


The Cost of Expertise

Humans versus AI-Enabled Automated Agents Using a Team of AI Agents to Streamline Risk Analysis of Complex Investment Portfolios

Mario Pardo

DOI: <https://doi.org/10.66241/5kbym>



Abstract

Ryse, Inc., a portfolio management software company founded in 2006, provides advanced AI-based tools to help clients analyze, simulate, and model complex portfolios across various asset classes. The platform supports financial analysis, risk management, and scenario simulation, empowering users to make informed investment decisions and optimize their strategies with the help

of artificial intelligence. The company's journey began with the Utah Retirement Systems pension fund, which manages over \$45 billion in assets today (up from \$20 billion in 2006). While this partnership marked a significant milestone, it also revealed major development needs. The platform's ability to generate more than 50,000 analyses across asset classes created an overwhelming workload for the client's investment team. Ryse had a choice—either redesign the platform and become an actionable tool that supported the investment management team instead of overburdening it, or hire an army of software support staff to continuously assist and manage client operations. This case study chronicles the steps—some carefully planned, others taken by chance—that Ryse followed to address these challenges. These steps ultimately culminated in the integration of AI agent technology, a promising advancement that transformed the platform's ability to automate complex processes. The adoption of AI agents was both a product of deliberate innovation and fortuitous developments, positioning Ryse to

offer a more efficient, cost effective, and powerful solution for portfolio management.

With the introduction of AI agents, Ryse has maintained its current service offerings while rapidly adding new features, such as news search, market research, key market event summaries, and earnings transcript interpretation. These tools leverage AI agent technology and provide essential support for regular investors who may have limited budgets, while also giving CIOs the ability to scenario test their portfolios. Through onboarding and testing the AI agent technology with existing clients and prospects, Ryse has successfully enhanced the overall service experience for its customer base.

Why did Ryse adopt AI Agent Technology?

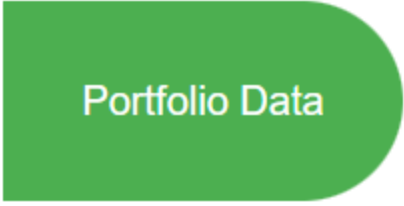
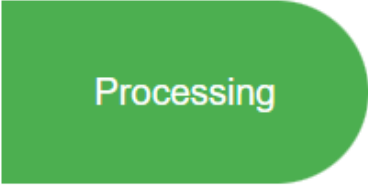
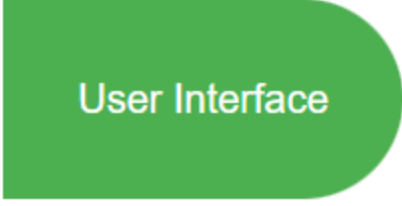
Mario Pardo, a founding member of Ryse and a graduate of Vanderbilt University as well as MIT's Management of Technology (MOT) Program— the first joint program between the Sloan School of Management and the School of Engineering — faced a significant decision. Should he pursue a traditional business model by engaging in well-known venture funding rounds available to companies of this size? A multi-billion dollar New York private equity firm approached Ryse with an offer to adopt a conventional model that involved hiring additional staff to assist clients with their data needs, bringing on more developers to add features, and shifting the focus toward an enterprise business model. However, this approach would ultimately compromise Pardo's vision for the company. Today, it is understood that the decision to concentrate on developing advanced computational tools, such as AI agents to assist users in interpreting data and integrating these tools into their decision-making processes, was not easy or straightforward. It has been a difficult and challenging path, fraught with many growing pains. Ultimately, time will tell if the market embraces this direction.

One major stepping stone for the company occurred back in 2016, before the ChatGPT era, when the newly developed “embeddings methodology” from Stanford and Google became public and showed promise. On the company's 10th anniversary, after successfully applying the tool with several clients, including pension funds, hedge funds, and family offices, Ryse leveraged its Harvard connections. A faculty member formerly of the Kennedy School of Government

advised Pardo to offer Ryse's technology to the Social Security Fund (SSF) of an East African nation through the school's highly regarded African initiative. The SSF proved to be an ideal candidate as it could serve as a testing ground for developing the kinds of AI tools needed to support the investments department without incurring the additional costs that a traditional risk service provider would demand.

SSF, a \$6B institution with investments across all asset classes and with focus on East Africa, connected with Ryse. The bidding process included all the world's major service providers in risk and portfolio management, with Ryse, Inc. competing alongside them. Ryse was no stranger to such competition, having first faced similar industry players when it won the mandate to provide services to Utah Retirement Systems in 2006.

