Financial Industry Explanations

Mike Bennett

Hypercube Limited, London, UK

Abstract

Accountability is an important requirement of the financial industry. Reporting and explaining can be the means by which accountability is achieved but only if the parties have shared meaning for the terms being used. What makes this difficult is that financial terminology frequently uses simple everyday terms in highly technical ways that differ from one context to another. This article examines the issues for ensuring that financial explanations are correctly understood by all parties, both at micro- and macro-economic levels. The proposed technique for solving this problem is to use ontologies. Several examples of successes with the use of ontologies and business rules in an ontological framework are presented in some detail. Since ontologies are a relatively new technique for the financial industry, it is necessary for the ontology itself to be explainable. This problem is also discussed. The conclusion is that the financial system could benefit from formal approaches to explainability based on ontologies.

1. Introduction

THE FINANCIAL INDUSTRY has a lot of explaining to do.

There are a lot of ways of looking at a financial institution such as a bank. It is an entity that deals with money. It is a data firm with complex data interactions. It is an entity that buys and sells risk.

These things all have a role in explaining. Banks may need to explain to customers why they didn't approve this loan or that line of credit. They need to explain to regulators and shareholders why they did. Regulators look both for microprudential and macroprudential risk – that is, what risks banks take on for themselves and what risks they present to the broader economy.

Explanations form one kind of a more general matter which is central to finance: accountability. Institutions have to account for themselves to their shareholders and to regulators and central banks. Regulators and central banks give an account of things to lawmakers and the public, and so on. The more specific notion of explanation may come into play at any point in this information lifecycle – the key thing is that the relevant data are available, timely, accurate and understandable. Central banks will apply a number of statistical analyses to the data that come in from individual banks in the economy for which they have oversight.

In this article we begin in Section 2 by discussing the relationship among explanations, reports and accountability. This will provide some motivation for why financial explanations are important, both for the financial services industry in general and for the retail finance sector in particular. For explanations and reports to be meaningful, the terms that are used to express the explanations and reports must be understood by all parties that are involved. In other words the parties must have commonly accepted shared meanings for the terms. This is more difficult to achieve than one would expect, and in Section 3, we explain why simple solutions such as glossaries and dictionaries are inadequate and introduce ontologies as a better solution. We then give a variety of examples of successes with the use of financial ontologies in Section 4. While some of the examples are still at the proof of concept stage, the results show great potential for solving difficult problems in financial services. Sections 5 and 6 discuss some of the challenges with the use of financial ontologies. Section 5 discusses the notion of a business rule and the difficulties involved in formulating and enforcing them in an ontological framework. While introducing ontologies to finance has the potential for solving significant problem for explainability, reporting, regulatory enforcement and so on, it adds another problem; namely, explaining the ontology itself. This challenge is discussed in Section 6. We end with some final observations and conclusions in Section 7.

2. Reporting and Accountability

Reporting itself covers a wealth of requirements. Broadly speaking the reporting requirements in finance fall under one of two basic motivations: ensuring that consumers and other participants in the financial system are fairly treated; and making sure that the system itself does not fall into some unstable state.

For risk there is both internal and external reporting, risk management assessment, and compliance. On the consumer protection side there are regulations setting out a number of reporting requirements to ensure fairness towards investors, including price transparency and fairness, compliance with investment guidelines, compliance with stated fund management objectives, and so on.

Much of the focus of domestic regulation before the 2008 Global Financial Crisis (GFC) could be assumed to have been driven by lawmakers, who are driven by their electorates. The GFC provided a wake-up call in that what was needed was neither more regulation nor less regulation, but different kinds of regulation, embodying fresh thinking on the nature of global systemic risk. The sum of regulations that aim to protect the consumer would not sum up to better protection of the financial system itself.

2.1 Macroprudential Risk

The global financial system can be viewed as a complex dynamic system. This means that there are emergent behaviors arising out of actions and interactions within the system. It is not realistic to sum up all the risks seen by each participant in the system and expect to understand the risks to the system as a whole. This became very apparent in the 2008 GFC. For this reason financial system regulatory bodies look for a number of different kinds of information from industry participants, ranging from consumer-level protections (avoiding mis-selling and the like) to macro-economic and systemic risk factors. However, part of the nature of emergent phenomena in complex adaptive systems is that what emerges can't always be explained – at best, we can know why we don't know what happens next. Financial systemic risk management is more about anticipating what risks might be starting to emerge than about accounting for what might have happened after the fact. For this reason there are also initiatives in macroprudential regulation such as the Basel Committee for Banking Supervision BCBS239 regulation (Bank for International Settlements, 2013). This regulation defines a category of 'Systemically Important Financial Institution' (SIFI) and sets out how SIFIs are to be able to submit reports in the future under conditions that are not known in the present. As one central banker put it, we can't expect to deal with the next financial crisis using the reports that were appropriate for the previous one.

The reports and information submitted to regulators are not explanations in the form defined elsewhere in this issue, but provide the raw material for providing such explanations. A further lesson from the GFC was that data on their own are only part of the requirements for understanding what was going on. Many firms had all the data they needed to understand their exposures to failing or at-risk institutions, but it still took several weeks to arrive at a knowledge of those exposures.

Data does not mean knowledge. For that you need the addition of some kind of meaning; the semantics of the data enables them to be understood and re-used as a source of knowledge. For this reason the financial industry, like many others, is starting to figure out how to introduce formal ontologies into these data workflows.

