Semantic Negotiation Among Autonomous AI Agents: Enabling Real-Time Decision Markets for Big Data-Driven Financial Ecosystems

[1]Goutham Kumar Sheelam, [2]Raviteja Meda, [3]Avinash Pamisetty, [4]Sai Teja Nuka, [5]Harish Kumar Sriram

[1]*IT Data Engineer, Sr. Staff, Qualcomm Inc., gouthamkumarsheelam@gmail.com*
[2]*Lead Incentive Compensation Developer, rvtjmeda@gmail.com*
[3]*Integration Specialist, Millsboro, avinaashpamisetty@gmail.com*
[4]*Sustaining Mechanical Engineer, Argon medical, saitejaanuka@gmail.com*
[5]*Lead software engineer, Global Payments, hariish.sriram@gmail.com*

Abstract: Weaved in every financial transaction are two key components - words and numbers. These make possible the semantic negotiation of the meaning of value in money markets and the commensuration of things with an associated price. But these are mostly carried out by human actors, and at the micro-level. Markets operating at nano-second intervals, the unprecedented velocity and frequency of these transactions in the current contemporary financial ecosystems make the former impossible. We introduce semantic negotiation talking in natural language about numeric value, enabled by Autonomous AI Agents. Briefly, we describe an architecture for the instantiation of Stock as a Service, a legislative framework for market microstructure that makes possible the automated technical, fundamental, and qualitative trading of financial assets by Autonomous AI Agents, making sense of the associated words and sentences repeated in the fabric of these markets, semantically negotiating the meaning of numeric value, its vagaries and uncertainties, and negotiate for inducing agents to subscribe to this information asymmetry - for a Fee. We combine foundational ideas developed in Natural Language Processing and Link Mining. We describe the envisioned self-organizing and self-regulating intelligent socio-technological ecosystems for semantic negotiation among Autonomous AI Agents in financial ecosystems that are responsive to policy and agent governance heuristics. The idea of Stock as a Service allows only those products made possible by this architecture and its market microstructure for asset price determination. For promoting needed transparency and disclosure for these products, and for promulgating the regulations needed for this market microstructure to create and enable autonomous trading by guiding new types of instruments like Digital Cowries.

Keywords: Autonomous AI Agents, Semantic Negotiation, Natural Language Processing, Numeric Value Interpretation, Financial Transactions, Market Microstructure, Stock as a Service, High-Frequency Trading, Value Commensuration, Information Asymmetry, Technical Trading, Fundamental Trading, Qualitative Trading, Link Mining, Self-Organizing Systems, Self-Regulating Ecosystems, Intelligent Financial Ecosystems, Policy-Responsive Agents, Governance Heuristics, Market Transparency, Regulatory Frameworks, Digital Cowries, Autonomous Trading Instruments, AI-Driven Price Discovery.

1. Introduction

Content-based semantic negotiation may be a step towards the realization of fully autonomous agents. Negotiation support has been addressed in several different areas, comprising business negotiations, virtual trading environments, multi-party/multi-layer negotiations, negotiation stability, negotiation services, and multi-agent societies, among others. These systems fall short of achieving full autonomy because negotiation strategies are usually hard-coded by the makers of agents, and resolve only limited types of negotiations. We explore semantic negotiation in the context of fully autonomous trading agents in a financial ecosystem.

Fully autonomous trading agents can act on behalf of individuals or corporations. No limitations on the way these agents may negotiate should be imposed. Most literature addressing trading agent negotiation focuses on predefined rules for each transaction. We argue that free negotiation over price, payment modality, and money transfer, as well as agent commissions, signal content, and signal transmission modality, both decrease the overall cost for the trading agents and allow smarter transaction handling. Our work focuses on enabling complex negotiation tasks; the objects of transactions are usually unique, and involve an object transfer in exchange for a transfer of ownership. This is the naming of one party to the other and an appropriate and unique identification of exchanged goods. Both flows are critical to the balance of the parties and are strictly conditioned by both parties. To achieve these complex transaction flows, autonomous trading agents need to negotiate using knowledge. Therefore, the data representations used during negotiation, and the protocols implemented that allow the parties to perform the exchange are crucial.

Traditional role-based negotiation models may not be sufficient for autonomous agents in an open electronic ecosystem setting. Human negotiators employ various approaches—not only price maximization and cost minimization—to satisfy their negotiation objectives. We study these problems and define helper components of a more complex negotiation model that allows for automated knowledge-based negotiation, based on agent task context and the semantics contained in data representations related to negotiation. This work explores the complex issue of how to approach and represent knowledge-based electronic negotiation for trading agents in the context of financial transaction ecosystems. The decision of whether and when to negotiate is driven by the actor’s sensing of positive or negative impulse; clarity of negotiation criteria; transitory or permanent user status; homogeneity of negotiated elements; agent budgets and thresholds; expected duration of negotiation; sensitivity of transaction prices; and the availability of negotiation resources.

Fig 1: Building a Global Ecosystem of the Decentralized Internet of AI Agents

1.1. Objectives and Scope of the Study

The proliferation of the Autonomous AI Agents technology, which enables independent operation, collaboration, and/or competition of autonomous AI agents, has robustly advanced in recent years thanks to advances in many AI areas, such as Natural Language Processing, Business Process Modelling, Intelligent Decision Support, Reinforcement Learning, and Multi-Agent Systems. Examples of such developments are the digitized use of large language models that allow natural, context-sensitive, and conversational interaction with semantic web data and systems that have produced a well-catchy phrase, such as 'Agents have become the new apps'. Furthermore, a promising development is the recent announcement of the autonomous agent, which couples advanced models with the external services it deploys supporting its task as instructed by a user.