However, given its startup status, Ryse recognized the need to provide solutions to the challenges faced by institutional managers. The chart (Figure 1) illustrates Ryse's key operational features. While institutional providers share similar features, they differ in the level of detail provided, pricing tools, asset class coverage, and more. One commonality among most portfolio and risk management platforms is that they tend to add more tasks to an already busy investments team. Pardo was well aware of this, having worked firsthand with various accounts at Ryse, Inc. for ten years.

Ryse, Inc.		
Initial Functional Service Offering		
		
<ul style="list-style-type: none"> • Access to custodian and market data through different APIs offered by the market, i.e. Bloomberg, etc. • Excel templates used to upload information not specified by the traditional custody services, i.e. simulations, etc. • SQL templates to access other data sources 	<ul style="list-style-type: none"> • Dotnet, C# application to process all portfolio positions • Risk proxy and Risk Transparency availability • Scenario Analyses (linear based) 	<ul style="list-style-type: none"> • Front end application that allows users to navigate through all layers of portfolio • Power-point/pdf Generator for presentations across any portfolio level
<p><i>Figure 1: Ryse’s Original Functional Service Offering</i></p> <p><i>Source: Ryse, Inc.</i></p>		

Consequently, Ryse approached the SSF’s bidding process with an innovative idea: leveraging embedding technology to explain the underlying risk factors driving both performance and volatility in their portfolio holdings. Ultimately, Ryse won the bidding process, and SSF has been collaborating with the company since 2016. The combination of embeddings, risk factor analytics, and risk profiling has helped analysts in SSF’s investments department better understand the numerous analyses and tables generated by Ryse’s platform over the past seven years.

The fact that Ryse's solution did not lead to additional labor costs in the investment team attracted the attention of SSF's Chief Investment Officer (CIO), a graduate of Harvard Business School's Executive MBA program. He played a pivotal role in facilitating Ryse's implementation of this new technology and ensuring ongoing access to their reports month after month. By 2022, SSF became the first institution in Africa to incorporate ChatGPT into its portfolio management, two years before AI's adoption in the US and other developed markets, significantly enhancing the fund's analytical capabilities. As of June 2024, the CIO was appointed as Deputy Managing Director.

Ryse, Inc.'s Product Evolution

Since its founding in 2006, Ryse, Inc. has been at the forefront of providing portfolio and risk analytics to financial institutions across the U.S., Latin America, Asia, and Africa. The company has cultivated a diverse roster of clients, including pension funds, hedge funds, and strategic consultants, each engaging with Ryse for an average of three to five years.

The evolution of Ryse's product began with an idea rooted in the founder's MIT thesis^[1] on causality, which sought to identify factors explaining the relationship between inflation in the U.S. and Latin America through the methodologies of causality and co-integration developed by C.W.J. Granger^[2]. This foundational concept laid the groundwork for the development of a product that would leverage an SQL database via a .NET application, integrating financial algorithms to identify the risk factors associated with investment portfolios. The technology was designed to empower clients to pinpoint the sources of risk within their portfolios at multiple levels. Additionally, the platform incorporated standard market computations, including performance metrics, Value at Risk (VaR), duration, convexity, and various drawdown analyses. Ryse's success with the Utah Retirement Systems (URS) pension fund paved the way for partnerships with other notable clients, such as the University of Connecticut endowment, and Strategy@ from Price Waterhouse, among other asset managers. This growing interest also attracted the attention of a New York-based private equity firm.

Despite these promising developments, Ryse faced a significant challenge to its growth—complexity. The company was forced to deploy analysts in order to assist

clients in interpreting and explaining what would soon be referred to as the “scary charts.” The complexity of data generated by the URS portfolio alone was staggering, with 92 sub-portfolios producing between 30,000 and 45,000 analyses each month on average. This volume of output required thorough analysis and clear communication of key findings to clients’ monthly investment committees. Only an experienced risk manager, or a power user, would be able to successfully navigate the universe of charts and tables generated by the platform and select the key analyses, every month. The sheer scale of this work posed a considerable barrier to the potential adoption of Ryse’s solutions by multiple institutional clients, as investment professionals untrained in the platform often struggled to identify the critical aspects relevant to understanding portfolio risk and performance.

