Pre-print version of the chapter included in E. Ippoliti, M. Vergara-Fernández, & F. Zennaro (Eds.), *Philosophy and finance: Ten open questions* (Synthese Library, Vol. 504, pp. 125-140), Springer.

# Why We Need a Computational Ontology of Finance (and How Philosophers Can Help Build It)

Gloria Sansò

**Abstract** The financial crisis of 2007–2008 showed, among other things, that the information technology in the financial sector was inadequate to assist institutions in producing reliable reports and keeping track of the securities issued. This, in turn, compromised the chances of correctly estimating financial risk. In this paper, I explain how a well-designed computational ontology of finance could help improve data management. I also argue that this ontology should be realism-based and conformant to Basic Formal Ontology (BFO). This choice, however, does not come without complications because it requires some preliminary reflections on the nature of the entities involved. Fortunately, some philosophical literature can be a valuable starting point for these reflections; in the paper, this is shown by considering money as a case study.

## 1. Introduction

One of the most astonishing facts about the financial crisis of 2007–2008 is that, while some people had a clue that the crisis was coming (Beattie, 2008), most economists did not realize, in time, what was going on. Why did that happen? There are different possible answers. One may argue, for example, that social sciences, and more specifically some branches of economics, are based on some sort of metaphysical misunderstanding about the building blocks of the social world, and this prevents scholars from making proper predictions and understanding when things are turning for the worse (Epstein, 2015, pp. 2–10). Others argue that the problem is more methodological than metaphysical: by working almost exclusively on models, economists neglect what is going on in real life (Sylos Labini & Caprara, 2017, p. 68).

Although there may be some truth in both these views, it must be said that, in general, a proper estimation of the systemic risk is challenging for another reason as well. The systemic risk is the potential collapse of the entire financial system (Bodie et al., 2018, p. 21), which entails the "disruption of the market's ability to facilitate the flow of capital that results in the reduction in the growth of the global GDP" (Fouque & Langsam, 2013, p. xxi). Its precise assessment is possible *only if* one has a reliable picture of the system as a whole, but this big picture requires a plethora of data that, in the years prior to the crisis, no one

could have. On the one hand, many transactions involved shadow banks, namely financial intermediaries, that were outside the scope of government regulations and, as such, characterized by little transparency (Schoen, 2017, p. 811). As a consequence, "the CDOs market had turned into a vast, *opaque* spider web of deals in which banks, shadow banks, and brokers alike had become dangerously ensnared, interlinked by fear" [emphasis added] (Tett, 2009, p. 226).

On the other hand, there was a more general problem afflicting all the financial institutions, including the largest banks, that is:

> banks' information technology (IT) and data architectures were inadequate to support the broad management of financial risks. Many banks lacked the ability to aggregate risk exposures and concentrations quickly and accurately at the bank group level, across business lines, and between legal entities. Some banks were unable to manage their risks properly because of the weak risk data aggregation capabilities and risk reporting practices (Basel Committee, 2013, p. 1).

In the financial sector, thus, informatic systems failed in two different ways. First of all, they were inadequate to assist each bank in producing reliable reports by collecting data about itself. Secondly, the informatic systems of each institution were incapable of working with those of the others and properly exchanging information; this drastically undermined the possibility of keeping track of the securities issued. When, for example, Lehman Brothers collapsed, no one knew how many CDOs were actually in the market, which of those were toxic, and who held them (Helleiner, 2011, p. 71; Dombret, 2013, p. 2; Chadha, 2016, p. 3).

There are several computational approaches that can be used to improve the informatic systems in the financial sector (Hughes et al., 2013); among them, semantic technologies seem to be particularly promising (Butler et al., 2015; Chen, 2015). In semantic technologies, ontology is the key component; therefore, creating a successful ontology of finance is of paramount importance in order to increase transparency and facilitate regulatory compliance by financial institutions.

In this paper, after introducing semantic technologies and the role of the ontology (Sect. 2), I present the benefits (and the challenges) of a realism-based ontology of finance (Sect. 3), and I show how philosophy can help in building it by considering money as a case study (Sect. 4).

## 2. Semantic Technologies

The expression "semantic technologies" denotes those methods and tools that make the content of the information accessible to machine processes. In doing so, they greatly enhance the capabilities of the information systems. If machines are able to interpret the content of a piece of information, then they can make inferences and derive knowledge that, although not in the system, can be inferred from the information in the system (Fürber, 2016, pp. 65–66).

