality Assurance

The final assembly relies on custom scripts to incorporate the generated content into a standardized LaTeX template. We maintain a consistent document structure by programmatically inserting AI-generated sections into the appropriate manuscript components, preserving academic formatting standards and internal consistency.

We maintain dedicated, signal-specific .bib files for references, merging newly introduced citations into a base bibliography. This step ensures that each manuscript is properly referenced and that the supporting literature is consistently integrated into the text.

The final documents are compiled using a multi-pass LaTeX build process to ensure proper formatting, stable referencing, and a professional appearance. Automated cleaning procedures remove extraneous auxiliary files and streamline file management. The end product is a fully formed, academically styled PDF suitable for journal submission. This fully integrated pipeline—encompassing data-driven predictor identification, rigorous statistical validation, AI-based content creation, and

[5]Appendices A.3 and A.4 contain our specific prompts for the Data and Conclusion sections, respectively.

automated document preparation—demonstrates a scalable approach to generating and disseminating academic finance research papers.

3 Results

3.1 Scale and Efficiency of AI-Powered Research

The automated pipeline successfully generated three versions of 96 complete academic papers, each documenting a novel return predictor. Figure 1 provides an example of a generated paper. Table 2 in Appendix B documents the signals and the resulting paper names. The process is remarkably efficient - while the data mining, validation, and generation of the PDF "template reports" from the "Assaying Anomalies" protocol takes about a day of computation time, the final paper generation takes minutes. This represents a dramatic acceleration compared to traditional research paper development.

3.2 Quality of Generated Content

The AI-generated papers have several notable characteristics. First, the signal names and acronyms are descriptive and show creative naming convention. For example, one of the signals is the ratio of current assets (COMPUSTAT item ACT) to EBITDA (COMPUSTAT item EBITDA). The name picked by GPT 3.5-Turbo for this signal is "Operating Liquidity Margin". Similarly, taxes paid (TXPD) to operating income (AO) is termed "Tax Efficiency". The naming algorithm suggestd in the prompt attempts to capture the economic relationships represented by the accounting variables while avoiding generic terms like "ratio" or "difference."

Second, the generated introductions show remarkable resemblance to academic papers. Claude consistently:

- Identifies plausible sounding economic mechanisms linking the signal to returns
- Expertly summarizes the empirical results by highlighting tests on which the signals perform particularly well
- Integrates the findings with existing literature through (mostly) appropriate citations
- Develops testable hypotheses that align with the empirical results
- Positions each study's contribution within the broader literature

Third, the data sections provide clear, technically accurate descriptions of signal construction. These sections successfully translate COMPUSTAT variable codes into meaningful economic quantities while maintaining precise documentation of the measurement process.

Finally, the conclusions effectively synthesize the findings by paraphrazing the abstract.

3.3 Further Content Evaluation

We are currently working on a systematic evaluation of the AI-generated content quality across several dimensions. First, we plan to examine citation accuracy by cross-referencing all citations in the generated papers against academic databases to identify hallucinated references. Preliminary analysis suggests that while most citations to foundational papers in top finance journals are accurate, the LLM occasionally generates fictitious references when attempting to cite more specific or recent work. We are working on quantifying both the rate of citation hallucination and on analyzing patterns in when and how these hallucinations occur.

Second, we are working on evaluating the alignment between generated hypotheses and empirical results through several quantifiable metrics. For each paper, we

will manually extract the main hypotheses stated in the introduction and compare it against the key statistical findings from the "Assaying Anomalies" protocol. For example, if a hypothesis predicts stronger effects among small firms or during market downturns, we would verify whether these specific cross-sectional or time-series patterns actually appear in the data. We would also track whether the LLM correctly incorporates the magnitude and statistical significance of the main portfolio sorts and factor model results when describing the findings. Additionally, we would examine if key analytical choices like portfolio construction methods and control variables are consistently referenced across the introduction, methods, and results sections of each paper. We are also considering complementing this analysis with expert evaluation from finance professors to judge whether the theoretical mechanisms proposed by the LLM represent economically sensible explanations for the documented return patterns.

4 Discussion

Our demonstration of AI-powered academic paper generation has broad implications for the future of financial research and raises important questions about research integrity, validation, and the role of theory in empirical finance. We organize our discussion around three key themes: methodological implications, challenges for research integrity, and future directions.

4.1 Methodological Implications and Research Production

The successful generation of 96 complete academic papers demonstrates both the potential and risks of automated research production in finance. First, our results show that AI can now develop hypotheses at an unprecedented scale. This capability fundamentally changes how we might approach the relation between empirical find-

ings and hypothesis development. While traditional research typically starts with hypothesis development followed by empirical testing, AI enables rapid iteration between empirical discovery and theoretical justification.

Second, this automated approach could democratize research production by reducing barriers to entry. However, it simultaneously raises concerns about research quality and validation. The ability to quickly generate and test multiple hypotheses could accelerate the discovery of market inefficiencies, but might also lead to their faster elimination through increased trading activity as these findings become more widely disseminated.

4.2 Challenges for Research Integrity

The integration of AI into research production presents several critical challenges for maintaining research integrity. Most significantly, our pipeline exemplifies the risk of industrialized HARKing (Hypothesizing After Results are Known). Although individual instances of post-hoc theorizing might be difficult to detect or may even reflect valid scientific practice, the systematic generation of hundreds of papers through automated processes fundamentally challenges traditional notions of theoretical contribution. This risk is dramatically amplified by recent advances in computational power and the increasing availability of machine-readable financial data, which enable researchers to test millions of potential predictors almost instantaneously. When combined with automated hypothesis generation, these technological capabilities could exponentially increase the scope of data mining and p-hacking already documented in the empirical asset pricing literature (Harvey et al., 2016; **?**). While traditional p-hacking might involve researchers consciously selecting favorable specifications, AI systems can now systematically explore and rationalize vast numbers of potential relationships, generating plausible theoretical justifications for any statistically significant pattern.