At Ryse, the team focused on developing two key aspects of their analytics tool. First, they enhanced its capability to handle both linear and non-linear analytics, setting the tool apart from traditional solutions that rely solely on linear modeling. Second, they aimed to create a powerful graphical user interface (GUI) that allowed users to access analyses at every level of their portfolio with just a few clicks. As financial engineers, the team had their work cut out for them. The implementation of linear and non-linear analytics necessitated the incorporation of new pricing models that could more accurately reflect market behavior. This decision was crucial as it enabled the tool to capture significant market events, such as the 2008 credit crisis and the high inflation of the 1970s. The principles underlying linear versus non-linear analytics are based on the widely held belief that asset behavior remains consistent across different market conditions. Many financial models rely on the premise of normality to evaluate asset performance over time.

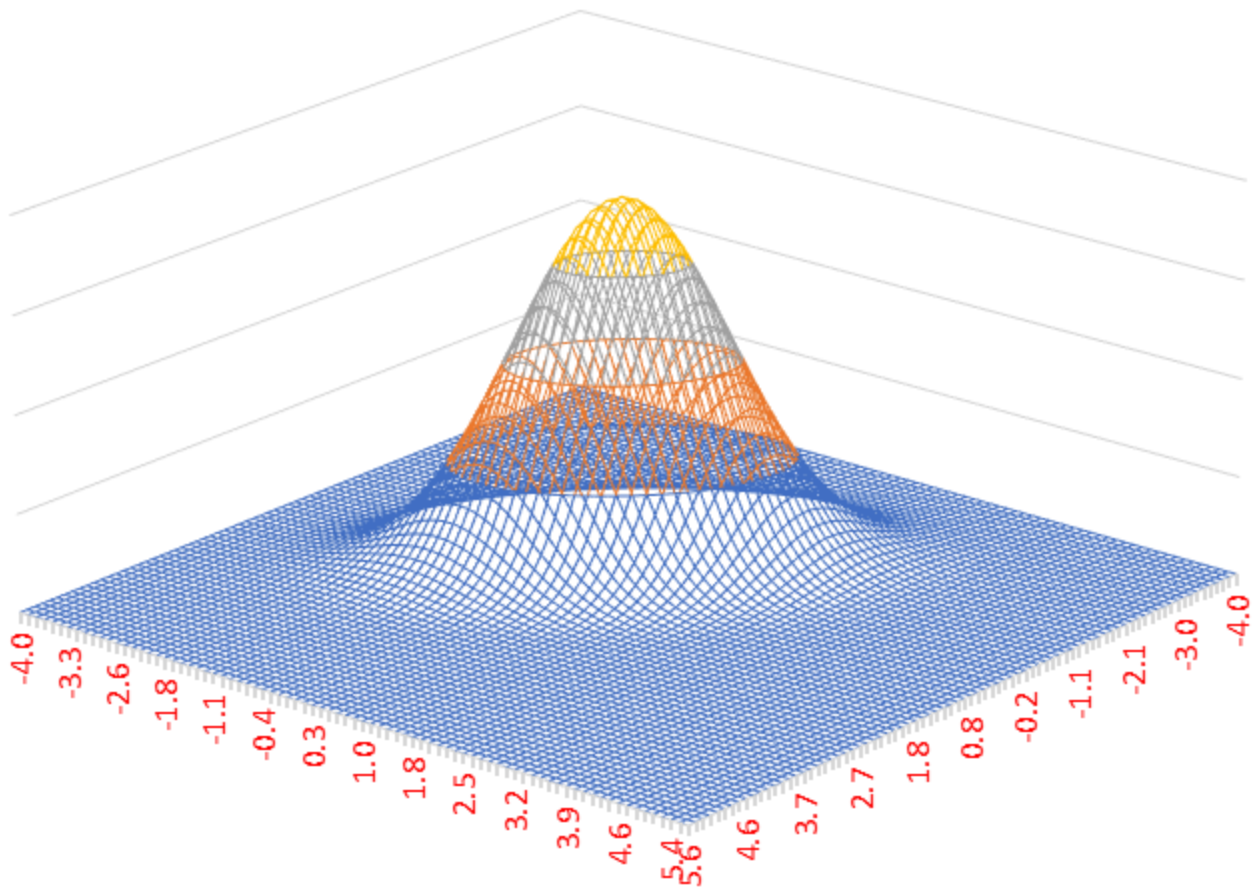
Having experienced the limitations of the Ryse cointegration methodology in predicting volatility during the 2008 credit crisis, the team chose to confront this challenge directly. They took the advice of Lisa Polsky, former Chief Risk Officer of Morgan Stanley and a member of the Risk Management Hall of Fame for her many contributions to the field. She advised Ryse to start over by developing a new method that could capture the non-linear aspects of risk. Pardo began working on this and subsequently tested its application with East African bonds. Emerging market bonds tend to exhibit non-normal behavior even in “normal