Yet the design of negotiation protocols and the consequent composition of the multi-agent team to interdict such protocol, for enabling the operation of autonomous agents over a range of interested application areas is not so robustly advanced. Negotiation is a crucial function by which many companies autonomously operate, either in collaboration, via supply chains, or even competing against others, to achieve their goals. Negotiation systems are an integral part of many areas in concurrent engineering processes, e-commerce, robotics, web services, and so forth. Thus, developing such protocols is an especially active area of research in multi-agent systems design. In this thesis and previous works, we have developed flexible study and design methodologies based on the dialog semantics formalization technique. A clarification kernel is defined at a higher level that then guides the various cycles of lower-level, more specialized clarification protocols, utilizing negotiation semantics and dialogue pragmatics.

$$SAU_{i,j} = \frac{\sum_{k=1}^{m} \mu_k \cdot sim_k(\alpha_i, \beta_j)}{m}$$

$SAU_{i,j}$ = Semantic agreement score between agents i and j

μ_k = Weight for semantic feature k (e.g., intent, context, ontology)

sim_k = Similarity function over semantic vectors α_i, β_j

m = Number of semantic dimensions

Equation 1: Semantic Agreement Utility (SAU):

2. Background and Motivation

The FinTech industry is progressively deploying AI and engaging in financial transactions without human intervention. These transactions involve an array of complex algorithms communicating to determine products, pricing, conditions, terms, and conditions of contracts. Due to the complexity of communication protocols and the size of datasets, such AI communication needs to be pre-programmed. However, we believe that the next step in that direction is to allow AI agents to negotiate on behalf of their human users. This involves understanding the semantics and subtleties of the intentions of the parties behind the contracts. The philosophy of autonomous negotiation is an alluring prospect, and to develop it we need to solve AI functionality that has not yet been addressed, which requires inter-agent communication that goes beyond simply pre-programming relatively simple codes to create intelligent agents that can interpret and communicate nuances of intention and meaning. These agents will adjust their architecture and language models to accommodate each specific negotiation context, and subsequently update them to allow for supervised learning of events to improve performance on subsequent negotiations. We refer to this process as semantic negotiation.

There are some characteristics of their functioning in financial ecosystems that we think would be necessary for these agents to become a reality. First, AI and NLP techniques should be able to predict outcomes of precedent negotiations and supervisory accounts of what were the events and the meanings behind the negotiations of rules and multiple appearances of financial contract types, by developing their repository of the act of solutions. Then these predictions should be deployed by the conversational agents to help them adjust their internal models of the preferences of the other party before the negotiation, participate in the negotiation to ask questions to clarify the preferences, and retrospectively analyze the negotiation events to learn from them, incorporating any illicit or unreasonable behavior in the modeling of the other party. Finally, the purposes of the negotiation should be described semantically through the use of different reasoning verbalization techniques and frameworks.

3. The Role of AI in Financial Markets

AI is now an essential component of operations in all fields of financial activities. It has a remarkable impact in various areas of finance such as asset management, retail banking, corporate finance, investment banking, regulation and supervision, credit scoring, market making, and algorithmic trading. In the business of asset management, AI algorithms perform competitions actively. They input massive quantities of market data in all areas of the world and perform statistical analyses upon it, developing quantitative trading strategies and providing statistical arbitrage opportunities, envisaging relevant decisions, being either long or short. Retail banks use various services of AI, from chatbots, recommending credit cards or loan amounts, to fraud detection and prevention.

AI methods are incorporated in market making and algorithmic trading techniques, relying on technological advances of high-frequency trading with the objective of profit maximization. Coping with competition within microseconds and milliseconds, such trading activity requires the use of AI-based suitable algorithms. They analyze data to exploit any inefficiencies created by microstructure noise within the market. Such algorithms are informed by the speed of technological and informational advances and uncertain structural stability-- a new data point becomes available when a major transaction occurs, during opening and closing, with larger volumes, and during periods of major shocks. Other applications of AI in finance include credit scoring, regulation, and prudential supervision. A challenge here is the opacity of results. The methods used are "black boxes" and the supervisory authorities follow a non-invasive approach regarding their dealings, verifiably monitoring the techniques without intervening.

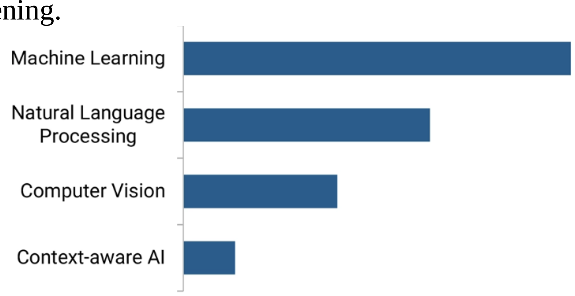

Fig 2: Artificial Intelligence Market Size

3.1. Overview of AI Applications in Finance

Artificial intelligence (AI) and machine learning (ML) have attracted growing interest over the last two decades, as they promise to significantly improve the accuracy and cost-efficiency of tasks that are of major importance for individuals and businesses alike. In finance, AI adoption has now reached a level of maturity at which several applications are commonplace. The financial AI market had an estimated size of USD 3 billion in 2021, with a projected compound annual growth rate (CAGR) of 21.6% for the period 2022-2027. In investing, just a few examples of typical applications are algorithmic trading, risk assessment, robo-advising for individual clients, sentiment analysis, financial forecasting, and analyzing financial documents, such as credit ratings and earnings reports.