Data accompanied by formal semantics do not in themselves form a set of explanations but they do provide the raw material for explainability. With this in mind we can look at a number of specific examples of explanations in the financial services space. Some of these are made available directly to the user, while others are directed to public regulatory authorities or to other financial ecosystem participants.

2.2 Explanations in Retail Finance

The common source of explanations requirements in retail is, of course, the customer. Whether in retail banking or credit cards, customers will want to know why they have been denied a new card or an increase in their credit. The standard set of answers include specific things like ‘your income is too low’ or ‘your existing credit balances are too high’ but may also include question-begging responses such as ‘You did not meet our criteria’ (what criteria?), ‘Too many credit enquiries’ (how many is too many?) or ‘You have not been at your current job long enough’ (how long is long enough?).

Many of the seeming explanations given to the retail customer will themselves beg follow-up questions that that customer feels they need to know. Set against this, the retail institution is often reluctant to hand over all the models and model inputs that they would use to make these decisions (thereby furnishing a full explanation) for the entirely understandable reason that someone in possession of all of these model parameters may use that information to game the system. Explainability leads to vulnerability.

At the same time, customers have the right to know that decisions made about them are made fairly and equitably. They have a right to know that the decisions made were based on current, up to date and accurate data about them. What if the job they are in is not the one on which the credit decision was based, or the debt balances ascribed to them have recently been paid off?

The challenge for the retail bank or credit institution is to ensure that the data they are working from are complete and coherent. Information they hold needs to be coordinated with that held by third parties such as credit reference agencies, and *vice versa*. Meanwhile their holding of data about each customer must comply with applicable regulations for disclosure on the one hand, such as 'Know Your Customer' (KYC), and non-disclosure on the other, such as the General Data Protection Regulations (GDPR) in the European Union.

Temporal mismatches need to be avoided, for example the lag time between data reported on at defined intervals and the data about things as they stand at the present time. Then there are variations in the usual pattern of credit card or *current* account holding, such as accounts with multiple holders, or those with holders who come under a status with specific protections, such as veterans. In some jurisdictions information held against an address, such as county court judgments, often causes subsequent occupants of that address to get down-graded – usually without a clear explanation that this is the case. Common complications often arise from situations of the borrower such as divorces, court orders, changes of address, or other personal circumstances. On the happier side of things there are unscheduled pre-payments by the borrower, where this is allowed. There may also be situations not under the control or knowledge of the customer such as concurrent fraud investigations, automatic payment system delays, differences or misunderstandings on month-end roll-over dates and so on.

This complex set of interlocking processes, disclosure and non-disclosure requirements and mismatches in knowledge between one party and another add up to a complex data management problem. It is also a complex problem of managing the relationships between the data and the things in the real world that the data are about.

The best-known of these is the issue of 'bi-temporality' in data management. This is the distinction between the date or time for which the data are about is available to the decision-maker, and the date or time that something occurred in reality for which the data are about. This matter of time is just one respect in which the state of things in reality and the state of the data that aim to reflect this may differ. Processes for data management and information supply chains need careful design and management, taking into account who needs to know what, how often and in what level of detail something needs to be reported, and how changes in circumstances are propagated across the entirety of the systems that hold relevant information, even understanding the impact on different data resources of a specific kind of change in the underlying reality.

Good customer service requires understanding the customer's journey through life. This sounds a bit like some set of buzz-words, but in reality it is important for the institution to be able to follow and understand the customer's changing situation, along with changes in statute law and in the institution's internal lending practices and longer term strategies, simply in order to avoid unnecessary misunderstandings, unhappy customers, reputational risk or legal exposures.

Explaining things in retail then is harder than it seems. Call center employees, the usual point of contact between the institution and the borrower, are effectively playing the role of knowledge workers, but all too often the knowledge they need is not available, or the knowledge is available but they are not skilled to the required level to make use of it in initial customer contacts. The tools they use may not be interoperable across different data sources, or the data they need may not be available in real time.

Meanwhile explanations are by their nature very contextual or scenario-dependent, and these dependencies also need to be understood. Decision support software meanwhile needs to balance the requirements for customer retention, profitability, and regulatory compliance.

These challenges will only get greater as innovations arise, both in technology, such as the increasing use of artificial intelligence in decision making, and in the financial marketplace itself, both in the emergence of new financial models in micro-finance and in the emergence of new technology-based systems such as distributed ledger (blockchain) technologies.

3. Shared Meaning

In the move towards more coherent use of data in financial reporting and decision making, one common theme has been the requirement for some way to represent common, shared business meanings. The common reaction to this requirement has been to try to establish business dictionaries or glossaries, where terms (words) are standardized to always mean the same thing, so this can be used as a point of reference across different data resources (Knight, 2018). A similar approach from the technology side is the use of 'data dictionaries'.

Both of these approaches suffer from a common weakness. In data dictionaries, each data model has textual information giving a 'definition' against each data element (field names and the like). It does not take long for these definitions to start to sprout extra qualifications, of the form "in the case of (XYZ) instrument, this field represents (some specific thing)." Soon extra business rules are added, reflecting logical statements about what sort of thing should be in a given field under different circumstances. The reason for this is that in any good data model design, data elements do not map precisely to single meanings of things in the real world. If they did, they would not be a design; there would be no data normalization or re-use.