markets." With access to data from Ugandan, Kenyan, Tanzanian, and Rwandan bond market datasets through their existing relationships, the team had the advantage of assessing the behavior of these bonds and evaluating the method's effectiveness. After extensive testing and observation of the method's application, the team engaged in a challenging technological endeavor to integrate this feature. Since 2019, Ryse has been an active participant in the Industry Projects in Analytics & Operations Research course at Columbia University's Operations Research Department. Throughout this time, Ryse has delivered high-quality projects in generative artificial intelligence, serverless computing, scenario analysis, and portfolio attribution, with a focus on market risk assessments for investment portfolios. These projects are known for their relevance, practical applications, and timeliness. Students consistently commend Ryse's projects in their 360 reviews. Lastly, they conducted tests at a multibillion-dollar family office, demonstrating the method's efficacy.

Normality vs Non-Normality


To illustrate its impact on an AI-based platform, the following presents a simple application of the method on a portfolio of two assets. The chart below depicts the joint sample space for this portfolio, assuming its behavior follows a Gaussian or "normal" distribution. As indicated in the chart (Figure 2), most observations of the portfolio's joint behavior cluster around the mean, extending out to three to four standard deviations from it. Outside this range, the model shows no observations. This Gaussian model is widely used and serves as the foundation for various analyses, including correlation and Value at Risk (VaR) amongst others. However, the increasing integration of markets and the accessibility of alternative assets—such as private equity, real estate, and hedge funds—continuously poses challenges for this traditional model, as alternative investments often exhibit non-linear payout structures that do not adequately capture tail events.

Figure 2: Bi-variate Joint Sample Space for 2-asset portfolio (Assuming Normality)



To address this, Ryse developed a model based on the concept of dependency. Although this is not the only solution, it was found to be the most useful for interactions with AI modeling. Focusing on dependency enabled the exploration of a wider range of possibilities that effectively capture tail events, as shown in the chart below. In this case, it can be observed that most of the cases appear on the tails, and not with the same probability. This represents one of the many possible states of the bi-variate portfolio, highlighting that reliance on traditional metrics based on normality will lead to an inaccurate capture of risk.

Figure 3: Non Linear Distribution (Frank Copula)



The primary advantage of utilizing dependency over correlation is its capacity to encompass complex relationships between variables. While correlation is limited to linear relationships, dependency captures a broader spectrum of interactions, offering a more comprehensive understanding of portfolio behavior.

To achieve this, additional levels of analytics needed to be developed to accurately reflect the complex non-linear relationships inherent in the data. This exploration led the team to Sklar's Theorem^[3] (1959), a foundational result in copula theory. Sklar demonstrated that for any multivariate distribution, its joint distribution can be separated into the marginal distributions of the individual variables, and there is a copula function that captures the dependence structure between these variables. This separation allowed the Ryse Model firstly to analyze the marginal distributions (which might follow non-normal or heavy-tailed distributions) independently, and secondly to capture the dependency structure (even non-linear and tail dependence) through the copula, which could then be adjusted to different scenarios or stress tests. Next, the team adopted the concept of Spearman's correlation, recognizing the necessity to recompute correlations

between risk factors due to shifting dependencies during stress periods. Spearman's rank correlation is particularly valuable in this context, as it does not rely on linearity, thereby better reflecting how relationships can change under non-standard conditions. For example, assets may display low or moderate correlation during normal market conditions; however, during a financial crisis, their correlations can spike—a phenomenon often referred to as correlation breakdown or correlation clustering^[4] is particularly valuable in this context, as it does not rely on linearity, thereby better reflecting how relationships can change under non-standard conditions.

This newly found flexibility enabled the team to leverage copulas, which are particularly useful for scenario-based analytics. Copulas allowed them to model various dependency structures observed during extreme events without being constrained by normality assumptions. For instance, using copulas, they could model strong lower-tail dependence, critical for stress-testing financial portfolios during market crashes. Additionally, copulas facilitated a focus on scenario sensitivity, enabling the tool to recompute how dependencies shift under new distributions or stress scenarios, thus providing the tool—and the AI agents—with greater insight into real risk exposure beyond the average behavior captured previously. The next step in their development cycle was to use Principal Component Analysis (PCA) to reduce the dimensionality of their system, which the tool needed to identify the key risk factors driving dependency more effectively. This approach represented these factors as uncorrelated components, highlighting the most sensitive parts of the system under various conditions—information that was more relevant for the AI agents. Once PCA decomposed the underlying risk factors, these components were mapped back to the copula framework, recalibrating dependencies for different scenarios. This process involved recomputing correlations within this lower-dimensional space using copulas, ensuring that the scenarios reflected nonlinear dependencies that dynamically shifted based on the underlying risk structure. The last step encompassed the generation of Monte Carlo simulations based on the dependencies captured through Spearman correlations. In summary, they enhanced the computational efficiency of their simulations by reducing the number of variables that needed to be simulated and focusing on the primary

drivers of risk. All this new information became critical data points for the active AI agents.