However, AI-generated predictions, signals, and decisions still expose firms to important risks, as AI algorithms can be and have been manipulated and deceived. For those reasons, regulation and risk management are critical areas related to AI in financial markets. For example, the proposed European Union AI Act lays down rules for high-risk AI applications in connection with specific industry sectors, including also financial services. The Basel Committee on Banking Supervision has also initiated discussions to issue guidelines on AI about banking supervision. It is of particular importance to understand how much of the risk for predictions, signals, or decisions should lie with the financial institutions if these were to use an AI algorithm as a tool. Would the financial AI solution be provided as a service, a common AI decision-making framework, or a policy-based autonomous agent? Would the decision-maker provide the decision framework but not the actual decision-making, as for AI-driven implementation or reinforcement learning? Or would full AI autonomy be permitted in closed-loop implementation?

3.2. Challenges Faced by AI in Financial Decision Making

These days it is difficult to avoid discussions about the role of Artificial Intelligence (AI) in modeling and decision-making seen as a pool of probabilistic methods that seem to provide answers to anything. It is claimed that if humans are not capable of managing a field of complexity, we should rely on AI. In finance, this is vigorously countered. During the meltdown, AI, neglecting its constraints, could not answer crucial decision-making challenges. Short-term decisions by AI relied on extrapolation of facts stated by market actors, and long-term decisions neglected the existential purposes of agents, or became misaligned with the social community. After the crisis, they required firms to prove that their AI was aligned with well-functioning markets. Following the market crash, volatility was largely attributed to an increase in HFT and algorithmic trading, feeding on information provided by public chat rooms.

Discussions can go two ways: whether AI should be applied in finance, and if yes, which shape should it take? It is nearly undisputed that data and regulatory-driven AI, provide surveillance of algorithmic trading behavior of large pools of liquidity. Yet, for decision-making by financial intermediaries or autonomous AI in AI-based auto-negotiation, we are at the beginning of a long, winding road; there are several challenges concerning AI-enabled financial decision-support systems, and the jury is still out on the decisions to come. In terms of market economy, a market is a computational tool that allows agents to find the optimal modality of exchange. If autonomous agents are cooperating through AI, there is a need for negotiation AI that allows the exchange of not only numerical results but also the underlying decision models. For decision-making companions that only provide recommendations, trust is one of the major challenges. Every decision agent in the assist/advise loop has a unique decision signature that involves their model and data for the explanations.

4. Semantic Negotiation: Concepts and Framework

When agents have to agree upon a solution to an open negotiation such as mediation, the term negotiation comes into explanation, as it is needed as an agreement and what is explained above is not sufficient. Negotiation is required when autonomous agents are not supposed to cooperate, or there is limited cooperation, with the other agents or humans, but still, it is required an agreement is reached through communication, bargaining, and arguments. People negotiate in their everyday lives all the time, for both granting and receiving services or commodities. Blinded by our anthropomorphic lenses, we consider computers and artificial agents in general as helpers; yet, at a more developed stage, they could be unbiased interlocutors. Negotiation has become an important mechanism for coordination between AI agents in an increasing variety of applications: asynchronous energy management, software package management, resource allocation in multiuser environments, distributed resource scheduling and management, auctioning, collaborative multimedia delivery, negotiation of security policy and security association protocol, purchase bidding in financial applications, just to list a few.
Formal models of the notion of ontologies, for agents to negotiate in an open context such as a mediator, being parties in the negotiation process, can be made accordingly. Our framework models such informal negotiations among agents and mediators in any open financial application, since it has its roots in knowledge representation and dialogue as a simulation of natural language dialogue, together with logical inference. An important aspect of our model is that, because of the heterogeneity of the parties, negotiations are tied to the mediator, which is open: mandatory if we want to preserve the open service or commodity semantics, yet logical coherence: by logical coherence, we mean that it must exist a logical space in which all parties agree on the ontologies, and properties governing logical primitives.