Suppose, for example, that a system has the information that a certain firm *A* owns firm *B*, and that *B* has a controlling interest in firm *C*, where having a controlling interest in *C* implies owning more than half of *C*'s voting shares. If this system has access to the content of "owning" and "having a controlling interest", then it can infer that *A* has a controlling interest in *C*. Secondly, if machines are able to interpret the content of a piece of information, then a system can use information present in another system (and *vice versa*), even when these systems may have different data structures, formats, and even vocabularies (SICoP, 2005, pp. 37–38). Semantic technologies can thus promote interoperability, that is, the ability of different systems to exchange information, search, query, and reason across multiple data sources. In order to better understand how all this works, it may be helpful to consider more details about the elements involved.

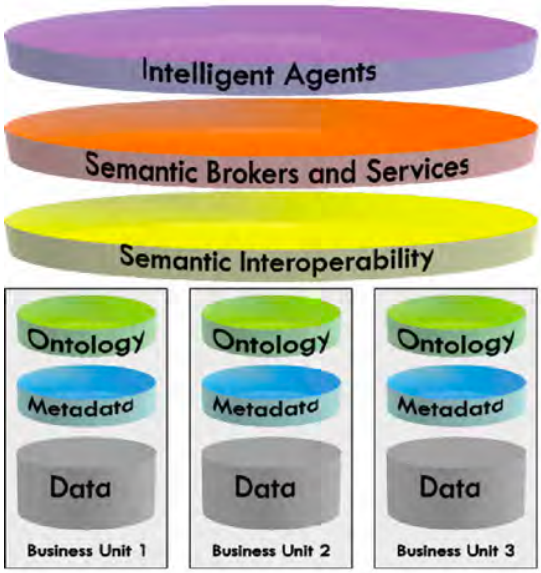


**Fig. 1** Semantic Technologies (Cregan, 2008, p. 59)

As shown in Fig. 1, in each business unit, there are data, metadata, and an ontology. Data can be both structured and unstructured. Structured data are quantitative and are stored in tables, such as an SQL database or an Excel file; some examples of structured data include dates, names, and prices. Unstructured data, instead, do not have a predefined format and are embedded in things like audio, video, or a Word file (Weglarz, 2004).

Metadata is data about data, meaning that it describes the relevant information about data. So, for example, the content of this essay is the data, whereas the size of the file, the author, and the time it was last edited are

metadata. Metadata serves as a bridge between data and ontology: by organizing metadata, an ontology helps identify and retrieve data.

In information and computer science, the term "ontology" denotes a formal representation of a portion of reality; such a representation includes a collection of terms: some of them refer to classes, whereas others refer to the relations between them. Among these relations, the *is_a* relation is particularly relevant in that it serves as a fundamental reference point for creating a taxonomy, namely a hierarchy consisting of classes and sub-classes, which is the backbone of any ontology (Arp et al., 2015, p. 1; Guarino et al., 2009, p. 2). Importantly, an ontology provides, for each of its terms, an accurate, unambiguous, and consistent semantics (Gruber, 1995, p. 11; Guarino, 1998, pp. 6–7) that is expressed in both human and machine-readable formats. Because of these characteristics, an ontology is said to have three main utilities: (i) to facilitate the communication between human beings; (ii) to promote interoperability between computer systems, and (iii) to improve the quality of pieces of software (Jasper & Uschold, 1999, p. 17).

Ontologies, thus, are the crucial element in semantic technologies because they make semantic interoperability possible; semantic interoperability, in turn, allows semantic brokers and semantic services to provide reasoning-based services; artificial agents, then, can compose these services to perform more complex tasks, such as analyzing documents and making decisions (see, for example, Joshi et al., 2019; Turner et al., 2019).

## 3. Ontology of Finance

Based on the above, it seems evident that a well-designed computational ontology of finance, by promoting semantic interoperability, would help in assessing systemic risk (see, for example, Ye et al., 2009; Bennett, 2013; Organ, 2018) and, in so doing, it could solve the problems mentioned at the beginning of this chapter. However, this is not all. Indeed, different portions of this ontology could also be used for several more specific tasks, such as providing support in financial planning (Bunnell et al., 2020), assisting in the identification of financial frauds (Kingston et al., 2004; Zhao et al., 2004; Chmielewski & Stapor, 2018), and rating companies based on their financial quality (Shue et al., 2009; Martin et al., 2011).