The ability to generate convincing theoretical frameworks that seamlessly integrate with existing literature creates new forms of potential academic arbitrage. Our process naturally generates citations to both existing and occasionally imaginary literature. When scaled to hundreds or thousands of papers, this automated citation generation could artificially inflate citation counts and create citation networks that appear legitimate but lack substantive theoretical foundations. Given the accessibility of large language models and the strong publication incentives in academia, there is a high likelihood that some researchers are already exploiting these capabilities to enhance their citation counts and publication records. This concern is particularly acute as the sophistication of language models makes such artificially generated content increasingly difficult to detect.[6]

Moreover, the flood of AI-generated papers could overwhelm traditional peer review processes. Even when papers contain statistically significant findings and seemingly plausible theoretical justifications, determining genuine scientific contribution becomes increasingly challenging. This suggests we need new standards for evaluating research contributions in the age of AI, focusing perhaps more on novelty and practical relevance rather than just statistical significance and theoretical plausibility.

4.3 The Role of Prompt Engineering

The quality of AI-generated research content depends heavily on prompt engineering—the art and science of designing effective instructions for language models. Our experience showed us that slight modifications to prompts can produce vastly different narratives, theoretical rationales, and levels of academic rigor. Well-crafted prompts proved essential for maintaining consistent academic writing standards, ensuring appropriate citation practices, developing logically structured hypotheses, and avoid-

[6]See, for example, a study that was retracted due in part to an apparent use of Generative AI in the introduction here.

ing speculative or unsubstantiated claims. By refining these prompts—incorporating stronger guidelines for citation relevance, greater skepticism in hypothesis framing, and explicit instructions to maintain theoretical restraint—we can mitigate some of the risks associated with HARKing and hallucinations. Prompt engineering thus represents a crucial skill set that must evolve alongside LLM technologies to enable more reliable and disciplined research outputs.

4.4 A Step in an Evolving Landscape

This study represents an early exploration of AI's capabilities in producing academic finance research at scale. The tools and methods demonstrated here are in their infancy, and we expect significant advances in the coming years. Just as the last decade witnessed breakthroughs in computational finance, we anticipate rapid evolution in AI-powered research tools.

Current limitations—such as citation hallucination, ambiguous theoretical framing, and rote replication of established narratives—may become trivial to solve as LLMs become more context-aware and evidence-based. Future systems will likely incorporate automated fact-checking mechanisms, dynamic citation verification, integrated replication capabilities, and self-improving research validation. As these technologies mature, the boundary between human-generated and AI-assisted research may become increasingly fluid, necessitating new frameworks for understanding and evaluating scholarly contributions.

4.5 Future Directions and Recommendations

Moving forward, we identify several crucial areas for development that could help maintain research integrity in an AI-enabled environment. The first priority should be the development of enhanced validation systems. We need automated tools that can verify citations, ensure reference accuracy, and validate theoretical frameworks.

These systems should be capable of detecting circular reasoning, redundant theorizing, and hallucinated citations.

The academic community must establish new quality control mechanisms for AI-assisted research. Future iterations of research automation should incorporate built-in quality controls, including automated checking of theoretical consistency, detection of overlapping or redundant hypotheses, validation of citation networks, and integration with replication databases. These mechanisms would help ensure that AI-generated content maintains high academic standards while contributing meaningful insights to the literature.

The finance community also needs to develop new evaluation standards that reflect the realities of AI-enabled research production. These standards should place greater emphasis on out-of-sample validation and focus on practical implementation and economic significance. Economic stories justifying observed phenomena should be evaluated, at least in part, by the novel testable predictions they make beyond the primary findings they were designed to explain. Transparently reporting AI involvement in research production would also help readers better evaluate the methodological rigor and theoretical contributions of each paper.

Implementation of these recommendations would require significant coordination within the academic finance community, or mechanism that incentives their adoption. Such efforts are essential, however, for maintaining research integrity as AI capabilities continue to advance. The focus should shift from merely identifying potential problems to developing practical solutions that can be implemented across the field. This includes establishing standardized protocols for AI disclosure in research, creating shared databases for validation, and developing community-wide standards for evaluating AI-assisted research.

5 Conclusion

Our findings suggest that the introduction of AI into academic research production represents more than just a technological advancement—it has the potential to be a fundamental shift in how we generate and validate knowledge in finance. The ability to automate hypothesis generation challenges us to reconsider what constitutes a meaningful research contribution.

The questions posed here have no easy answers, but demand careful consideration as we enter an era where AI becomes an increasingly integral part of the research process. The future of financial research may depend less on our ability to generate hypotheses and more on our capacity to distinguish meaningful insights from statistically significant but theoretically hollow findings.

References

Christopher A Bail. Can generative ai improve social science? *Proceedings of the National Academy of Sciences*, 121(21):e2314021121, 2024.

Andrew Y. Chen and Chukwuma Dim. High-throughput asset pricing. *Working Paper*, 2024.

Andrew Y Chen and Tom Zimmermann. Open source cross-sectional asset pricing. *Critical Finance Review*, 11(2):207–264, 2022.

Andrew Y. Chen, Alejandro Lopez-Lira, and Tom Zimmermann. Does peer-reviewed research help predict stock returns? *Working Paper*, 2024.

Eugene F Fama and Kenneth R French. Choosing factors. *Journal of Financial Economics*, 128(2):234–252, 2018.