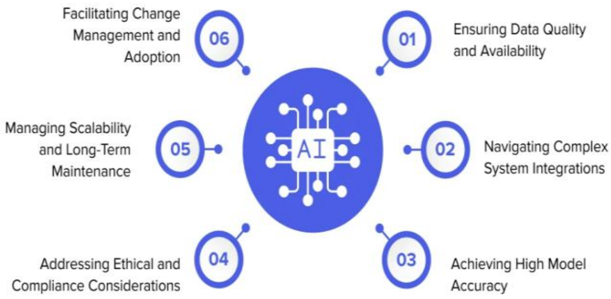

Fig 3: Semantic AI for Business is Redefining Search Strategy

4.1. Definition of Semantic Negotiation

Negotiation has traditionally been understood as a joint process in which two or more stakeholders exchange offers and counteroffers until they reach an agreement. In recent years, this view has evolved to accommodate more complex commercial interactions. In the classic sense, the final result is usually the result of several successive stages that lead the agents to negotiate more and more precise offers, for example by formally defining their requirements or their commercial interests. Nevertheless, negotiation can also be viewed as a "staged decision-making process in which autonomous agents, possessing conflict and/or mutually beneficial preferences over a set of possible outcomes, engage in an iterative mechanism of communication and deliberation."
In the business-to-business environment, the negotiation usually takes place between two or more partners, who have precisely defined the terms and conditions and the product or service that will be exchanged. These are usually present in the form of Electronic Marketplaces and V-commerce, where Internet Service Providers offer their services, goods, or products. However, this involves the presence of people and a large number of negotiations. In this way, the Intelligent Agent Systems are introduced as automated software agents that allow for the reduction of the time and costs in these processes.
In this scenario, the negotiation is seen as an adversarial game, where agents usually negotiate the prices only, without taking into account other attributes. These mechanisms have shown to be effective, although they did not advance the solution for service selection or other complexity domains that cannot be tackled through a pure negotiation game.

4.2. Framework for Autonomous AI Agents

We present a semantic negotiation framework aimed at helping autonomous AI applications find suitable solutions for their clients as well as autonomously create value and increase their own and, possibly, others' utility within financial ecosystems. The framework enables agents to efficiently employ secure AI-driven solutions for sensible tasks that can include imparting substantial services, and simultaneously reducing and efficiently managing various kinds of risks, including security, privacy, monetary, and many others. The framework contains the following elements.

Description of agent motivation and utility. The motivation can be created by pushing agents to increase their utility to the maximum within the concise but complete description of how the utility is evaluated. The description should include what costs and how should be compensated by the utility and what benefits should be expected from the services offered for compensation. Components of motivation can include monetary costs, privacy and security costs, etc. Provided that agents can communicate voluntarily, can accept or decline job or service requests, and can notify their clients of varying success on the way to building a durable relationship, agents will employ the utility to satisfy clients in a way that will lead to satisfactory rewards or compensations of expenses at a minimum for both sides. For such agents, the utility will include the satisfaction of providing quality services in addition to compensating payments, privacy and security risks, and task-related resource consumption.

$$AMBF_i(t) = \theta_i \cdot \frac{U_i(t)}{C_i(t) + \epsilon}$$

$AMBF_i(t)$ = Real-time bid strength of agent i

θ_i = Agent-specific negotiation aggressiveness coefficient

$U_i(t)$ = Predicted utility of market outcome at time t

$C_i(t)$ = Estimated cost or risk of action at time t

ϵ = Smoothing constant

Equation 2: Agent Market Bidding Function (AMBF):

5. Big Data in Financial Ecosystems

1. Sources of Big Data in Finance

Data is the core of AI-based financial systems. In a financial ecosystem populated by a large number of autonomous AI agents engaged in semantic negotiation cycles, greater volumes of data are created, exchanged, interpreted, solicited, and validated. Enabling value utilization at the ecosystem level (with AI agents autonomously generating new knowledge) is reliant on the ecosystem structure via stronger interaction links, and mutual trust, followed by an effective layer of solutions facilitating the interactions. There is a distinction between data that is 'pushed' into the market when it is created or generated by a decision maker, versus data that is 'pulled' from the market by the decision maker at a point in time.

We have entered the era of Big Data, which has meant an exponential increase in the amount of data created and made available. The very nature of vast amounts of data has become multidimensional, both from the sources and valuation methods' points of view. The changes in company governance and regulatory transparency requirements, for instance, lead to traditional fundamental data publishing patterns with cyclic and seasonal characteristics becoming less impactful. Some traders and trading book managers even argue that they can trade more easily on quantitative data than purely fundamental financial data based on analyst forecasts. New alternatives to traditional data providers have appeared in the market, crowd-sourcing alternatives promising to provide new data value-generating alternatives for traders and financial analysts.

2. Impact of Big Data on Financial Decision Making

The availability of numerous alternatives has removed the constraints of traders' reliance on traditional data services. Furthermore, regulatory transparency measures have also introduced more qualitative types of data, referred to as 'alternative' data. These new sources include web scraping of customer reviews and news topics being used to identify publicly perceived company themes and their impact on stock performance; the scraping of financial portfolios and web transactions and financial statements posted on company websites and social media pages; mention of fluctuations associated with communication; credit card data; satellite imagery used to count customer traffic visiting retail stores;

geolocation data used to estimate consumer traffic into retail businesses; and text mining aimed at detecting fund flows from mutual fund company releases.

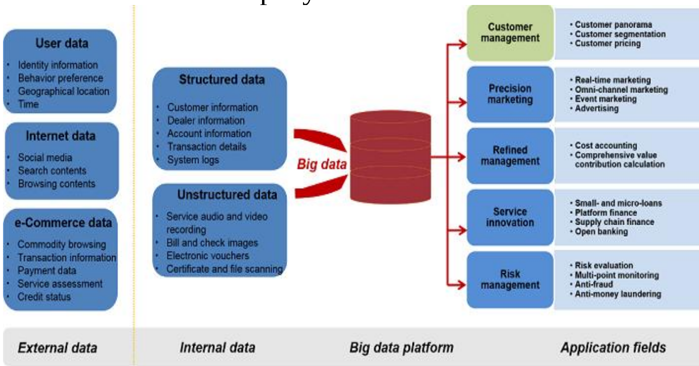

Fig 4: Big Data & Financial Service Innovation

5.1. Sources of Big Data in Finance

With the advent of smartphones and the increase in internet availability, many sectors of the economy are now wired, allowing large amounts of real-time transactional data generated every second. Internet of Things devices connected to the internet are sharing significant insights on consumer behavior and preferences of choice. Such data can be used in the finance sector to optimize business operations and for analytics-driven hiring, trading, and other strategic decisions. Thanks to modern communication, information that is available through daily news and other channels can be connected and stored in one location for future reference. Social media allow conversations related to any topic trend in real-time, often leading to a collective voice that can impact or reflect public sentiments. Financial organizations are concerned about conducting enhanced analysis and prediction of market data and movement trends. Such an assessment is far more effective when combined with social news data.