The question, now, is: how can one build such a computational ontology? Different methodologies can be employed: the attempts made so far to create an ontology of finance[1] are concept-based. There are, however, some general concerns about the conceptualist approach; for that reason, in what follows, I argue that a reality-based ontology of finance would be an exciting alternative and that philosophy can help overcome some of the challenges involved in this endeavor.

[1] I am referring to the Financial Industry Business Ontology (FIBO), a well-known and comprehensive project, and to the Ontology Network in Finance and Economics (OntoFINE).

### 3.1 Concept-Based Ontologies

In computer and information science, the conceptualist approach is grounded on the idea that, in order to build a valuable ontology, it is not necessary to commit to a definite metaphysic of the world. The categories employed in an ontology are simply "cognitive artifacts ultimately dependent on human perception, cultural imprints and social conventions" (Gangemi et al., 2002, p. 167). Consequently, an ontology does not represent things in reality; it just represents our concepts of these things.

Unfortunately, the conceptualist approach is problematic for several reasons.[2] One of them is that such an approach deprives us of a standard to evaluate ontologies. Concepts are indeed subjective, and if an ontology merely represents concepts, it is unclear on what basis to evaluate them. Of course, an ontology can be evaluated based on the coherence of its concepts or, alternatively, on how common these concepts are among people. However, this evaluation would still have a certain degree of subjectivity. Things entirely change if one assumes that an ontology represents things in reality because, if this is the case, then one can assess an ontology based on how accurately it represents what there is out there (Smith, 2004, p. 76).

One may counterargue that it is possible to evaluate ontologies based on their utility, meaning based on their ability to perform the task they are built for. Therefore, what ontologies represent (concepts or things in reality) and, more generally, the underlying metaphysical arguments, are secondary factors when assessing them. Ontologies are good when they successfully do what they are supposed to do and bad when they do not (Guarino, 2017, p. 11).

While an ontology certainly has to be able to perform the task it is built for, there are at least two reasons why this criterion, alone, is insufficient to evaluate ontologies properly. First, in information and computer science, as shown earlier, the aim is not simply to have ontologies performing specific tasks but to have systems able to communicate with each other. Interoperability can be achieved more easily if ontologies are consistent, meaning if they are created by referring to the same benchmark, which is reality (see, for example, Smith et al., 2007; Arp et al., 2015, p. 48; Kulvatunyou et al., 2018). Second, it is only sometimes possible to know in advance for which task an ontology will be used in the future, and this is especially true for reference ontologies. Indeed, the purpose of a reference ontology is to represent all there is in a specific domain under the assumption that, at some point, someone will use part of it to build an application ontology, an ontology with a specific task. The utility of a reference ontology lies in its ability to stand the test of time and to be reused by as many application ontologies as possible. Interestingly, it has been proven that realism-based ontologies better serve this goal than those representing concepts (Merrell et al., 2021, p. 3).

[2] For a more comprehensive overview of the main criticisms of the conceptualist approach, see Smith (2004) and Arp et al. (2015, pp. 5–15).

### 3.2 Building a BFO-Conformant Ontology of Finance

Basic Formal Ontology (BFO) is the upper-level ontology[3] developed by Barry Smith and associates; its last version, released in 2020, fulfills the requirements for top-level ontologies in ISO/IEC 21838-1, and it is thus generally acknowledged as being able to support interoperability among heterogeneous information systems.[4] Since philosophical theories were (and still are) of great support for developing BFO and its extensions, BFO can be considered the first product of philosophy to become an industry standard (Jansen & Brochhausen, 2022).

In BFO, entities are divided into two main categories: *continuants* and *occurrents*. *Continuants* are entities that continue to exist through time (some examples include my desk, the content of a book, the function of a clock); *occurrents* are entities that occur or happen (some examples include a ceremony, a war, a storm) (Arp et al., 2015, p. 87). Continuant and occurrent entities are further categorized (see Fig. 2).