I. M. Harking. Operating liquidity margin and the cross section of stock returns. *Working Paper*, 2024a.

I. M. Harking. Inventory efficiency ratio and the cross section of stock returns. *Working Paper*, 2024aa.

I. M. Harking. Stock ownership contrast and the cross section of stock returns. *Working Paper*, 2024ab.

I. M. Harking. Equity weighted debt scale and the cross section of stock returns. *Working Paper*, 2024ac.

I. M. Harking. Equity debt differential and the cross section of stock returns. *Working Paper*, 2024ad.

I. M. Harking. Stock depreciation difference signal and the cross section of stock returns. *Working Paper*, 2024ae.

I. M. Harking. Stock depreciation gradient and the cross section of stock returns. *Working Paper*, 2024af.

I. M. Harking. Stock dividend relationship index and the cross section of stock returns. *Working Paper*, 2024ag.

I. M. Harking. Stock dividend impact and the cross section of stock returns. *Working Paper*, 2024ah.

I. M. Harking. Employees per share sensitivity and the cross section of stock returns. *Working Paper*, 2024ai.

I. M. Harking. Stock-gross profit contrast and the cross section of stock returns. *Working Paper*, 2024aj.

I. M. Harking. Shareholder capital efficiency difference and the cross section of stock returns. *Working Paper*, 2024ak.

I. M. Harking. Stock-intangible disparity and the cross section of stock returns. *Working Paper*, 2024al.

I. M. Harking. Stock inventory delta and the cross section of stock returns. *Working Paper*, 2024am.

I. M. Harking. Stock liability differential signal and the cross section of stock returns. *Working Paper*, 2024an.

I. M. Harking. Equity liability differential and the cross section of stock returns. *Working Paper*, 2024ao.

I. M. Harking. Stock-ppe scale signal and the cross section of stock returns. *Working Paper*, 2024aq.

I. M. Harking. Net asset utilization gap and the cross section of stock returns. *Working Paper*, 2024ar.

I. M. Harking. Stock and receivables relationship and the cross section of stock returns. *Working Paper*, 2024as.

I. M. Harking. Inventory adjusted cash flow and the cross section of stock returns. *Working Paper*, 2024at.

I. M. Harking. Stock sales impact and the cross section of stock returns. *Working Paper*, 2024au.

I. M. Harking. Stock equity imbalance scale and the cross section of stock returns. *Working Paper*, 2024av.

I. M. Harking. Tax-adjusted stock difference and the cross section of stock returns. *Working Paper*, 2024aw.

I. M. Harking. Stock-impact ratio and the cross section of stock returns. *Working Paper*, 2024ax.

I. M. Harking. Operating expense normalized common stock difference and the cross section of stock returns. *Working Paper*, 2024ay.

I. M. Harking. Stock-rental discrepancy signal and the cross section of stock returns. *Working Paper*, 2024az.

I. M. Harking. Intangibles-to-ebitda and the cross section of stock returns. *Working Paper*, 2024b.

I. M. Harking. Revenue efficiency factor and the cross section of stock returns. *Working Paper*, 2024ba.

I. M. Harking. Debt asset differential and the cross section of stock returns. *Working Paper*, 2024bb.

I. M. Harking. Debt issuance impact factor and the cross section of stock returns. *Working Paper*, 2024bc.

I. M. Harking. Debt issuance efficiency and the cross section of stock returns. *Working Paper*, 2024bd.

I. M. Harking. Debt surplus delta and the cross section of stock returns. *Working Paper*, 2024be.

I. M. Harking. Capital expenditure to long-term debt issuance differential and the cross section of stock returns. *Working Paper*, 2024bf.

I. M. Harking. Debt funding efficiency and the cross section of stock returns. *Working Paper*, 2024bg.

I. M. Harking. Debt-equity liquidity gap and the cross section of stock returns. *Working Paper*, 2024bh.

I. M. Harking. Equity-debt imbalance factor and the cross section of stock returns. *Working Paper*, 2024bi.

I. M. Harking. Capital debt depreciation delta and the cross section of stock returns. *Working Paper*, 2024bj.

I. M. Harking. Debt depreciation scale and the cross section of stock returns. *Working Paper*, 2024bk.

I. M. Harking. Debt capacity shift and the cross section of stock returns. *Working Paper*, 2024bl.

I. M. Harking. Debt issue impact on ebit and the cross section of stock returns. *Working Paper*, 2024bm.

I. M. Harking. Debt-issuance gross profit delta and the cross section of stock returns. *Working Paper*, 2024bn.

I. M. Harking. Debt capital gap and the cross section of stock returns. *Working Paper*, 2024bo.

I. M. Harking. Debt impact efficiency score and the cross section of stock returns. *Working Paper*, 2024bp.

I. M. Harking. Capital funding efficiency margin and the cross section of stock returns. *Working Paper*, 2024bq.

I. M. Harking. Debt-issuance-ppe scale offset and the cross section of stock returns. *Working Paper*, 2024br.

I. M. Harking. Debt impact factor and the cross section of stock returns. *Working Paper*, 2024bs.

I. M. Harking. Debt impact on sales growth and the cross section of stock returns. *Working Paper*, 2024bt.

I. M. Harking. Equity impact divergence and the cross section of stock returns. *Working Paper*, 2024bu.

I. M. Harking. Debt-efficiency score and the cross section of stock returns. *Working Paper*, 2024bv.

I. M. Harking. Rent-scaled debt emission deviation and the cross section of stock returns. *Working Paper*, 2024bw.

I. M. Harking. Property machinery nonop income discrepancy and the cross section of stock returns. *Working Paper*, 2024bx.