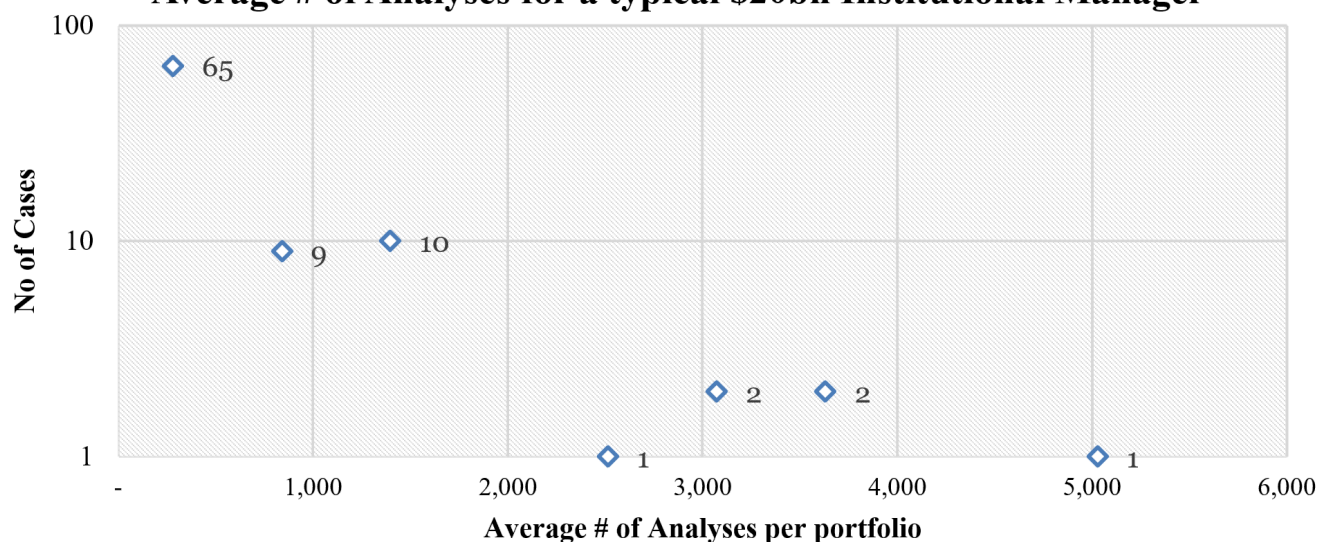
Complexity

While the addition of advanced analytics significantly improved the platform's capabilities and enhanced users' understanding of their portfolios, it did not resolve the underlying issue of complexity. Users would still require extensive support from data scientists, engineers, IT developers, and other professionals to effectively navigate and utilize the platform.

Based on their experience with a former client who managed a \$20 billion portfolio across various asset classes, the complexity of analysis becomes evident. On average, a client this size must be prepared to comprehend 65 different portfolios, each requiring an average of 279 analyses. Additionally, their findings indicated that between 20 to 30 portfolios would contain 800 to 1,400 analyses each, while 5 to 10 portfolios could have analysis counts ranging from 2,500 to 5,000. Altogether, this could result in over 50,000 different analyses for such a \$20 billion portfolio—an unmanageable task for any single analyst and a formidable challenge for even a team of portfolio analysts (See chart below, figure 4).

*Figure 4: No of Analyses generated for a typical \$20bn Institutional Manager,
Source: Ryse, Inc.*

Average # of Analyses for a typical \$20bn Institutional Manager



This complexity has significant and far-reaching implications for how industry players operate. Larger, established institutions often deploy extensive teams of engineers, data scientists, risk analysts and financial experts to assist their clients in evaluating intricate portfolios. Consequently, these firms are compelled to charge exorbitant fees for risk and portfolio management services, further widening the disparity in access to such services. As a result, only well-established asset managers can afford these comprehensive analytical tools, limiting access for smaller investors. This restricts smaller investors from gaining the same market insights that larger institutions enjoy, perpetuating inequality in the investment landscape. In response to these challenges, Ryse successfully deployed the embeddings methodology, allowing users to actively search through available analytics via a web interface. This method enabled users to quickly access specific analyses by typing their requirements into an input box, with the tool directing them to the corresponding analysis.