Currently, BFO is used as a foundation for 420 ontologies, and it is involved in many projects[5] belonging to different fields, such as biomedicine, national security, and industry. BFO and BFO-conformant ontologies share the same methodology called ontological realism. According to this view, a good ontology is *not* a representation of concepts but a representation of reality as our best current science describes it (Smith, 2004, p. 76; Arp et al., 2015, p. 48). While it is true that, sometimes, scientific statements are false, meaning that they do not describe reality as it really is, scientific statements are typically accurate, and they are thus the best candidates for being the benchmark to create and evaluate ontologies.

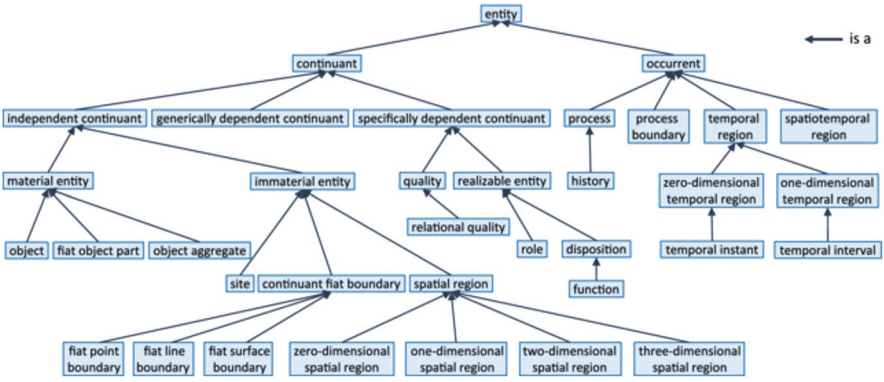


**Fig. 2.** BFO 2020: is_a hierarchy

[3] An upper-level ontology is an ontology providing very general terms, such as "object", "event", and "property", which are common to any domain. In contrast, domain ontologies are ontologies representing specific portions of reality; they thus provide terms that are peculiar to those areas. A domain ontology is typically an extension of an upper-level ontology.

[4] https://www.iso.org/standard/74572.html (last accessed May 17, 2023).

[5] http://basic-formal-ontology.org/users.html (last accessed: May 17, 2023).

Each term in a realist ontology refers to a universal or an individual. "Universals are entities in reality that are responsible for the structure, order, and regularity - the similarities - that are to be found there. To talk of universals is to talk of what all members of a natural class or natural kind such as *cell*, or *organism*, or *lipid*, or *heart* have in common" (Arp et al., 2015, p. 13). So, for example, the term "heart" refers to the collection of qualities that an object must have in order to be a heart, while a specific heart, such as Mary's heart, is an individual instantiating *that* universal. It is important to point out that a universal exists only insofar as at least one individual is instantiating it. This means that, although one can perhaps list the qualities that a thing must have in order to be a unicorn, "unicorn" cannot be a term in a realist ontology because no existing individual instantiates that universal.

The above quotation refers to natural kinds and natural entities in general. However, ontological realism can also be used as a methodology to represent social entities (see Arp et al., 2015, p. 44). Thus, the BFO framework lends itself to being a useful tool to build an ontology of finance and avoid the issues of the conceptualist approach. On the other hand, this choice does not come free of complications. Creating realism-based ontologies is *per se* more complex than building concept- based ontologies: because of their ontological commitment to entities in reality, realist ontologies cannot include, for example, terms such as "prevented abortion", "absence of metastases", or "surgical biopsy not taken", since they do not, strictly speaking, denote existing things. However, in some contexts, it may be necessary to represent these states of affairs in order to keep track of the related data; in order to do that, one has to come up with a solution and, for example, introduce new relations, such as the *lack* relation (for more details about this, see Ceusters et al., 2007; Schuler & Ceusters, 2018).

In addition to this general difficulty, there may be field-specific challenges, such as those that one must face in representing the social world. While it may be enough to leaf through a science textbook in order to understand what there is in the natural world and how to represent its phenomena, the same does not apply to the social world. Although an economics textbook explains many phenomena, such an explanation is often based on implicit presuppositions about the bricks forming the economic reality (Mäki, 2001, pp. 3–14). Therefore, to build a realist ontology of finance (and benefit from it), it is necessary to take some preliminary steps and make these presuppositions and the entities one needs to include in the ontology explicit. After that, it is necessary to explain how to represent these entities, which are primarily non-physical, in a framework, which is BFO, designed to represent especially physical and observable entities.