I. M. Harking. Asset financing impact and the cross section of stock returns. *Working Paper*, 2024by.

I. M. Harking. Profitable investment flow and the cross section of stock returns. *Working Paper*, 2024bz.

I. M. Harking. Net asset impact to depreciation and the cross section of stock returns. *Working Paper*, 2024c.

I. M. Harking. Capital scale nonop diff and the cross section of stock returns. *Working Paper*, 2024ca.

I. M. Harking. Inventory payment pressure margin and the cross section of stock returns. *Working Paper*, 2024cb.

I. M. Harking. Inventory efficiency operating factor and the cross section of stock returns. *Working Paper*, 2024cc.

I. M. Harking. Nonop liability contrast and the cross section of stock returns. *Working Paper*, 2024cd.

I. M. Harking. Equity-debt slant and the cross section of stock returns. *Working Paper*, 2024ce.

I. M. Harking. Asset utilization impact and the cross section of stock returns. *Working Paper*, 2024cf.

I. M. Harking. Net property plant and equipment to nonoperating income scale and the cross section. *Working Paper*, 2024cg.

I. M. Harking. Receipts impact and the cross section of stock returns. *Working Paper*, 2024ch.

I. M. Harking. Equity impact scale and the cross section of stock returns. *Working Paper*, 2024ci.

I. M. Harking. Operating efficiency margin and the cross section of stock returns. *Working Paper*, 2024cj.

I. M. Harking. Tax shield sensitivity factor and the cross section of stock returns. *Working Paper*, 2024ck.

I. M. Harking. Tax deprec profit impact and the cross section of stock returns. *Working Paper*, 2024cl.

I. M. Harking. Cash flow to equity and the cross section of stock returns. *Working Paper*, 2024cm.

I. M. Harking. Cash flow sustainability index and the cross section of stock returns. *Working Paper*, 2024cn.

I. M. Harking. Tax-effectiveness yield and the cross section of stock returns. *Working Paper*, 2024co.

I. M. Harking. Tax dividend efficiency score and the cross section of stock returns. *Working Paper*, 2024cp.

I. M. Harking. Tax dividend coverage metric and the cross section of stock returns. *Working Paper*, 2024cq.

I. M. Harking. Tax efficiency and the cross section of stock returns. *Working Paper*, 2024cr.

I. M. Harking. Growth impact efficiency metric and the cross section of stock returns. *Working Paper*, 2024d.

I. M. Harking. Profitable liquidity score and the cross section of stock returns. *Working Paper*, 2024e.

I. M. Harking. Cash earnings proportion and the cross section of stock returns. *Working Paper*, 2024f.

I. M. Harking. Cash profitability index and the cross section of stock returns. *Working Paper*, 2024g.

I. M. Harking. Efficiency of expense allocation and the cross section of stock returns. *Working Paper*, 2024h.

I. M. Harking. Asset income spread and the cross section of stock returns. *Working Paper*, 2024i.

I. M. Harking. Acquisitions efficiency ratio and the cross section of stock returns. *Working Paper*, 2024j.

I. M. Harking. Acquisition adjusted receivables current and the cross section of stock returns. *Working Paper*, 2024k.

I. M. Harking. Asset nonop impact and the cross section of stock returns. *Working Paper*, 2024l.

I. M. Harking. Equity to cash scale and the cross section of stock returns. *Working Paper*, 2024m.

I. M. Harking. Cash liquidity impact and the cross section of stock returns. *Working Paper*, 2024n.

I. M. Harking. Equity scale diff and the cross section of stock returns. *Working Paper*, 2024o.

I. M. Harking. Asset efficiency margin and the cross section of stock returns. *Working Paper*, 2024p.

I. M. Harking. Stock asset delta and the cross section of stock returns. *Working Paper*, 2024q.

I. M. Harking. Stock-to-asset spread and the cross section of stock returns. *Working Paper*, 2024r.

I. M. Harking. Equity dilution factor and the cross section of stock returns. *Working Paper*, 2024s.

I. M. Harking. Equity efficiency and the cross section of stock returns. *Working Paper*, 2024t.

I. M. Harking. Equity adjustment impact and the cross section of stock returns. *Working Paper*, 2024u.

I. M. Harking. Capital stock utilization delta and the cross section of stock returns. *Working Paper*, 2024v.

I. M. Harking. Stock investment efficiency signal and the cross section of stock returns. *Working Paper*, 2024w.

I. M. Harking. Net ownership stake and the cross section of stock returns. *Working Paper*, 2024x.

I. M. Harking. Equity share deviation and the cross section of stock returns. *Working Paper*, 2024y.

I. M. Harking. Stock cash differential and the cross section of stock returns. *Working Paper*, 2024z.

Campbell R. Harvey, Yan Liu, and Heqing Zhu. ... and the cross-section of expected returns. *Review of Financial Studies*, 29(1):5–68, 2016.

John J Horton. Large language models as simulated economic agents: What can we learn from homo silicus? *Working Paper*, 2023. MIT & NBER.

Kewei Hou, Chen Xue, and Lu Zhang. Replicating anomalies. *The Review of Financial Studies*, 33(5):2019–2133, 2020.

Norbert L Kerr. HARKing: Hypothesizing After the Results are Known. *Personality and Social Psychology Review*, 2(3):196–217, 1998.

Anton Korinek. Language models and cognitive automation for economic research. *NBER Working Paper*, (30957), 2023.

Weixin Liang, Yaohui Zhang, Zhengxuan Wu, Haley Lepp, Wenlong Ji, Xuandong Zhao, et al. Mapping the increasing use of llms in scientific papers. *arXiv preprint arXiv:2404.01268*, 2024.

Chris Lu, Cong Lu, Robert Tjurko Lange, Jakob Foerster, Jeff Clune, and David Ha. The AI scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint*, 2024. Equal Advising.

Benjamin S Manning, Kehang Zhu, and John J Horton. Automated social science: Language models as scientist and subjects. *SSRN Electronic Journal*, 2024.

Robert Novy-Marx and Mihail Velikov. Assaying anomalies. *Working Paper*, 2024.

Chenglei Si, Diyi Yang, and Tatsunori Hashimoto. Can llms generate novel research ideas? a large-scale human study with 100+ nlp researchers. *arXiv preprint arXiv:2409.04109*, 2024.