Business dictionaries or glossaries have the same weakness but for a different reason. The way that humans use words is very contextual. I can use a word like 'bank' and you will know what I mean by whether I am talking about investment or fishing. The same words mean different things in different contexts. Even within finance itself there are subtle differences, for example the term 'over the counter' might refer to a derivative trade that is struck directly between parties, or it might refer to securities that are traded directly rather than through an exchange. A subtle difference but again any dictionary would need to add qualifying terms to state what concept is referred to by the words in different contexts. Meanwhile different parts of the industry and different functions within a firm may use different words to refer to the same concepts, for example 'coupon' or 'interest' on a debt instrument.

With both data dictionaries and business dictionaries this is not a fault but a feature – neither human words nor data field names map directly to concepts. The push for 'Why can't we all agree on the same terms' inevitably comes up against what Wittgenstein in his later work (Wittgenstein, 1953) calls 'language games'. Words play games.

People in the financial industry have long sought to solve the question of shared meaning for data elements, for example under the guidance of the Enterprise Data Management Council (EDM Council, n.d.). It was during one such meeting in which people were trying to agree on common terms for ‘Critical Data Elements’ in securities clearing and settlement, that someone thought to ask the question: “While we are disagreeing on what words to use, do we at least agree about what the concepts are?” The answer was yes – everyone agreed what the real world meanings, the concepts, were.

From this realization, the idea for a common ontology for the financial industry was born. The Financial Industry Business Ontology (FIBO) (Bennett, 2010) was initially conceived as a source of common shared meaning for the industry, to provide a point of reference for data models, integration, reporting, and other requirements. This was subsequently standardized as a series of machine-readable ontologies for use with financial industry data in a range of applications.

A formal ontology provides a simple account of the meanings of things. It does this by means of declarative statements of the form ‘there exists’ and qualifications such as ‘for all’. This falls under what is defined as first-order logic, and most published ontologies for use with data (including FIBO) use a sub-set or variation on FOL called Description Logic (DL). This is the sub-set of logic for which it is possible mathematically to prove that the assertions in the ontology are consistent and can be reasoned over in a finite period of time.

An ontology simply sets out a logical definition of what kinds of things there are, and what features distinguish one thing from another. FIBO defines a range of financial instruments in these terms, along with kinds of business entities and the relationships between these.

This way of saying things about something in the real world is necessarily limited but useful; in plain English terms, this first-order kind of ontology defines what there is, what kind of a thing something is, and what features or characteristics of a thing distinguish it from other things; that is, what are the necessary and sufficient characteristics for something to be considered to be the member of a particular set of things. This set-theoretic notion effectively defines a ‘concept’ (Odell, 2011). More specifically this is an ‘intentional’ definition of a concept. Some ontology languages also allow for extensional definitions where a set of things is defined by explicitly specifying all of its members.

4. Financial Ontology Examples

Some uses of ontology rely on the technical deployment of these as part of a solution to a specific problem, for example to draw inferences from available data. Other uses, in data management, integration, and reporting as well as artificial intelligence, rely on the provision of common meaning, via formally defined concepts, to streamline the information supply chain for reporting, for example, clarifying the meaning of each item in a report. This aids in the accountability of data in financial regulatory reports and thereby the explainability of the information contained.

This basic reflection of reality, as exemplified by FIBO, can be used to gain insights from data that would normally be sitting in different data silos under different schematic structures. For example, one set of data might contain information about a series of derivatives transactions, while another data source would carry information about the ownership and control relations between corporations or other business entities.

The underlying abstract model for many ontologies, the Resource Description Framework (RDF) (World Wide Web Consortium, 2014) coupled with the Web Ontology Language (OWL) (World Wide Web Consortium, 2012) provides a common syntax, so that terms from different data sources are framed using the same underlying technology language. In the case of RDF, this is the language of ‘triples’ (relations of the form subject-predicate-object). A collection of triples is a graph in which each triple is an edge in the graph.

The OWL language sits on top of RDF, and if the terms from different data sources in RDF are defined with reference to a single OWL ontology or a single mutually coherent set of ontologies, then data comparisons, formal inferences and semantic queries can be made against that data. FIBO provides one such set of mutually consistent ontologies, covering financial instruments, business entities and entity ownership and control relationships.

This kind of graph-based representation of data, coupled with a common schema in the form of an ontology, is known as a ‘knowledge graph’. A precise definition of the term ‘Knowledge Graph’ is a matter of ongoing discussion in the industry, see for example Ontology Summit (2020). For a workable definition see Yu (2020).

The potential for this basic knowledge graph framework is that if existing data across instrument transactions or holdings and business entity ownership and control hierarchies can be ingested from their existing data habitat and reframed as RDF data under a suitable ontology, we can ask new questions such as, “what is this bank’s exposure to that other bank, based on the trading positions it has open with not only that bank but its subsidiaries, parents and affiliates”.

4.1 Counterparty Exposures Proof of Concept

A proof of concept to demonstrate this usage was initially carried out at Wells Fargo using FIBO with indicative dummy data (Newman & Bennett, 2012). This was later repeated at another major US bank, State Street, with real data on interest rate swap transactions (David, 2016).

In this proof of concept a set of data about swaps transactions in a standard XML messaging format called FpML (International Swaps and Derivatives Association, n.d.) was fed into the triple data store and each data element was linked to the corresponding term in FIBO to define its meaning. A further set of data, available from U.S. Securities and Exchange Commission (SEC) filings and company registry information, was fed in to define the various ownership and control relations across the institutions that were the counterparty to each swap transaction in the swaps transactions data.