This is not an easy task, but in the last 30 years, social ontology as a branch of philosophy has been rapidly developing, and thus, some philosophical literature can offer (in this case, too) a valuable starting point to discuss

the nature of the financial entities and overcome the challenges implied by a realist approach. In what follows, I shall show this by considering money as a case study: although providing a complete theory goes beyond the scope of this chapter, I outline what can be seen as a theoretical foundation for the representation of money in a BFO-compatible ontology of finance.

## 4. Money

The nature of money has been investigated over the centuries by prominent scholars of various disciplines, such as philosophy, anthropology, economics, and sociology. In the last few decades, money has also been one of the most debated topics in social ontology; in what follows, I discuss some of the views that emerged.

### 4.1 Money and Its Representations

John Searle is well-known for introducing the formula *X counts as Y in C*, through which he aims to account for institutional facts (or, at least, for some of them) (1995, 2010). A certain object is a dollar bill because, he says, people collectively assign a new status to that object, that is the status of dollar bill, and, in virtue of this new status, this object can perform certain functions in a given context, such as the function of means of exchange in the United States. One may thus say that the piece of paper in my pocket *counts as* a dollar bill in the United States.

Although this formula has some plausible appeal, it has been challenged in many ways; one of the most compelling objections is advanced by Barry Smith. Smith (2003, pp. 24–27) points out that there are many social entities that do not coincide with any physical object X; some examples include debts, obligations, rights, and digital money. "The blips in the bank's computer merely *represent* money, just as the deeds to your property merely *record* or *register* the existence of your property right. The deed is not identical with your property right and nor does it count as your property right" (Smith, 2003, p. 20). These entities are called by Smith "free-standing Y terms" or "quasi-abstract" because, like abstract entities, they are non-physical and non-psychological but, on the other hand, they are tied to time, so they are historical (Smith, 2008, p. 37).

Based on that, it seems that, in addition to things that are actually money, such as coins and banknotes, there are things, such as the blips in the bank's computer or the ink on a bank statement, that just represent money. If one assumes that "representing" means being about some portion of reality (Arp et al., 2015, p. 3) and that coins and banknotes are money, then "representing money" means being about some banknotes and coins (maybe those banknotes and coins that I have the right to withdraw from the bank based on my most updated bank statement). The problem with this view is that, as it is, it cannot accommodate the case of Central Bank Digital Currencies (CBDCs), which could, over the long term, completely supplant paper cash.[6]

[6] The same applies to cryptocurrencies, such as Bitcoin (assuming that one wants to consider them as money).

Although governments do not have concrete plans in that direction, one can suppose that, at some point, banknotes and coins will become obsolete, and transactions will occur electronically only. If so, then it is not clear what a bank statement would represent since banknotes and coins would cease to exist as money. A possible way to solve this problem is to say that money is something other than the objects used as a means of exchange. Therefore, a bank statement does not represent banknotes and coins, but something else. It remains to be clarified, though, what "something else" is.

### 4.2 Institutions and Causal Powers

Guala (2020) argues that institutions are ontologically more fundamental than the devices used within them; this means, for example, that money-as-an-institution has ontological priority over money-as-an-object, that is the object used as money, such as coins and banknotes. Consequently, to understand what money is (and thus, what a bank statement represents), it is a good idea to consider money-as-an- institution first. Institutions are defined as "*rules in equilibrium*, or rules that people are motivated to follow" (Guala, 2020, p. 272). Indeed, while there is no doubt that rules form an institution and that some sort of authority explicitly issues some of them, these rules alone do not make people observant. In order to follow the rules, people need an incentive, as the fact that following the rules makes their lives easier. A person living in the US, for example, is motivated to use dollar bills not (just) because there are some laws prescribing this but because, since most people in the US are using dollar bills, this person finds it convenient to use dollar bills as well. In this context, using dollar bills allows this person to enjoy a number of benefits, such as participating in smooth transactions, sharing a unit of account, and so on. "If the system of expectations is coherent, the institution is self-sustaining, and our lives are a lot simpler if we follow the rules" (Guala, 2020, p. 273).