Alan D. Sokal. Transgressing the boundaries: Toward a transformative hermeneutics of quantum gravity. *Social Text*, (46/47):217–252, 1996.

Emma van Inwegen, Zanele Munyikwa, and John J Horton. Algorithmic writing assistance on jobseekers' resumes increases hires. *Working Paper*, 2023. MIT & NBER.

Xuemin Sterlin Yan and Lingling Zheng. Fundamental analysis and the cross-section of stock returns: A data mining approach. *Review of Financial Studies*, 30(4): 1382–1423, 2017.



Figure 1: Thumbnail grid of the PDF document.
This figure displays a combined thumbnail overview of a generated PDF document, showing all pages arranged in a 7x4 grid. Each page is represented as a scaled-down image, allowing for a quick visual reference to the content and layout of the entire document.

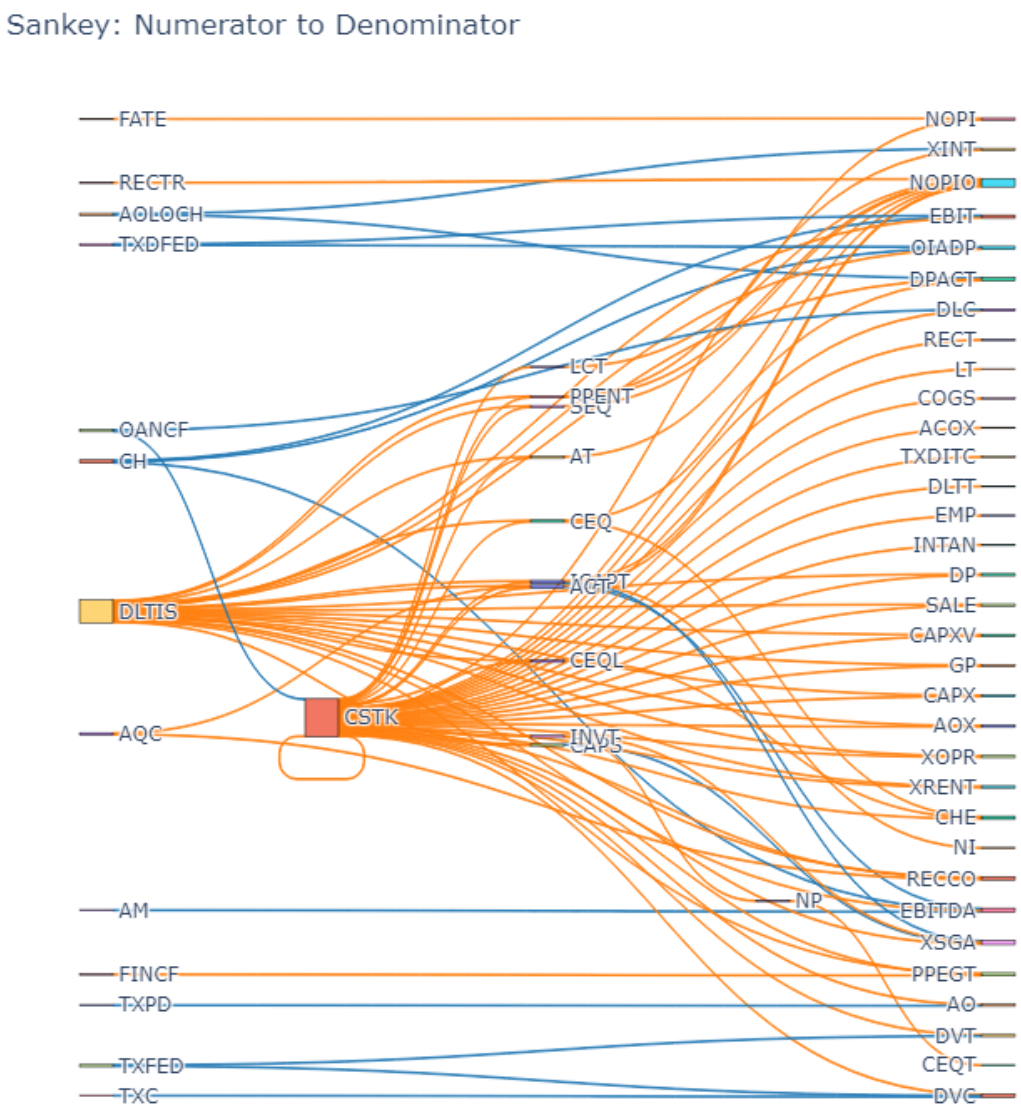


Figure 2: Sankey plot
This figure displays a Sankey plot showing the mapping of accounting variables used to construct stock return predictor signals. The figure illustrates the flow from numerator variables (left) to denominator variables (right). Orange connections represent difference-type signals (constructed as the difference of a COMPUSTAT variable and its lag, scaled by a lagged COMPUSTAT variable), while blue connections represent ratio-type signals (constructed as the ratio of two COMPUSTAT variables).

Table 1: Data mining summary
This table outlines the filtering process for the signals chosen to demonstrate the scale of paper generation possible with AI.

Filter	# of Signals	% of Filtered
Initial set	31460	
Exclude redundant signals	29315	
Require 30 stocks	25852	
Require data until 12/2023	19834	
Require 360 months	17074	100.0%
Panel B: Cumulative significance criteria		
+ $\lvert t_{\hat{r}_{\text{(decile, name, EW)}}} \rvert > 1.96$	7,102	41.6%
+ $\lvert t_{\hat{r}_{\text{(quintile, name, EW)}}} \rvert > 1.96$	1,249	7.3%
+ $\lvert t_{\hat{r}_{\text{(quintile, NYSE, EW)}}} \rvert > 1.96$	808	4.7%
+ $\lvert t_{\hat{r}_{\text{(quintile, NYSE, VW)}}} \rvert > 1.96$	640	3.7%
+ $\lvert t_{\hat{\alpha}_{\text{(quintile, NYSE, VW)}}} \rvert > 1.96$	183	1.1%
+ $\lvert t_{\text{Assay, Close Span}} \rvert > 1.96$	96	0.6%

A LLM Prompt Engineering

This appendix details the exact prompts used to generate the papers' content. Each prompt was designed to ensure consistent, high-quality academic writing while maintaining appropriate structure and theoretical development.