The business motivation for this work was what happened in the Global Financial Crisis: what would happen to the positions at this bank, if a certain other bank were to fall into bankruptcy? Given that each of these institutions has quite complex ownership and control hierarchies, this was not simply a matter of what trades this bank has with that other bank, but what trades it has with its parents and subsidiaries and what the knock-on effect would be on just one of those entities going down, on the complex network of relationships and positions it is tied to.

The resulting knowledge graph was queried using semantic queries, to return data about the monetary amount of each instrument position, the relative capitalization of each institution and the relationships between the relevant institutions. These results, in the form of data, were fed into a visualization program to provide a graph in which the relative trade positions were reflected by the thicknesses of lines between the entities, the capitalization was shown as the size of a circle representing each entity, and ownership and control relationships were additionally represented as lines between entities in different formats.

This proof of concept shows what can be done with ontology, not acting on its own but as something from which to feed graphical visualization techniques. A similar framework could also be used to feed mathematical models or other programmatic solutions. This is not ontology working alone but as a means to integrate data across a range of data sources and carry out operations across that data.

This also shows the use of a particular kind of ontology. In this case the ontology reflects the common meanings of instruments, transactions and business entities, but does so in a way that is directly applicable to data itself.

4.2 Bank of England Proof of Concept

At the Bank of England a pilot project was undertaken to show what could be done with ontology in the reporting chain (Bholat, 2016).

In the existing state of affairs the bank sends out a number of forms to those banks that fall under its jurisdiction. Each box in each form asks for a response, for example to give the amount of debt held by the reporting institution in US Dollars with 3 to 5 years residual maturity. The reporting bank looks to its various internal systems to find the answer to that question and puts it in the form.

Two issues were apparent in this approach. One was that of finding the right information in the right system, an inefficient and potentially time-consuming process. The other was that having received these forms from all these banks, the Bank was not fully confident that each reporting entity had assumed the same intended meaning for each entry in the form; they lacked confidence in the ability to compare like with like.

The premise of the proof of concept was that it should be possible to save time and cost for the reporting entities and at the same time increase the central bank's confidence in the reported information, if reports were made using granular, semantically-aligned data.

A further potential benefit was flexibility for the central bank. If data could be reported in a granular way with clear semantics for each element, then they should not need to send out forms at all. Instead each box on the report would be a semantic query against that granular data. This meant that if the bank wanted to introduce a new box on the form – for example if they wanted to isolate US Dollar debt holdings at different maturities (say, everything with up to 2 years residual maturity) then they did not need to redesign a form with this new box and send it around, they merely needed to write a new semantic query internally and apply this to the same data.

The first step in this proof of concept was to select three forms at random, analyze each line entry, and define the meaning of the data in that box. For example, you may have wondered what the term 'residual maturity' means in the above examples. This is the length of time, on any given day, until the debt in question has been paid off. That is not the same as 'original maturity', which is the length of time to maturity (debt repayment) at the time the security was issued. These may both be given the label 'maturity' in different data models, where the context is obvious by the function of that particular model. The original maturity of a 5 year bond will always have been 5 years, while the residual maturity (or current maturity, or some other label) is the amount of time from today until it matures. A five-year bond issued four and a half years ago has a residual maturity of less than one year, and this determines what box it should be reported in for this example.

The result of this phase of the proof of concept work was a formal ontology of the concepts that the bank had in mind when defining the information that they required in each box. Reporting against this ontology would enable the bank to recreate these forms locally using semantic queries. This would use a data querying language designed for this purpose, called SPARQL (pronounced 'sparkle').

4.3 Regulatory Proof of Concept

More complex reporting requirements call for more complex solutions, including the use of formal business rules.

Regulation W is a US Federal Reserve regulation that establishes terms for transactions between banks and their affiliates (U.S. Electronic Code of Federal Regulations (2002). It was enacted by Congress as part of the Federal Reserve Act and applies to all federally-insured depository institutions.

The Reg. W Proof of Concept initiative (Grosof *et al*, 2015) was formed by the Enterprise Data Management Council (EDM Council) and included Wells Fargo Bank, Coherent Knowledge Systems, SRI International, and the Governance Risk and Compliance Technology Centre (GRCTC) of Ireland (Governance Risk and Compliance Technology Centre, n.d.), with participation from other members of the EDM Council. This combination of participants was selected in order to have access to a range of rules-based technology solutions alongside the basic semantics expertise for the use of FIBO.

Regulation W defines a set of limitations against an illicit market practice called 'front-running'. Front running is the practice of buying or selling a security with advance knowledge of pending transactions that could influence the price, in such a way as to capitalize on that knowledge. To explain or account for whether a given trade does or does not count as a front-running trade, banks that could potentially carry out such trades are required to report all potentially applicable trades under a Regulation W reporting requirement.

In addition to the requirement to account for transactions as not falling foul of the Reg. W requirement, affected banks have an obvious need for internal decision support to determine, ahead of carrying out some potential trade, that it would not fall foul of the Reg. W requirements. This is an explanation requirement.

Core concepts used in this regulatory requirement include 'bank affiliate', 'covered transaction' (a potential transaction covered by the regulation), 'collateral requirements' and the notion of 'low quality assets'. Each of these terms needed to be defined and those definitions acted upon in decision making and explanations. These concepts are used to define limits to potential investments by the firm itself. The regulation stipulates that covered transactions with an affiliate cannot exceed 10 percent of a bank's capital stock and surplus, and transactions with all affiliates combined cannot exceed 20 percent of the bank's capital stock and surplus.