If that is the case, then there is little point in investigating the nature of banknotes and coins to understand what money is because (i) they are just two monetary devices among many others, and (ii) monetary devices become ontologically interesting once they are placed within an institution since that would help to understand what they do. Instead of merely focusing on banknotes and coins, it is much more promising to zoom out and better understand the relations between these objects and the institution to which they belong. In a way, this is what Mäki tries to do: rather than specific instantiations, he focuses on "moneyhood", the collection of properties shared by *all* monetary devices, and defines it as "a set of *institutionally* and *practically* sustained causal powers" [emphasis added] (Mäki, 2021, p. 253).

In metaphysics, causal powers are dispositional properties, meaning properties that, under certain circumstances, make objects behave in a certain way; some examples of dispositional properties include being inflammable, being fragile, and being soluble. Causal powers are typically used to offer an

account of causation: the solubility of a sugar cube, that is, the disposition of a sugar cube to dissolve itself, accounts for the change involving the cube when it is placed in non- saturated water.

Mäki believes that something similar can be claimed of the objects instantiating moneyhood, and this is also reflected by the fact that money-as-an- object is commonly said, for example, to have purchasing *power*. Unlike the powers a sugar cube possesses, the powers exhibited by monetary devices do not derive, of course, from their physical structure but rather from their placement within an institution, which attributes those powers to these objects. In this framework, a bank statement does not represent banknotes and coins but the collection of powers somehow carried by these monetary devices.

This view has, I think, two main benefits. First, it recognizes the role of money- as-an-institution in determining the properties possessed by money-as-an-object. Second, it accounts for the modal dimension of monetary devices: like sugar cubes, these devices have potentialities, such as that of facilitating people's exchanges, that manifest themselves only when certain conditions occur. On the other hand, Mäki's view seems to imply the existence of a "thing" carrying these powers, and this is highly problematic in the case of digital currencies, such as CBDCs, where there is no money-as-an-object. In addition (and more generally), appealing to causal powers in order to account for social phenomena is a controversial choice that brings with it several legitimate concerns.[7] The most relevant, for the purpose of this chapter, is that causal powers (and dispositional properties in general) are typically said to be internal,[8] meaning that they are possessed by an object just in virtue of itself. This is clearly incompatible with the case of money since the objects instantiating moneyhood possess their qualities just in virtue of an institution and not in virtue of themselves; these qualities are thus external. While this does not seem to be the biggest problem of all (Mäki himself, as well as a minority of philosophers, does not believe that causal powers *have* to be internal properties), it prevents one from using Mäki's theory as a theoretical foundation to represent money in the BFO framework. Indeed, in BFO, like in most of the philosophical literature, an entity is a disposition if it is internally grounded, meaning grounded in the physical make-up of an object (Arp & Smith, 2011, pp. 6–7) and, clearly, the powers of money do not fulfill this condition.

### 4.3 Information and Roles

My proposal is to give an account of money by considering *deontic* powers, meaning abilities that an individual has (or does not have) in virtue of her social role and some regulations. According to Searle (1995, pp. 84–85), some of these deontic powers are positive, whereas others are negative. Rights and permissions are instances of the former; they are *positive* because they make new courses of action possible. Obligations and prohibitions are, instead,

[7] These concerns are discussed by Wahlberg (2018).

[8] See, for example, Harré (1970); Molnar (2003); Bird (2007); Marmodoro (2017); Williams (2019).

instances of the latter; they are *negative* because they limit people's actions.

Deontic powers allow us to account for modality in the social world without, however, the complications implied by causal powers. Consider, for example, my right to vote: in virtue of this deontic power, when certain conditions are obtained, meaning that there are elections in my home country, I have the possibility to perform a certain action, that is, voting. Something similar can be said about money. In virtue of the right to use a certain currency, when certain conditions are obtained, meaning that there is a commercial opportunity and I have a certain amount of that currency, I have the possibility to perform a certain action, that is, purchasing. Like Mäki, thus, I claim that money involves powers that are "institutionally and practically sustained"; unlike Mäki, I claim that these powers are (i) *deontic,* not *causal*, and (ii) possessed by individuals, not by monetary devices.

