PhD Candidate: Davide Magnanimi

Advisor: Prof. Stefano Ceri

"ONTOLOGICAL REASONING ON FINANCIAL KNOWLEDGE GRAPHS"

National Central Banks, such as the Bank of Italy, play a crucial role in maintaining economic and financial stability. They carry out various functions, including economic research, supervising financial intermediaries, and enforcing anti-money laundering regulations. The growing complexity of the financial landscape, marked by intricate corporate structures and cross- and indirect shareholdings, requires the adoption of advanced analytical tools to support institutional responsibilities and decision-making processes. While useful for specific tasks, traditional approaches are inadequate for automatically extracting hidden insights (e.g., corporate control chains), and promptly identifying critical risks (e.g., hostile takeovers). Moreover, the ever-evolving nature of financial markets requires continuous updates and the harmonisation of data from various sources to ensure accurate and timely assessments.

In this dissertation, we propose an innovative framework to address these challenges by introducing new methodologies based on Knowledge Representation and Reasoning. The framework integrates the representation of the financial landscape, in the form of a Knowledge Graph, with the logical formalisation of the domain knowledge as reasoning rules in the Vadalog language, a powerful language belonging to the Warded Datalog+- family that provides high expressivity with low computational complexity. Hence, the framework leverages ontological reasoning on Knowledge Graph to offer valuable and comprehensive insights on the complex financial landscape, to enable simulations and consequence analyses, to support the reactive monitoring of specific economic dynamics, and to enhance the prediction of future financial phenomena. It represents a solid tool for supporting data-driven decision-making processes and enhancing research, regulatory, and supervisory activities in the financial context.

A first contribution lies in constructing the Italian Companies Knowledge Graph, leveraging the data available at the Bank of Italy, which gathers granular information about Italian companies, their shareholdings, and relationships. The Knowledge Graph also integrates a logical formalisation of fundamental financial concepts, such as the integrated ownership, company control, ultimate controller, and close-links, thus enabling reasoning techniques to derive new implicit and valuable insights in advanced domain analyses. This includes the identification of control concentration in corporate groups or the resilience assessment of the national corporate landscape to economic shocks, as seen during the COVID-19 pandemic. To support regulatory authorities in protecting national strategic assets from foreign acquisitions, we demonstrate the framework can be employed to proactively identify defensive strategies and reactively monitor the implication of potentially hostile acquisitions, even suggesting the application of the regulatory tool Golden Power to block them.

As the financial context rapidly evolves, we extend the framework to make the Company Knowledge Graph update process more timely and efficient. We present an advanced methodology that incrementally updates control relationships by efficiently localising ownership variations, thus avoiding costly recomputation on the entire Knowledge Graph. This contribution marks a considerable advancement in the usage of automated reasoning in the context of financial supervision as its employment significantly enhances interactive "what if" analysis, in which analysts simulate shareholding variations of one or more companies to figure out potential regulatory implications in advance.

A further contribution of this thesis is the definition of an innovative approach for the prediction of corporate takeovers based on the combination of logical reasoning techniques on Knowledge Graphs and predictive statistical models. By identifying and formalising determinants of historical takeovers, the framework can systematically assess the presence of these factors across the entire Company Knowledge Graph and thus provide accurate estimates of the probability of each company being the subject of an acquisition. We

recognise the framework to be an advanced and particularly effective tool for the analysis of the financial domain. In order to make it more accessible to financial analysts in their daily activities, we present an interactive web application for Knowledge Graph analysis and exploration, which abstracts the complexity of the logical formalism of business definitions through the metaphor of reasoning widgets. Widgets allow the user to dynamically enrich the analysis, implicitly requesting the underlying reasoning framework to infer new relationships and concepts in real-time.

As a final contribution, we present an innovative method, based on the translation of logical rules into active triggers, to enable reasoning techniques within graph databases, such as Neo4J. This solution extends the applicability of reasoning techniques in scenarios that do not have or do not require the use of full-fledged reasoning engines.