The term 'affiliate' here presented some definitional challenges, since it is both broader and narrower than the normally understood concept of 'affiliate' as being some entity that is either a parent or subsidiary of a given entity. It is broader because Reg. W 'Affiliate' includes firms to which the bank gives certain kinds of investment advice, and narrower since it refers not to the affiliates of all kinds of entity, but only to those that are affiliates, in this broader sense, to a bank. This means that the ontology needed to define the Reg. W concept of ‘affiliate’ as a sub-set of the union of affiliation and investment advisement in relation to the bank for whom the calculations are being carried out.

This is a good (if niche) example of why words alone cannot be used as the basis for meaning. It also illustrates how the meanings of words as defined in specific legislation texts are not necessarily suitable as a source of meaning for those words more generally. What a word means in the context of a given regulation may or may not also be what it means in some different context or some broader set of contexts. This is the reason that institutions need to navigate the realm of meaning by means of concepts and not by words.

The kind of trade that would fall foul of the Reg. W anti-front-running regulation thus required some fairly complex logic to describe it. The use of a formal ontology such as FIBO provides part of the solution to this reporting, in terms of common shared meanings, but there need to be more complex, or higher-order, logical operations on the data in order to determine and explain whether or not each trade comes under the Reg. W limitations.

The aim this proof of concept was to unambiguously understand and automatically comply with regulatory rules. The project used the FIBO in combination with advanced semantic rules defined in the rules languages Rulelog (Grosof, 2013) and Flora-2 (Yang *et al*. 2003), to automatically keep a bank in compliance as transactions were being processed. The intended result was to address the question ‘Am I in compliance?’ This had the associated explanatory requirement ‘Why / why not?’

FIBO and Rulelog/Flora-2 were used to make Reg.-W requirements explicit and applied to sample transaction data to automate compliance assessment. Detailed explanations were provided so that humans could understand the reasoning and facts that led to the conclusions. GRCTC provided expertise in controlled natural language for rule authoring via OMG's Semantics of Business Vocabulary and Business Rules language (SBVR) standard (Object Management Group, 2019) using the SBVR form of structured English. Coherent Technology and SRI technology provided automated reasoning capabilities using the Episto and Sunflower languages, with detailed explanations in English. SRI's technology provided automatic import of knowledge graph data in OWL, into the Flora-2 engine.

The rules engines defined a number of types of transaction that are defined as ‘covered transactions under the regulation, and a number of exemptions that were applicable. The proof of concept demonstrated that using these facilities a bank was able to enter into a transaction with a Counterparty, check if the counterparty is an affiliate, check if the transaction type is covered by regulation W and verify if the amount and total amount are permitted.

The structured English in SBVR was used to capture the business domain, specifically terms referring to business concepts, relationships between concepts and definitional constraints on these relationships. The ‘Rules’ part of the standard was used to capture the business behavioral constraints, obligations, prohibitions and so on. This formed a kind of bridge between the concept language of FIBO for the basic instrument concepts and the technology-based rules languages mentioned earlier.

The methodology developed by GRCTC delivered a system that could follow reference chains and produce self-contained sentences; define terms iteratively until all confusions were clarified; identify, describe and constrain links/relationships between terms, and capture regulatory requirements using the interlinked vocabulary elements from these other steps.

This proof of concept demonstrated the ability to deliver improved confidence in the correctness of compliance checks both for banks and for regulators. This was largely because understandable explanations were provided. This can reduce cost and risk due to the ability of this approach to adapt more easily and quickly to changes in regulations, since these are now framed using a common financial language, aligned with industry standards.

4.4 Explanations in Accounting: Tax Filing Example

Another particularly striking example of explainability makes use of knowledge graphs in an innovative way. This is the system developed by Intuit (Yu, 2020), the software vendor behind QuickBooks and other accounting applications for small businesses and consumers. In this patented innovation, users are able to file tax returns and interrogate each line entry to determine how that figure was derived.

This product uses a knowledge-graph (KG) based solution to determine the values to be placed in each line entry in a tax return. The basic

KG structure is enhanced by the addition of arithmetic functions such as ‘add’ or ‘subtract’, these functions being included in the graph structure.

Explanations for a given line entry are then provided to the user by means of traversing the graph to identify each of the inputs and functions used. These can be traversed iteratively so that the input to one function is traced to the output of an earlier function. So, for example, if a tax withholding entry is based on the following rule:

> **20.** If Line 19 is more than line 16, subtract line 16 from line 19.
> This is the amount you **overpaid.**

Then the user can interrogate the line entry for line 20 and see the amounts for lines 16 and 19. They can also traverse the graph by interrogating the line entry for line 19 and determine the line items that went into this and the fact that (in this example) these were added. And so on.

4.5 Ontology in Understanding Data

The range and complexity of requirements for financial institutions to be able to provide accountability and explanations leads to a set of complex data management requirements. One use of ontologies is in assisting the data management function within such firms to have a better understanding of their own data; better explanations internally of the data that they hold and the conclusions that are derived from this data.

These internal explanations make use of something called a ‘semantic data catalog’ (Newman, 2020). This enables the user, in this case, someone in the data management function within the bank, to pose questions such as “What types of customers are in the Customer table?” or “Where can we find organizational names in this database?”