The ever-increasing growth in data has brought with it many challenges for storage, processing, and preliminary analytics. Traditional database systems are ill-equipped to store and process such immense and often unstructured data volumes. The insight obtained from this data can be used to impact policy and for decision-making. Financial markets are often volatile, reacting to every new piece of information be it qualitative or quantitative. This was accentuated by the pandemic. The availability of a large, diversified, and complete information set has led to the concept of Big Data in Finance. The enormous amount of data today's financial services rely on comes from sources inside and outside the organization. Transactional internal data from ERP systems will deal with core processes.

5.2. Impact of Big Data on Financial Decision Making

With the rapid development of big data in the financial field, the reliance on traditional financial statements and projections has been heavily challenged. More importantly, with the development of decision behavior research in recent years, the traditional assumption of the irrational behavior of market participants has also been challenged. Therefore, the imbalance theory still applies, even with different theories, and thus it is necessary to study big data and the decision-making process of market participants. Here we summarize several important aspects of market behavior, to review not adventurers. Traditional investors perform differently in both the information acquisition and decision phases, which enables arbitrage fund managers to withstand higher risks and implement trading strategies at various intervals. Diverse market participants approach big data from different perspectives and make use of diverse models to exploit unique idea generation. It is only after we enrich the ideal set with more diverse market participants who employ better models of idea generation that we can reduce arbitrage opportunities to zero. Moreover, the development of AI also makes sense, and non-linear models have been proposed for idea generation. Hence, the competitive advantage of AI with operating speed over humans becomes smaller with the development of AI. For a perennial equilibrium, the trading frequency and speed of arbitrary funds must be similar, while ordinary investors are detached from tasks of arbitrage trading. When big data enters a bubble stage, however, the balance will certainly be disrupted. For high-quality news about a company's operational conditions, at different times the reaction of capital should be quite different, but during big data anomalies, it elicits price reactions several times greater than the incoming signal.

6. Real-Time Decision Markets

Although an element of decision-making is found in almost every economic transaction, decision-making itself is seldom the object of a transaction. Decision markets are a novel concept. They address a gap in current markets; provide market participants with new products, services, and income opportunities; increase the efficiency with which decisions are made, remove the need for recourse to traditional intermediaries; provide an enabling infrastructure for decentralized collective action; and enrich decision dialogue. Decision markets have two main functions; they provide a mechanism for real-time decision-making, ensuring that key decisions are made in time; and they offer a platform for the development of market products and services such as restricted solicitation, virtual steering, reporting, alerts, and bounties. To fully understand the market we created, we first define the concept of decision markets from which the market is derived.

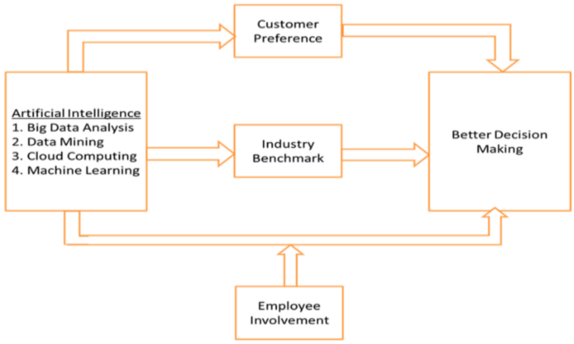

Fig 5: Artificial Intelligence and Better Entrepreneurial Decision-Making

6.1 Concept of Decision Markets

The concept of decision markets extends applied market theory beyond the exchange of material products and services to the element of decision-making. Decision markets are mechanisms for the actual operation of the decision-making core function of the economy. They are not just additional structures, facilitating execution, reporting, and clearing functions of different types of other transactions. Although they contain these functions, decision markets are in essence collections of incentive structures that determine outcomes of who decides, when a decision is made, what alternatives are considered, and how the decision is made. The basic premise is simple: every decision influences the choices of a large number of other market participants. Introducing an incentive structure that influences the economic decision-making of an individual will thus change the decisions of all other individuals.

6.2 Concept of Decision Markets

While automated, machine-driven trade in various financial instruments is a common thing today, investor-firm-driven, active financial decision-making is decentralized, and performed by individual firms separately. Several decision markets exist presently - to name a few popular - the stock, the bond, the commodity market, etc. The agents involved in decision-making perform their investments with their financial means in expectation of private gains. The investment gains are made based on the decision of buying/selling in an asset market. It is believed that during the short period of an asset trade, the action of each agent is motivated primarily by self-interest rather than altruistic consideration. Moreover, it is not expected that an agent's choice will be significantly affected by share price movements associated with trades by other agents. The naive, simplistic premise, of many economists for such an asset futures market to work is that all agents involved are trying to make money-fund transfers from one group of agents into the hands of another group.