The Ryse AI Agents

In 2020, Ryse unveiled its first AI agent. An agent is traditionally defined as a self-contained software process that runs concurrently, maintains its own state, and communicates with other agents through message passing. This concept is considered a natural evolution of the object-oriented concurrent programming paradigm^[5]. Agents^[6] communicate with each other by sending messages using a specific language designed for that purpose. They can be simple, like basic

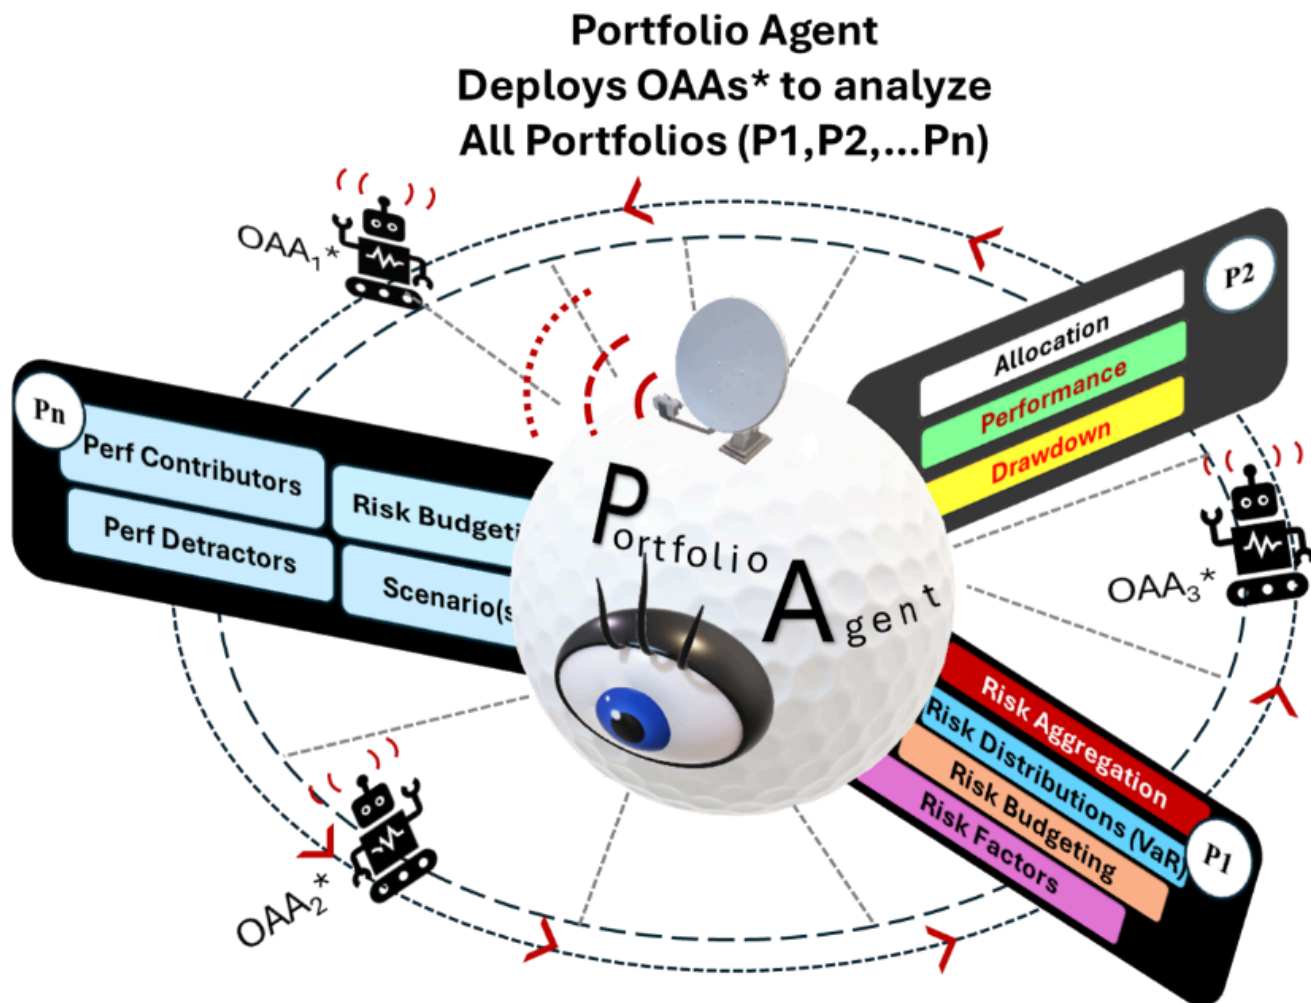
functions in a program, but they are usually more complex and have some form of ongoing control or management. An AI agent is then a software entity that uses artificial intelligence (AI) techniques to perform tasks autonomously. The agents themselves can be programmed to make decisions or interact with the data and, if possible, mimic human behavior.