The knowledge graph provides ‘chains of meaning’ relating to the real-world subject matter, such as ‘Customer has identity some Person’, ‘Person has name some Personal Name’ and ‘Personal Name has First Name some string’. These chains of connections make up the ontology and this can be applied to the data held in various databases and mapped to these meanings to provide answers to the questions posed by the data owners.

Similarly the data administrator can ask questions about what information is held in a given data resource, for example “What information is held in the Marketing Database Customer table?” or “Where would I find personal contact information in the Marketing database?” The Semantic Data Catalog is organized in such a way that the user can ask a number of broad based questions about the data.

The ability to address such questions of the data relies on semantic search capabilities, that is, the ability to frame questions in a semantic query form. While this is not explanation, being able to return data based on the semantics of a question is a prerequisite to accessing the right data for accountability or explainability further down the line. Turning user requirements for explanations into formal semantic searches will itself make use of a number of techniques. These include predicting concepts from string values or predicting a vector of concept plus predicate. Predicting string values may use lexical predictions, where mistyped text is replaced with text corresponding to the entries in the knowledge base, using metrics such as the Levenshtein Distance (Levenshtein, 1966); or it may use a concept vector approach, for example replacing the search text ‘vanilla interest rate swap’ with the synonymous term in the ontology, ‘fixed float interest rate swap’. Semantic search using concept and predicate combinations would use a standard semantic querying language directly, for example to return the concept of ‘agreement’ for a search on ‘contract is a type of?’.

This ability to link an ontology to internal data structures also provides the user with metrics on data quality, for example ensuring that stored data conforms to specific patterns for identifiers and the like, or that information on something like a person or a corporation contains all the relevant information as expected for that category of thing, as defined in the ontology. This makes use of another Semantic Web standard called SHACL (World Wide Web Consortium, 2017), which allows the user to define allowable patterns within the data in a knowledge graph.

Given data that may or may not conform to the pattern set out in this pattern (shapes) language, a validation report is produced which flag whether or not the information held in that data source conforms with the requirements for such data – for example that data about a human shall have only one date of birth, a given set of identifiers and so on. Where the report indicates that the data is not conformant to the required pattern an explanation is also generated, showing where the data diverges from the stated requirements.

This ability to validate available data can potentially be applied not only to the management of data, but in managing the requirements for explanations to bank customers, regulators, compliance officers and other end users. For example, a report on some business activity or proposed action, such as investment or lending decisions, may flag up an indication of whether the proposed activity would be conformant to a particular internal rule or external regulation. Simply putting up a flag to say conformance is ‘true’ or ‘false’ is not enough; there also needs to be some explanation for this result. These explanations can be derived from the semantic representations of the data in the ontology, as seen in the earlier example for front-running. In practice the various techniques of formal ontology, business rules, semantic queries and semantic ‘shapes’ can be used in combination to provide explanations to end users, data managers and other stakeholders.

5. Explanations and Logic

When considering the application of rules engines to business compliance and explanations, it is important to define what a business rule really is. It is easy to define technology-based rules engines and apply these to data in some technical ecosystem. It is also easy to claim these are ‘business rules’. However, in order for a rule to be a business rule, we should consider the nature of rules: a rule is some piece of logic, applied to something. The logic might be ‘if this, then that’ or it might be simply ‘don’t do that’. But what is the ‘something’ to which the rule is applied? If rules are applied to raw data – that is if the predicates of the rules are data in some system, then there is no guarantee that the rules represent business relationships among business concepts. For this to be the case, the predicates to which the rules are applied must themselves reflect real-world items – that is, concepts in an ontology. The predication of rules determines what kinds of rules they are – application-specific or business rules. Regulatory compliance, accountability and explainability, when they use rules, must use rules that are predicated on some ontology in order to have genuine explanatory power.

The example from Intuit goes beyond the normal usage of a ‘first-order’ ontology that simply defines what things there are and extends the knowledge graph paradigm, to connect mathematical and logical operations – similar to what we have seen with business rules and semantic shapes, but based in mathematical and arithmetic formalisms (addition, subtraction and so on). Because these features are all aligned with a formal description of 'real world' items (the ontology, assuming a quality ontology that does capture core senses of reality), any good graphical or textual representation of these relationships can be understood by any stakeholder. That is, anyone can derive explanations from a combination of first-order assertions about things in the world, formal rules that operate on those assertions, and mathematical operations on assertions that are of a numerical nature. This is the language of explanation: to the extent that end users can identify what they are seeing with the reality of their world, they should be able to interrogate semantically-enabled data to arrive at their own understanding for the reasons behind data results, decisions and other assertions that are based on that data.

6. Explaining Ontologies

The range and nature of accountability and explainability requirements in finance is indicative of the challenges and opportunities in any data-intensive industry. Formal ontologies of the business concepts are a key component to tying the data to reality and thereby to making data-driven decisions and 'what-if' analyses of proposed actions explainable and furnishing coherent explanations of a financial institution's activities to regulatory authorities.

For this kind of connection to underlying reality to work however, it is important that the ontologies used truly represent the concepts in the business domain. This cannot be approached as 'yet another data modeling' exercise; arguably it should not be approached from within the IT discipline at all. Ontologies need to reflect specialist subject matter. This means that any such ontology needs to be presented to subject matter experts in the relevant business domain, for them to validate and ideally to formally sign off. That means that the content of the ontologies needs to be explained in human-facing ways, whether through tabular views or diagrams. Ontologies are simple declarations of 'What kind of thing is this?' and 'What distinguishes it from other things?' – the necessary and sufficient characteristics for something in the world to belong to a given set of things, even if details of some relationships are harder to formalize. This can be explained relatively simply using set-theoretic notions: set membership, logical unions and intersections and so on.