Despite the substantial capabilities that computers have demonstrated in interpreting different types of signals, the decision investors' expertise in evaluating the desire of society for particular products and services across periods and forecasting the ability of a firm to provide these when needed and the resilience of the firm when unexpected shocks occur is irreplaceable. Automated systems can only act as secondary tiers of the overall decision system. In what follows, we describe, in informal terms, how the notion of decision markets can be generalized to both the problem of time and the more general investment decision. We consider these generalizations to be logically necessary within the context of the overall theory outlined.

6.3 Mechanisms for Real-Time Decision Making

The concept of decision market describes the setting of open, large, real-time information markets where market participants can freely put up and accept bets concerning future events, as well as trade event contracts on events that other participants have offered to bet on. This is a general model that allows to construction of markets for all types of events and diverse financial instruments. However, the decision market does not prescribe what market mechanisms should be used to ensure continuous trading of event contracts so that the market participants can express their beliefs, and decisions, at any point in time.

We may use the following market mechanism. All bets that are accepted will be settled at the time of expiration. At expiration time, the market processes the event, finds its realization, and pays all winners. Of course, we need to make sure that, at the expiration time, the market has sufficient reserves for paying all event contracts that have been settled by non-zero amounts. If each market event has a time to expiration long enough so that the market can accumulate sufficient reserves, and if the market maintains reserves that are proportional to the total value of the event contracts in the market, this market mechanism assures continuous decision-making. The market needs to estimate the total value of the event contracts in the market to maintain sufficient reserves. This market mechanism allows all participants wishing to enter a bet contract at any point in time to do so as long as such bet contracts are fair. Further, the market settles the bets, which will not be zero payoffs at expiration such that all required payments to the winners can be made by the market, other than one riskless payoff.

7. Interoperability of Autonomous AI Agents

Research in IT systems, particularly the fields of distributed artificial intelligence and multi-agent systems, has long been focusing on very similar problems regarding the design of interoperable autonomous agents. The goal is to achieve seamless interoperability and collaboration on a decentralized peer-to-peer basis without central control, the need for prior specific agreements, negotiation costs, administrative overhead, or difficulties in adaptation by the participation partners. Autonomy does not negate the requirement for interoperability, because, like humans, these technology agents need to find consensus with other agents for purposes of communication and interaction. The consensus-building processes take place via the specification of machine-readable interoperability knowledge standards formed into communication and interaction protocols.

An agent system receives agent communication standards in machine-readable agent ontologies that focus on the specification of dealing with autonomous agents. They facilitate messaging and modeling activities, e.g., the definition of trade terms within a domain or business process specification for trade execution. For financial trading domains, semantic interoperability among negotiation-participating agents requires domain ontologies to cover the fundamental purpose of the system, which is to execute a trade to satisfy the utility-defined objectives of its market participants. Trade transactions are purely utility-maximizing interactions. Information about each participating agent's functional activities and private parameters is defined by algorithms as agent configurational ontologies in the agent descriptor. Autonomy of intelligent agents declares who defines what in terms of specification of the agent properties in the data structures. Unlike classical trading agent typologies such as brokers, dealers, or market-makers, no general rules exist for providing models and configurations of typical agents with predefined roles involved in a market transaction.

Different information needs definitions and parameter data setups create difficulties in the agent negotiation concerning the predefined trading definitions in the trade terms ontology. Such domain-related agent role distinctions act as preset triggers for the description and discovery of trading agent configurations via the configuration agent descriptor before the trading negotiation itself. Other machine-readable capability locators at the activity-sharing business process level provide protocols and checklists for function and request message exchange during trading interaction.

7.1. Standards for AI Communication

At the communication level, interoperability is achieved through the use of standards, which enable agents implemented at different platforms to have a common reference for the meaning and the structure of the data and knowledge exchanged. For intelligent agents to be able to interoperate, at least the following standards are needed:

• agent communication languages (ACL). These languages provide a set of performative primitives that convey the intention of the information contained in the message.

• ontologies. These provide the conceptualization of domains allowing agents to understand the content of the messages.
• knowledge representation formalisms. These define the structure in which the content of messages is represented.
• (temporal, spatial, and legal) annotations. These define the circumstances upon which the content of messages is valid.

Originally, a consortium put forth the Knowledge Query and Manipulation Language (KQML), and further developed it to be a complete standard for agent communication. KQML is one of a number of agent communication languages that have been proposed for use by agents. These include Agent Communication Language; Agent Language for Multi-Agent Systems; Diegos; M-APL; KQML; CHECKERS; FIPA-ACL; and Agent Communication Language.

In parallel with these developments, the formalization of ontologies has also seen active interest, and notably an amazing variety of often overlapping approaches have been undertaken. Some of the more prominent proposals include Description Logic, Ontolingua, Ontology Environment, and OntoWeb. Other proposals include Knowledge Interchange Format, Lexical Knowledge Base, Medical Logic, and Resource Description Framework.