A.1 Signal Naming Prompt

The GPT-3.5-turbo prompt for generating signal names was:

Signal Naming Prompt

Create a descriptive and short name, as well as an acronym for a financial signal where the signal type is `'[signal_type]'`. Avoid using the words ratio and difference in the name. The acronym should only include capital letters.

For ratio-type signals, the prompt continued with:

Signal Naming Prompt

It is the ratio of `'[numer_full]'` to `'[denom_full]'`.

For difference-type signals:

Signal Naming Prompt

It is the difference in `'[numer_full]'` scaled by `'[denom_full]'`.

For signals involving negative transformations:

Signal Naming Prompt

The signal should be the negative of the computed value.

A.2 Introduction Generation Prompt

The primary prompt for generating paper introductions was:

Finance Paper Introduction Prompt

Write an introduction for a finance academic paper discussing the signal `'[signal_name]'` that predicts stock returns. Please follow these detailed guidelines:

- 1. Motivation (2 paragraphs, ∼200 words total):
  - Open with a broad statement about market efficiency or asset pricing.
  - Identify the specific gap or puzzle in the literature.
  - Use active voice and declarative statements.
- 2. Hypothesis Development (3 paragraphs, ∼300 words total):
  - Present economic mechanisms linking the signal to returns.
  - Draw on established theoretical frameworks.
  - Build logical arguments step by step.
  - Support each claim with citations to foundational papers in LaTeX format.
- 3. Results Summary (3 paragraphs, ∼300 words total):
  - Lead with the strongest statistical finding.
  - Present results in order of importance.
  - Use precise statistical language.
  - Include economic significance.
  - Mirror exactly the terminology used in the results section.
- 4. Contribution (3 paragraphs, ∼300 words total):
  - Position relative to 3-4 most closely related papers.
  - Cite papers from: *Journal of Finance, Journal of Financial Economics, Review of Financial Studies, Journal of Accounting Research, Journal of Accounting and Economics*.
  - Highlight methodological innovations.
  - Emphasize novel findings.
  - End with broader implications.
- Use active voice (e.g., "We find" instead of "It is found").
- Maintain formal academic tone.
- Include 2-3 citations per paragraph on average in LaTeX format.
- Use `\citep{AuthorsYear}` for parenthetical citations.
- Use present tense for established findings.
- Use past tense for your specific results.
- Avoid speculation beyond the data.
- Make clear distinctions between correlation and causation.

A.3 Data Section Generation Prompt

For signals constructed as ratios, the data section prompt was:

Data Section Prompt

The data section should include a description of the construction of the signal, '[signal_name]', which is constructed as the ratio of COMPUSTAT variable '[numer]' and COMPUSTAT variable '[denom]'.

Our study investigates the predictive power of a financial signal derived from accounting data for cross-sectional returns, focusing specifically on the ratio of current assets to earnings before interest, taxes, depreciation, and amortization (EBITDA). We obtain accounting and financial data from COMPUSTAT, covering firm-level observations for publicly traded companies. To construct our signal, we use COMPUSTAT's item ACT for current assets and item EBITDA for earnings.

Current assets (ACT) represent the firm's short-term assets, which are expected to be converted to cash or consumed within a year, including cash, receivables, and inventories. EBITDA, on the other hand, provides a measure of core operating performance by isolating operating income from non-operating expenses and tax effects.

The construction of the signal follows a straightforward ratio format, where we divide ACT by EBITDA for each firm in each year of our sample. This ratio captures the relative scale of a firm's liquid or short-term assets against its operational income, offering insight into how efficiently the firm utilizes current assets to generate earnings.

By focusing on this relationship, the signal aims to reflect aspects of liquidity management and operational efficiency in a manner that is both scalable and interpretable. We construct this ratio using end-of-fiscal-year values for both ACT and EBITDA to ensure consistency and comparability across firms and over time.

For difference-based signals, the prompt was modified to:

Data Section Prompt

The data section should include a description of the construction of the signal, `'[signal_name]'`, which is constructed as the difference of COMPUSTAT variable `'[numer]'` and its lag, scaled by lagged COMPUSTAT variable `'[denom]'`.

A.4 Conclusion Generation Prompt

The conclusion section was generated using the following prompt:

Conclusion Section Prompt

Write a conclusion for a financial research paper analyzing the signal `'[signal_name]'` in predicting stock returns. Summarize the key findings of the analysis, discussing the significance of the signal in terms of predictive power and practical implications. Conclude with suggestions for future research and limitations of this study. The conclusion should be based on the following abstract:

`'[abstract_text]'`

B Generated papers

Table 2: Signal Descriptions and References - Part 1

Numerator	Denominator	Signal	Title	Citation
ACT	EBITDA	ratio	Operating Liquidity Margin and the Cross Section of Stock Returns	Harking (2024a)
AM	EBITDA	ratio	Intangibles-to-EBITDA and the Cross Section of Stock Returns	Harking (2024b)
AOLOCH	DPACT	ratio	Net Asset Impact to Depreciation and the Cross Section of Stock Returns	Harking (2024c)
AOLOCH	XINT	ratio	Growth Impact Efficiency Metric and the Cross Section of Stock Returns	Harking (2024d)
CH	EBITDA	ratio	Profitable Liquidity Score and the Cross Section of Stock Returns	Harking (2024e)
CH	EBIT	ratio	Cash Earnings Proportion and the Cross Section of Stock Returns	Harking (2024f)
CH	OIADP	ratio	Cash Profitability Index and the Cross Section of Stock Returns	Harking (2024g)
CAPS	XSGA	ratio	Efficiency of Expense Allocation and the Cross Section of Stock Returns	Harking (2024h)
ACT	NOPIO	diff	Asset Income Spread and the Cross Section of Stock Returns	Harking (2024i)
AQC	ACT	diff	Acquisitions Efficiency Ratio and the Cross Section of Stock Returns	Harking (2024j)
AQC	RECCO	diff	Acquisition Adjusted Receivables Current and the Cross Section of Stock Returns	Harking (2024k)
AT	NOPIO	diff	Asset Nonop Impact and the Cross Section of Stock Returns	Harking (2024l)
CEQ	CHE	diff	Equity to Cash Scale and the Cross Section of Stock Returns	Harking (2024m)
CEQL	CHE	diff	Cash Liquidity Impact and the Cross Section of Stock Returns	Harking (2024n)
CEQ	NOPIO	diff	Equity Scale Diff and the Cross Section of Stock Returns	Harking (2024o)
CSTK	ACOX	diff	Asset Efficiency Margin and the Cross Section of Stock Returns	Harking (2024p)
CSTK	ACT	diff	Stock Asset Delta and the Cross Section of Stock Returns	Harking (2024q)
CSTK	AOX	diff	Stock-to-Asset Spread and the Cross Section of Stock Returns	Harking (2024r)
CSTK	AO	diff	Equity Dilution Factor and the Cross Section of Stock Returns	Harking (2024s)
CSTK	AT	diff	Equity Efficiency and the Cross Section of Stock Returns	Harking (2024t)
CSTK	CAPS	diff	Equity Adjustment Impact and the Cross Section of Stock Returns	Harking (2024u)
CSTK	CAPXV	diff	Capital Stock Utilization Delta and the Cross Section of Stock Returns	Harking (2024v)
CSTK	CAPX	diff	Stock Investment Efficiency Signal and the Cross Section of Stock Returns	Harking (2024w)
CSTK	CEQL	diff	Net Ownership Stake and the Cross Section of Stock Returns	Harking (2024x)
CSTK	CEQ	diff	Equity Share Deviation and the Cross Section of Stock Returns	Harking (2024y)

Table 2: Signal Descriptions and References - Part 2

Numerator	Denominator	Signal	Title	Citation
CSTK	CHE	diff	Stock Cash Differential and the Cross Section of Stock Returns	Harking (2024z)
CSTK	COGS	diff	Inventory Efficiency Ratio and the Cross Section of Stock Returns	Harking (2024aa)
CSTK	CSTK	diff	Stock Ownership Contrast and the Cross Section of Stock Returns	Harking (2024ab)
CSTK	DLC	diff	Equity Weighted Debt Scale and the Cross Section of Stock Returns	Harking (2024ac)
CSTK	DLTT	diff	Equity Debt Differential and the Cross Section of Stock Returns	Harking (2024ad)
CSTK	DPACT	diff	Stock Depreciation Difference Signal and the Cross Section of Stock Returns	Harking (2024ae)
CSTK	DP	diff	Stock Depreciation Gradient and the Cross Section of Stock Returns	Harking (2024af)
CSTK	DVC	diff	Stock Dividend Relationship Index and the Cross Section of Stock Returns	Harking (2024ag)
CSTK	DVT	diff	Stock Dividend Impact and the Cross Section of Stock Returns	Harking (2024ah)
CSTK	EMP	diff	Employees per Share Sensitivity and the Cross Section of Stock Returns	Harking (2024ai)
CSTK	GP	diff	Stock-Gross Profit Contrast and the Cross Section of Stock Returns	Harking (2024aj)
CSTK	ICAPT	diff	Shareholder Capital Efficiency Difference and the Cross Section of Stock Returns	Harking (2024ak)
CSTK	INTAN	diff	Stock-Intangible Disparity and the Cross Section of Stock Returns	Harking (2024al)
CSTK	INVT	diff	Stock Inventory Delta and the Cross Section of Stock Returns	Harking (2024am)
CSTK	LCT	diff	Stock Liability Differential Signal and the Cross Section of Stock Returns	Harking (2024an)
CSTK	LT	diff	Equity Liability Differential and the Cross Section of Stock Returns	Harking (2024ao)
CSTK	PPEGT	diff	Stock-PPE Scale Signal and the Cross Section of Stock Returns	Harking (2024aq)
CSTK	PPENT	diff	Net Asset Utilization Gap and the Cross Section of Stock Returns	Harking (2024ar)
CSTK	RECCO	diff	Stock and Receivables Relationship and the Cross Section of Stock Returns	Harking (2024as)
CSTK	RECT	diff	Inventory Adjusted Cash Flow and the Cross Section of Stock Returns	Harking (2024at)
CSTK	SALE	diff	Stock Sales Impact and the Cross Section of Stock Returns	Harking (2024au)
CSTK	SEQ	diff	Stock Equity Imbalance Scale and the Cross Section of Stock Returns	Harking (2024av)
CSTK	TXDITC	diff	Tax-Adjusted Stock Difference and the Cross Section of Stock Returns	Harking (2024aw)
CSTK	XINT	diff	Stock-Impact Ratio and the Cross Section of Stock Returns	Harking (2024ax)