These basic notions can be presented in a number of possible visual formats – typically simple boxes (or blobs) and lines showing the classes (things) and the relationships between them.

Some more complex logic features used to define the necessary and sufficient conditions are harder to explain. Instead, the ontologist needs to frame specific questions with reference to the classes, for example ‘Is it always the case that one of these has to have this relationship to one of those?’ or ‘Must there necessarily be this property or relationship in order for something to be one of those?’

Unfortunately many of the available tools tend to produce more technology-oriented visualizations, for example auto-generated pictures of ‘bouncy balls’, none of which remain where you left them last time the diagram was generated. This makes it hard to get SME confidence in the basic structure of the model. Some specialist tools do exist that provide a persistent view of the subject matter.

The alternative, which is much to be avoided, is to play into the assumption among many domain experts, that somehow words can be used to solve the problems of meaning. Vocabularies can be generated from an ontology, giving the various words that may be used to reflect a given concept in different contexts. Doing it the other way around – trying to use words as a starting point to represent concepts to the domain experts, is not a good idea. Words play games, and sets of unconnected definitions will give rise to fuzzy, overlapping and incomplete sets of things in the subject matter representation.

The matter of explaining ontologies themselves is currently very immature. Few tools exist, and many ontologies are developed to address specific data-focused application problems (drawing data inferences from existing data, answering specific questions and so on) but are often touted as though they contain some coherent account of meaning. Until these problems are better understood, it is unlikely that we will see the kinds of tools that are needed to ensure that ontologies are adequately explained to business stakeholders, and therefore adequately reflect the business reality that is needed to understand and draw explanations from the wealth of data in large institutions such as in finance.

7. Conclusions

In the world of finance there is much that would benefit from explanation. Formal approaches to explainability are not well established in the industry, but from one perspective the entire financial system can be envisioned as the ebb and flow of data; the raw material of explanations. Some of the requirements for explanations are fairly simple: the reasons for advancing or withholding credit from consumers, for example. Here the industry is moving towards better ways to partner with customers on their life journeys or business trajectories, treating explanations as something to which customers should be entitled. Other explanation requirements are more complex, dealing with the economy, macro-economics, issues of money supply and so on, along with the risk factors that accompany each of these.

Meanwhile the 2008 Global Financial Crisis reinforced an understanding that certain aspects of the global financial system form a complex adaptive (or maladaptive?) system, of the kind from which the phenomenon of ‘emergence’ gives rise to events and structures that cannot easily be anticipated in advance. Sometimes the best we can do within the parameters of complex systems theory is to explain why we cannot explain something.

In the meantime the massive flows of data in the industry admit of a couple of different purposes – as material for explanations and as something to react to. They provide the source material for accountability, a pre-requisite for explainability. Information reported to regulators can be used to understand and analyze the details that went into why someone or something made a particular decision. Similar data is analyzed by the statistics functions of central banks and used to inform those holding the economic levers of power when and whether to raise or lower interest rates, adjust the money supply and so on. Separately these data flows provide for understanding emergent risks in a system that can never be fully understood, much less explained, but that can be reacted to, given sufficient information.

A common theme in all of these uses of shared data is the need for formal business semantics. In any industry where information technology is extensively used – and many of the issues explored here can be as easily applied in healthcare and elsewhere – the information technology ecosystems used are somewhere in a transition between an older world in which each application had its own data formats and structures, feeding screens or tapes or other things read by humans and not needing to be consumed by different machines, and a future world in which every machine speaks the same language. At this point we have common syntactical formats for exchanging data but not much in the way of common understanding or language.

Confucius was asked what he would do if he was a governor. He said he would "rectify the names" to make words correspond to reality. We now find ourselves in a world where there is more data than there are words to go around, so we need to apply more sophistication to questions of meaning, something that is not an IT function at all but requires deeper business engagement. Even where 'semantics' is already being used as a term, it is often in relation to self-contained applications for drawing inferences over limited amounts of data; semantic technology stovepipes replacing rigid database stovepipes, but contributing little to the kind of common language that will be needed to furnish detailed formal explanations of the sort explored in this edition, at every level from consumer protection, through to micro- and macro-economic regulatory oversight and systemic risk mitigation.

References

Bank for International Settlements (January 2013). Principles for effective risk data aggregation and risk reporting. Retrieved September 19 2020, from https://www.bis.org/publ/bcbs239.pdf

Bennett, M. G. (2010). EDM Council Semantics Repository Next Steps – Exploring a Consensus Semantics Framework. In Proceeding: ODiSE'10 Ontology-Driven Software Engineering, Article No. 3 ACM New York, NY, USA 2010. ISBN: 978-1-4503-0548-8. DOI: 10.1145/1937128.1937131

Bholat, D. (2016). Modelling metadata in central banks. ECB Statistics Paper Series. No 13/April 2016. Retrieved 18 September 2020, from https://www.ecb.europa.eu/pub/pdf/scpsps/ecbsp13.en.pdf?a66e0b0448ff62df28d39643573b2b3d