7.2. Protocols for Semantic Negotiation

Regulatory interfaces can facilitate the iterative give-and-take communications required for semantic negotiations intended to reach agreements to cooperate on a contract. The terms of legally enforceable contracts are formalized as business rules in the executable business code of a business process. The business rules express the mutual obligations of each of the parties over time, including what is to be done, when it is to be done, in what manner it is to be done, and the payments to be made. Business rules must be written down in an unambiguous way that sets clear expectations and is generally accepted, and such rules are often specified by regulations. Whereas computational ontologies provide the semantic structure to the contract terms, contract negotiation protocols provide the semantics and pragmatics of the language used in contract negotiations.

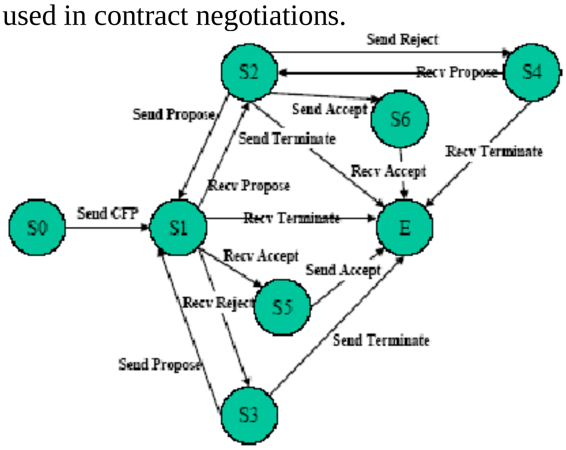

Fig 6: Negotiation Protocol

It is important to point out that the optimization strategy used by the agents jointly influences the strategic proposals they make as well as the responses they return. Therefore, the design of the negotiation strategy is one of the most effective ways for the designer to influence the agent's behavior and put restrictions on the problem space. Conditional offers – offers that are contingent on what action the other party takes – help both parties to converge to optimal deals and aid convergence at any time. A flexible and intelligent negotiation protocol enables intelligent agent systems to reach an agreement in cases when human agents reach no agreement. The challenge in the design of the negotiation protocols is to balance the complexity involved in the communication between the parties and the results achieved. The initial knowledge is often biased, requiring a more dynamic and elaborate negotiation to reach a satisfying conclusion.

Explicit model sharing allows more elaborate negotiations with greater amounts of shared information about what is known and what knowledge is currently being shared. Task sharing refers to what tasks are being performed, shared, and variously attributed. Different agents generally share different models of the world and the other parties in the negotiation, as well as different commitments, beliefs, and abilities. However, agents are capable of learning, updating, and broadening those models, and agent interactions can be used as a vehicle for model exchange and updates.

$$DCS = \frac{1}{1 + \frac{1}{n(n-1)} \sum_{i \neq j} |SAU_{i,j} - S\bar{A}U|}$$

where:

- DCS = Group-level agreement metric among agents
- $SAU_{i,j}$ = Pairwise semantic agreement
- $S\bar{A}U$ = Mean agreement across all agent pairs
- n = Total number of agents

Equation 3: Decision Convergence Score (DCS):

8. Conclusion

In this essay, we presented a vision of a novel intelligent financial ecosystem, where autonomous AI agents engage in concurrent negotiations, with the goal of efficiently and flexibly deciding on a financial transaction. We hypothesized reasons for favoring indirect negotiation, through an intermediate Phygital Exchange - e-transactional - Market, and built on the established theory of Semantic Towers to define what semantics are required for the underpinning intelligent services to be able to support Semantic Agent Negotiation.

Although research in distributed AI, multi-agent systems, agent-based modeling and deployment, and the financial domain has been underway for decades, this proposal raises several unanswered questions to be explored further. We have provided answers that guided the design of the MATS prototype, however, that research was primarily top-down, resulting in a rather complex framework. Thanks to these new explorations and insights, evidence seems to suggest a novel bottom-up approach may lead to the necessary insights to allow further simplification of the proposed architecture, while still supporting the stated functional and non-functional requirements.

Future directions for more research include: the enabling financial ontologies, the aggregator acts, agent value-agent demands to equivalence, the introduction of signals on 'uncertain' ontological relations, the mechanism to equilibrate two agents' negotiations exchanges over time (to reverse micro-fraud situations), and the level of acceptable 'uncertainty' for several user interactions (notably invitations to meetings for confirming signals or equilibrating exchanges). More generally, we have thus far not considered how to support anomalous situations detection and the definition of expert feedback into the infrastructure to enable reinforcement learning and agent specialization. Along with these questions, several Phygital Market and Exchange - e-transactional - Market 'real' operation use-cases, starting with the classical finance - investments, loans, financial guarantees, and waged transaction market - need to be further explored, to reinforce the completion towards possible solutions for the hypothesized innovative, reliable, simpler and sustainable financial ecosystem.

8.1. Final Reflections and Future Directions

Practical implementation of the foundations of a systematic theory of semantic negotiation conducted among autonomous AI agents, namely the conception of that negotiation as an interdependent game among those agents and the theoretical resources introduced in the chapters of this essay is an important future objective of our research group. It aims to reach a proof of concept of how the resources for semantic negotiation shape the behavior of agents in a real-world context. For that, we plan to use a simplified version of a financial ecosystem, i.e. a multi-agent system in which a small set of heterogeneous AI autonomous agents exchanges simple financial products with different initial resources, behavioral attributes, and interaction models. More precisely, the agents will conduct simulated operations in a synthetic but supervised environment in which the rules of the agent's interaction are predefined but their investment strategies are autonomously learned in a reinforcement learning fashion over time. Due to practical limitations, we will focus our implementation on the dimensionality of the agent's attribute vectors, the dimensionality of the data space, and the dimensionality of the product space. Initial findings suggest that as we increase these dimensionalities, the plurality of strategies learned by agents tends to grow, increasing the desired emergent properties of the financial system simulated in our experimentation.