Table 2: Signal Descriptions and References - Part 3

Numerator	Denominator	Signal	Title	Citation
CSTK	XOPR	diff	Operating Expense Normalized Common Stock Difference and the Cross Section of Stock Returns	Harking (2024ay)
CSTK	XRENT	diff	Stock-Rental Discrepancy Signal and the Cross Section of Stock Returns	Harking (2024az)
CSTK	XSGA	diff	Revenue Efficiency Factor and the Cross Section of Stock Returns	Harking (2024ba)
DLTIS	ACT	diff	Debt Asset Differential and the Cross Section of Stock Returns	Harking (2024bb)
DLTIS	AOX	diff	Debt Issuance Impact Factor and the Cross Section of Stock Returns	Harking (2024bc)
DLTIS	AT	diff	Debt Issuance Efficiency and the Cross Section of Stock Returns	Harking (2024bd)
DLTIS	CAPS	diff	Debt Surplus Delta and the Cross Section of Stock Returns	Harking (2024be)
DLTIS	CAPXV	diff	Capital Expenditure to Long-term Debt Issuance Differential and the Cross Section of Stock Returns	Harking (2024bf)
DLTIS	CAPX	diff	Debt Funding Efficiency and the Cross Section of Stock Returns	Harking (2024bg)
DLTIS	CEQL	diff	Debt-Equity Liquidity Gap and the Cross Section of Stock Returns	Harking (2024bh)
DLTIS	CEQ	diff	Equity-Debt Imbalance Factor and the Cross Section of Stock Returns	Harking (2024bi)
DLTIS	DPACT	diff	Capital Debt Depreciation Delta and the Cross Section of Stock Returns	Harking (2024bj)
DLTIS	DP	diff	Debt Depreciation Scale and the Cross Section of Stock Returns	Harking (2024bk)
DLTIS	EBITDA	diff	Debt Capacity Shift and the Cross Section of Stock Returns	Harking (2024bl)
DLTIS	EBIT	diff	Debt Issue Impact on EBIT and the Cross Section of Stock Returns	Harking (2024bm)
DLTIS	GP	diff	Debt-Issuance Gross Profit Delta and the Cross Section of Stock Returns	Harking (2024bn)
DLTIS	ICAPT	diff	Debt Capital Gap and the Cross Section of Stock Returns	Harking (2024bo)
DLTIS	OIADP	diff	Debt Impact Efficiency Score and the Cross Section of Stock Returns	Harking (2024bp)
DLTIS	PPEGT	diff	Capital Funding Efficiency Margin and the Cross Section of Stock Returns	Harking (2024bq)
DLTIS	PPENT	diff	Debt-Issuance-PPE Scale Offset and the Cross Section of Stock Returns	Harking (2024br)
DLTIS	RECCO	diff	Debt Impact Factor and the Cross Section of Stock Returns	Harking (2024bs)
DLTIS	SALE	diff	Debt Impact on Sales Growth and the Cross Section of Stock Returns	Harking (2024bt)
DLTIS	SEQ	diff	Equity Impact Divergence and the Cross Section of Stock Returns	Harking (2024bu)
DLTIS	XOPR	diff	Debt-Efficiency Score and the Cross Section of Stock Returns	Harking (2024bv)

Table 2: Signal Descriptions and References - Part 4

Numerator	Denominator	Signal	Title	Citation
DLTIS	XRENT	diff	Rent-scaled Debt Emission Deviation and the Cross Section of Stock Returns	Harking (2024bw)
FATE	NOPI	diff	Property Machinery Nonop Income Discrepancy and the Cross Section of Stock Returns	Harking (2024bx)
FINCF	PPEGT	diff	Asset Financing Impact and the Cross Section of Stock Returns	Harking (2024by)
ICAPT	NI	diff	Profitable Investment Flow and the Cross Section of Stock Returns	Harking (2024bz)
ICAPT	NOPIO	diff	Capital Scale Nonop Diff and the Cross Section of Stock Returns	Harking (2024ca)
INVT	NP	diff	Inventory Payment Pressure Margin and the Cross Section of Stock Returns	Harking (2024cb)
INVT	XSGA	diff	Inventory Efficiency Operating Factor and the Cross Section of Stock Returns	Harking (2024cc)
LCT	NOPIO	diff	Nonop Liability Contrast and the Cross Section of Stock Returns	Harking (2024cd)
NP	CEQT	diff	Equity-Debt Slant and the Cross Section of Stock Returns	Harking (2024ce)
PPENT	NOPIO	diff	Asset Utilization Impact and the Cross Section of Stock Returns	Harking (2024cf)
PPENT	NOPI	diff	Net Property Plant and Equipment to Nonoperating Income Scale and the Cross Section	Harking (2024cg)
RECTR	NOPIO	diff	Receipts Impact and the Cross Section of Stock Returns	Harking (2024ch)
SEQ	NOPIO	diff	Equity Impact Scale and the Cross Section of Stock Returns	Harking (2024ci)
ICAPT	XSGA	ratio	Operating Efficiency Margin and the Cross Section of Stock Returns	Harking (2024cj)
TXDFED	EBIT	ratio	Tax Shield Sensitivity Factor and the Cross Section of Stock Returns	Harking (2024ck)
TXDFED	OIADP	ratio	Tax Deprec Profit Impact and the Cross Section of Stock Returns	Harking (2024cl)
OANCF	CSTK	ratio	Cash Flow to Equity and the Cross Section of Stock Returns	Harking (2024cm)
OANCF	DLC	ratio	Cash Flow Sustainability Index and the Cross Section of Stock Returns	Harking (2024cn)
TXC	DVC	ratio	Tax-Effectiveness Yield and the Cross Section of Stock Returns	Harking (2024co)
TXFED	DVC	ratio	Tax Dividend Efficiency Score and the Cross Section of Stock Returns	Harking (2024cp)
TXFED	DVT	ratio	Tax Dividend Coverage Metric and the Cross Section of Stock Returns	Harking (2024cq)
TXPD	AO	ratio	Tax Efficiency and the Cross Section of Stock Returns	Harking (2024cr)