The actions of the agent could be understood by considering its choices and obligations, treating it as if it had its own mind. The attitudes that represent an agent fall into two main types: information attitudes, which are related to what the agent knows or understands about its environment—such as its knowledge and awareness of its surroundings—and pro-attitudes, which guide the agent's behavior and decision-making, including its choices and obligations. To effectively represent an agent, it must have at least one information attitude and one pro-attitude.

To design the agents, the team created a series of software components with data flowing through them, allowing the agents to exert control over the different components. The chart below maps out the major interactions between all the modules, which are then used for overall portfolio processing. Through each interaction, the portfolio agent learns first through the "One Analysis Agent (OAA)." Witnessing the success of the first agent, they developed the "Ryse Portfolio Agent (PA)." These agents were initially tested with the East African client, SSF, where they were welcomed as valuable support to the company's investment department. In fact, SSF became the first institutional manager in all of East Africa to integrate AI into its portfolio of investments.

These two agents serve different objectives but are intimately connected. The OA Agent's role is to analyze all the data generated across all the portfolios by the back-end tool, while the PA Agent, which ultimately controls the OAAs, reviews the different areas of the portfolio and lastly feeds a chatbot that allows users to interact with and ask questions about their portfolio across its various layers. Having successfully deployed their first two agents, three more AI agents were added to the tool, a news research agent, an editing agent, and a design agent, all of which were successfully tested during the onboarding process of a major U.S. asset manager.

Figure 5: Ryse AI Agents, OAA (One Analysis Agent), Portfolio Agent that controls the OAAs



Ryse has gained a competitive edge in developing AI agents that communicate with one another, perform different tasks, and actively support portfolio managers. The Ryse AI agents are designed to process both the information generated by the platform and market data. Since they integrate these two sources, their behavior must be carefully controlled by senior agents to ensure the accuracy of analytics and to prevent potential hallucinations. Powered by AI-Large Language Models (AI-LLMs), these agents learn from experience and take actions to achieve specific goals without human intervention. Some of the key characteristics of Ryse’s AI agents include autonomy, perception, the ability to leverage financial models to learn from and adapt to data, and decision-making. By design, all of the Ryse agents are developed with the goal of providing an output that is either utilized by a superior agent or the end user. The multiple

analyses generated by the back-end application are transformed into JSON (JavaScript Object Notation), a lightweight data-interchange format that is easy for agents to parse and evaluate. Users interact with the agents through a series of chatbots integrated into the platform, which allow for continuous engagement. As users interact with these agents, the agents themselves adapt to the users' requests autonomously.

To achieve these results, the team embarked on a painstaking optimization process, addressing one analysis at a time. They incorporated different data characteristics that ensured compatibility with standard AI-LLM models. This required reengineering and, in many cases, updating old pricing and risk engines. The revamped risk engine now includes tags across all computations, allowing the autonomous agents to efficiently navigate the analyses, saving time and ensuring accuracy when interacting with users. The improved engines also addressed several key issues, such as simplifying complex "scary charts," reducing data size while improving processing time, and enhancing user interaction.

As a result, the portfolio agents can process a \$20 billion portfolio with over 50,000 analyses in under five minutes, depending on the AI-LLM model used and the platform of choice, such as OpenAI, Llama, Gemma, or the Groq platform, which leverages LPUs (Language Processing Units) to achieve faster response times. In fact, Groq's architecture is optimized for parallel processing, making it well-suited for execution on LPUs. At the core of their platform is a unique AI engine built to process massive datasets in parallel, which allows Groq to execute queries much faster (up to 100x) than traditional GPU-based query engines, especially for large datasets. With AI-supported tools, user interaction becomes more natural, leading to a deeper understanding of the entire analysis in real time. These advancements, when fully implemented, are expected to have a significant impact on helping achieve the goal of providing retail investors with the ability to understand their portfolios on par with large institutional investors at no cost.