David, E., (3 June 2016). FIBO Marches Forward: A Look Inside State Street's FIBO Proof of Concept. Waters Technology. Retrieved 18 September 2020, from https://www.waterstechnology.com/data-management/2459451/fibo-marches-forward-a-look-inside-state-streets-fibo-proof-of-concept

Enterprise Data Management Council (n.d.). Website. Retrieved September 19 2020, from https://edmcouncil.org/default.aspx

Governance Risk and Compliance Technology Centre (n.d.). Website. Retrieved September 19, 2020, from https://www.cubsucc.com/research-centres/governance-risk-and-compliance-technology-centre--grctc-/

Grosof, B. (2013). Rapid Text-Based Authoring of Defeasible Higher-Order Logic Formulas, via Textual Logic and Rulelog. In: Morgenstern, L., Stefaneas, P., Lvy, F., Wyner, A., Paschke, A. (eds.) Theory, Practice, and Applications of Rules on the Web, Lecture Notes in Computer Science, vol. 8035, pp. 2–11. Springer Berlin Heidelberg. Retrieved on 19 September 2020, from http://dx.doi.org/ 10.1007/978-3-642-39617-5_2

Grosof, B., Bloomfield, J., Fodor, P., Kifer, M., Grosof, I., Calejo, M., Swift, T. (2015). Automated Decision Support for Financial Regulatory/Policy Compliance, using Textual Rulelog. Proceedings of the RuleML 2015 Challenge, the Special Track on Rule-based Recommender Systems for the Web of Data, the Special Industry Track and the RuleML 2015 Doctoral Consortium. 9th International Web Rule Symposium (RuleML 2015). Retrieved September 19 2020, from http://ceur-ws.org/Vol-1417/paper8.pdf

International Swaps and Derivatives Association (n.d.). Financial products Markup Language (FpML). Retrieved 19 September 2020, from https://www.fpml.org/the_standard/current/

Knight, M. (January 24, 2018). What is a Business Glossary? Dataversity. Retrieved September 18 2020, from https://www.dataversity.net/what-is-a-business-glossary/

Levenshtein, V. I. (1966). Binary Codes Capable of Correcting Deletions, Insertions and Reversals. *Soviet Physics Doklady*, Vol. **10**, p.707. Bibcode: 1966SPhD...10..707L

Newman, D., & Bennett, M. (2012). Semantic Solutions for Financial Industry Systemic Risk Analysis. Next Generation Financial Cyberinfrastructure Workshop, Robert H. Smith School of Business, University of Maryland. Retrieved 18 September 2020, from https://wiki.umiacs.umd.edu/clip/ngfci/images/8/80/BennettNewman.pdf

Newman, D., (2020). ‘The Case for a Semantic Data Catalog’. CS 520 Knowledge Graphs - How should AI explicitly represent knowledge? Department of Computer Science, Stanford University, Spring 2020. Retrieved 19 September 2020, from https://web.stanford.edu/class/cs520/abstracts/newman.html

Object Management Group (2019). Semantics of Business Vocabulary and Business Rules, version 1.5. OMG Document Number: formal/2019-10-02. Retrieved on 18 September 2020, from https://www.omg.org/spec/SBVR/1.5/PDF

Odell, J. (2011). Ontology. CSC Catalyst White Paper. Computer Sciences Corporation. Retrieved September 19 2020, from http://www.jamesodell.com/Ontology_White_Paper_2011-07-15.pdf

Ontology Summit (2020). Ontology Summit 2020: Knowledge Graphs. Retrieved 18 September 2020, from https://ontologforum.org/index.php/OntologySummit2020

U.S. Electronic Code of Federal Regulations (2002). Title 12, Chapter II, Subchapter A, Part 223.

Wittgenstein, L. (1953). *Philosophical Investigations.* Macmillan.

World Wide Web Consortium (11 December 2012). OWL 2 Web Ontology Language Structural Specification and Functional-Style Syntax (Second Edition). W3C Recommendation. Retrieved 19 September 2020, from https://www.w3.org/TR/2012/REC-owl2-syntax-20121211/

World Wide Web Consortium (25 February 2014). Resource Definition Framework (RDF) Schema 1.1. W3C Recommendation. Retrieved 19 September 2020, from https://www.w3.org/TR/2014/REC-rdf-schema-20140225/

World Wide Web Consortium (20 July 2017). Shapes Constraint Language (SHACL). W3C Recommendation. Retrieved 19 September 2020, from https://www.w3.org/TR/shacl/

Yang, G., Kifer, M., Zhao, C. (2003). FLORA-2: A Rule-Based Knowledge Representation and Inference Infrastructure for the Semantic Web. Lecture Notes in Computer Science 2888:671-688. Springer. DOI: 10.1007/978-3-540-39964-3_43

Yu, J. (2020). Knowledge Graph Use Cases @ Intuit. CS 520 Knowledge Graphs - How should AI explicitly represent knowledge? Department of Computer Science, Stanford University, Spring 2020. Retrieved 19 September 2020, from https://web.stanford.edu/class/cs520/abstracts/yu.pdf

BIO

Mike Bennett is the director of Hypercube Limited, a company that helps people manage their information assets using formal semantics. Mike is the originator of the Financial Industry Business Ontology (FIBO) from the EDM Council, a formal ontology for financial industry concepts and definitions. Mike provides mentoring and training in the application of formal semantics to business problems and strategy, and is retained as Standards Liaison for the EDM Council and the IOTA Foundation, a novel Blockchain-like ecosystem.