References

1. Polineni, T. N. S., & Seenu, A. (2025). The New Frontier of Healthcare and Industry: Subash's Expertise in Big Data and Cloud Computing for Enhanced Operational Efficiency. Cuestiones de Fisioterapia, 54(2), 271-283.
2. Maguluri, K. K., Ganti, V. K. A. T., Yasmeen, Z., & Pandugula, C. (2025, January). Progressive GAN Framework for Realistic Chest X-Ray Synthesis and Data Augmentation. In 2025 6th International Conference on Mobile Computing and Sustainable Informatics (ICMCSI) (pp. 755-760). IEEE.
3. Koppolu, H. K. R. Deep Learning and Agentic AI for Automated Payment Fraud Detection: Enhancing Merchant Services Through Predictive Intelligence.
4. Nampalli, R. C. R., & Adusupalli, B. (2024). AI-Driven Neural Networks for Real-Time Passenger Flow Optimization in High-Speed Rail Networks. Nanotechnology Perceptions, 334-348.
5. Chakilam, C. (2022). Generative AI-Driven Frameworks for Streamlining Patient Education and Treatment Logistics in Complex Healthcare Ecosystems. Kurdish Studies. Green Publication. https://doi. org/10.53555/ks. v10i2, 3719.
6. Sriram, H. K. (2023). Harnessing AI Neural Networks and Generative AI for Advanced Customer Engagement: Insights into Loyalty Programs, Marketing Automation, and Real-Time Analytics. Educational Administration: Theory and Practice, 29(4), 4361-4374.
7. Burugulla, J. K. R. (2025). Enhancing Credit and Charge Card Risk Assessment Through Generative AI and Big Data Analytics: A Novel Approach to Fraud Detection and Consumer Spending Patterns. Cuestiones de Fisioterapia, 54(4), 964-972.
8. Chava K. Dynamic Neural Architectures and AI-Augmented Platforms for Personalized Direct-to-Practitioner Healthcare Engagements. J Neonatal Surg [Internet]. 2025Feb.24 [cited 2025Mar.24];14(4S):501-10. Available from: https://www.jneonatalsurg.com/index.php/jns/article/view/1824
9. Challa, K. (2024). Neural Networks in Inclusive Financial Systems: Generative AI for Bridging the Gap Between Technology and Socioeconomic Equity. MSW Management Journal, 34(2), 749-763.
10. Sondinti, K., & Reddy, L. (2025). The Future of Customer Engagement in Retail Banking: Exploring the Potential of Augmented Reality and Immersive Technologies. Available at SSRN 5136025.
11. Malempati, M., & Rani, P. S. Autonomous AI Ecosystems for Seamless Digital Transactions: Exploring Neural Network-Enhanced Predictive Payment Models.
12. Pallav Kumar Kaulwar. (2023). Tax Optimization and Compliance in Global Business Operations: Analyzing the Challenges and Opportunities of International Taxation Policies and Transfer Pricing. International Journal of Finance (IJFIN) - ABDC Journal Quality List, 36(6), 150-181.
13. Vankayalapati, R. K. (2025). Architectural foundations of hybrid cloud. The Synergy Between Public and Private Clouds in Hybrid Infrastructure Models: Real-World Case Studies and Best Practices, 17.
14. Nuka, S. T. (2025). Leveraging AI and Generative AI for Medical Device Innovation: Enhancing Custom Product Development and Patient Specific Solutions. Journal of Neonatal Surgery, 14(4s).
15. Rao Suura S. Agentic AI Systems in Organ Health Management: Early Detection of Rejection in Transplant Patients. J Neonatal Surg [Internet]. 2025Feb.24 [cited 2025Mar.24];14(4S):490-50.
16. Kannan, S. (2025). Transforming Community Engagement with Generative AI: Harnessing Machine Learning and Neural Networks for Hunger Alleviation and Global Food Security. Cuestiones de Fisioterapia, 54(4), 953-963.
17. Srinivas Kalisetty, D. A. S. Leveraging Artificial Intelligence and Machine Learning for Predictive Bid Analysis in Supply Chain Management: A Data-Driven Approach to Optimize Procurement Strategies.
18. Challa, S. R. Diversification in Investment Portfolios: Evaluating the Performance of Mutual Funds, ETFs, and Fixed Income Securities in Volatile Markets.
19. Vamsee Pamisetty. (2023). Intelligent Financial Governance: The Role of AI and Machine Learning in Enhancing Fiscal Impact Analysis and Budget Forecasting for Government Entities. Journal for ReAttach Therapy and Developmental Diversities, 6(10s(2), 1785–1796. https://doi.org/10.53555/jrtdd.v6i10s(2).3480
20. Komaragiri, V. B. (2022). AI-Driven Maintenance Algorithms For Intelligent Network Systems: Leveraging Neural Networks To Predict And Optimize Performance In Dynamic Environments. Migration Letters, 19, 1949-1964.
21. Annapareddy, V. N., & Rani, P. S. AI and ML Applications in RealTime Energy Monitoring and Optimization for Residential Solar Power Systems.