The New Way Forward

The team had the opportunity to demonstrate their tool to the former Chief Risk Officer for Global Private Banking at a leading European bank. He noted that the

work completed by the agents in 30 minutes was equivalent to what his team of engineers at this institution could accomplish in one week using their traditional risk management platform. During his tenure, the bank used a prominent third-party risk and portfolio management system. Following this success, the Ryse team began marketing their capabilities, and in the spring of 2024, they initiated a successful onboarding process with one of the world's largest and most reputable asset managers.

Ryse's ultimate objective is to empower retail investors with the same level of analytical capabilities traditionally reserved for institutional investors—at no cost. The platform will allow product and portfolio managers from established financial institutions to offer their products for analysis on their websites, supported by Ryse's AI agents. With simple drag-and-drop functionality, retail investors will be able to access the same tools as family offices, endowments, foundations, pension funds, and other institutional players.

This paper has explored the core concepts and challenges related to the theory and application of intelligent agents as applied to portfolio management. The article covered a broad spectrum of topics on AI, explaining what an agent is, how they are designed, and what their capabilities are. It also discussed the architecture used to enable the agents to perform their functions. Lastly, and perhaps more importantly, it examined how this application benefits from agent-based solutions and the transformation this technology could bring to the overall industry. Agents will ultimately empower smaller investors, both at the institutional and retail levels, to access the same kinds of capabilities that larger institutional investors and the wealthiest family offices have. In other words, they will be instrumental in leveling the playing field and closing the widening disparity in access to such services. The importance of AI agents is underscored by the growing consensus in both academia and industry that intelligent agents will be crucial as computing systems become more distributed, interconnected, and open.

References

1. Rodriguez, L. A. (1997). *Testing for Causality of Inflation in Latin America from the U.S. Economy* [Thesis, Massachusetts Institute of Technology, Sloan School of Management]. <http://hdl.handle.net/1721.1/17442>
2. Granger, C. W. J. (1988). Some Recent Developments in a Concept of Causality. *Journal of Econometrics*, 39(1-2), 199-211. <https://ideas.repec.org/a/eee/econom/v39y1988i1-2p199-211.html>
3. Sklar, A. (1959). Fonctions de Répartition à n Dimensions et Leurs Marges. *Publications de l'Institut de Statistique de l'Université de Paris*, 8, 229-231. <https://hal.science/hal-04094463v1/document>
4. Schweizer, B., & Wolff, E. F. (1981). On Nonparametric Measures of Dependence for Random Variables. *The Annals of Statistics*, 9(4), 879-885. <https://doi.org/10.1214/aos/1176345528>
5. Agha, G. (1986). *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT Press. <https://dl.acm.org/doi/book/10.5555/7929>
6. Genesereth, M. R., & Ketchpel, S. P. (1994). Software Agents. *Communications of the ACM*, 37(7), 48-53. http://logic.stanford.edu/publications/genesereth/software_agents.pdf

About the Author



MP

Mario Pardo

Founder and CEO at Ryse, Inc.

Mario Pardo is an industry practitioner affiliated with Ryse Inc., where he works at the intersection of business strategy and artificial intelligence-enabled innovation. His professional and scholarly interests focus on applied machine learning, analytics-driven growth, and the governance practices required to deploy AI responsibly in organizational contexts. Drawing on experience translating technical capabilities into operational outcomes, he contributes to initiatives that prioritize measurable value, scalable delivery, and stakeholder alignment. Pardo's work emphasizes the integration of data infrastructure, process redesign, and human-centered change management to support sustainable AI adoption. He is particularly interested in how generative AI and automation can augment decision-making, improve productivity, and strengthen competitive advantage across industries. Through ongoing engagement with emerging technologies and enterprise implementation challenges, he brings a practice-informed perspective to research and dialogue on business and artificial intelligence.

ORCID: [0000-0001-5110-7674](https://orcid.org/0000-0001-5110-7674) · [LinkedIn](#)

Copyright Notice

All articles are published in the Journal of Business and Artificial Intelligence under the Creative Commons 'CC BY' ("Gold Open Access") license, where authors retain the copyright of their articles. The author grants JBAI a license to publish the article under a Creative Commons 'CC BY' license, which allows the work to be freely accessed, shared, and used under certain conditions.

About the Journal

The Journal of Business and Artificial Intelligence (ISSN: 2995-5971) is the leading publication at the nexus of artificial intelligence (AI) and business practices. Our primary goal is to serve as a premier forum for the dissemination of practical, case-study-based insights into how AI can be effectively applied to various business problems. The journal focuses on a wide array of topics, including product development, market research, discovery, sales & marketing, compliance, and manufacturing & supply chain. By providing in-depth analyses and showcasing innovative applications of AI, we seek to guide businesses in harnessing AI's potential to optimize their operations and strategies.