The nature of these deontic powers remains to be clarified. Based on our ordinary talk, it seems that expressions such as "right to vote" or "property right" denote (somehow abstract) things that people can actually own. Following Donohue (2020, pp. 139–158), however, I opt for a different approach, which is, I believe, ontologically more reasonable insofar as it refrains from asserting the existence of metaphysically controversial entities. I argue that deontic powers are not individual things but states of affairs, which include regulations, people holding certain roles, and the relations between these two things. The sentence "I have the right to vote in Italy", for example, does not refer to a situation where there is a thing, the right to vote, owned by me. Rather, this sentence refers to a situation where (i) I fulfill the conditions to have the role of voter, (ii) there is some regulation, such as the Constitution, that defines the actions I am allowed to perform in virtue of this role, and (iii) people around me are motivated to follow that regulation and allow me to perform the actions provided for.

The same can be applied to money. The sentence "I have the right to use currency X" does not refer to a situation where there is an individual thing that can be actually owned by people. Such an expression refers, instead, to a state of affairs that includes the rules forming the institution of money as well as those governing the commercial exchange, people holding certain roles, and the relations between these things. Therefore, the sentence "I have the right to use currency X", refers to a situation where (i) I fulfill the conditions to have the role of owner of currency X, (ii) there is some regulation that allows me to perform some actions, such as purchasing, in virtue of my role and by means of currency X, and (iii) people around me are motivated to follow that regulation and allow me to perform the actions provided for.

What I have said so far about deontic powers accounts for money-as-an- institution, but what about money-as-an-object? I argue that monetary devices, such as coins and banknotes, can be seen as carriers of information: they specify an amount of currency and, together with other information, such as the purchasing power of that specific currency, define the range of things that

one can acquire by means of them. In the case of a CBDC, where there are no banknotes and coins, one may assume the existence of documents such as bank statements that, similarly to banknotes and coins, convey information about the amount of currency possessed by each individual. Although these documents are not instances of money-as-an-object, and they cannot be exchanged as banknotes and coins, they represent the possibilities implied by the amount of currency the account holder possesses. This, along with the rules forming the institution (and the willingness of people to follow them), is enough to put into existence a state of affairs where people have (or do not have) the ability to participate in transactions, even in the absence of banknotes and coins.

My proposal, which prioritizes money-as-an-institution and characterizes money-as-an-object in terms of carrier of information, effectively addresses the challenges related to the absence of money-as-an-object illustrated in 4.1 and 4.2. It also clarifies what documents such as bank statements represent. The second advantage of my proposal is that it relies on types of entities that can be easily included in a BFO-conformant ontology. More specifically, my account relies mostly on pieces of information carried by physical objects and roles held by people. In the context of BFO, a piece of information can be seen as an abstract pattern, such as a pattern of letters. This abstract pattern is concretized in some of the qualities possessed by the material support that is carrying the information; some examples include ink marks, in the case of a piece of paper, and magnetic traces, in the case of the disk of a computer (Arp et al., 2015, pp. 105–107). Also, a piece of information is always about something (Ceusters & Smith, 2015, p. 2).

## 5. Conclusions

In this paper, I have shown that, in the financial sector, a more efficient information technology would facilitate data management and assist institutions in producing more reliable reports, keeping track of the securities issued, and better assessing systemic risk. One way to make this improvement is by building a computational ontology of finance. I have argued that since realism-based ontologies have often proven more successful than concept-based ontologies, building a realism-based ontology of finance would be an interesting endeavor. In particular, I have suggested Basic Formal Ontology (BFO) as the appropriate framework for building such an ontology. This choice, however, carries with it some challenges insofar as it requires some preliminary steps in order to identify a way to account for financial entities that is satisfactory as well as BFO-conformant. By considering money as a case study, I have shown that some philosophical literature is a valuable support in that process. Based on some recent work on the ontology of money, I have accounted for money and its modal dimension in terms of deontic powers and deontic powers in terms of states of affairs populated by two main entities: the roles possessed by the individuals and the pieces of information carried by regulations and monetary devices.

**References**


Arp, R., & Smith, B. (2011). Function, role and disposition in Basic Formal Ontology. *Nature Precedings*, 1–4.

Arp, R., Smith, B., & Spear, A. (2015). *Building ontologies with basic formal ontology*. MIT Press.

Basel Committee. (2013). *Principles for effective risk data aggregation and risk reporting*. https:// www.bis.org/publ/bcbs239.pdf

Beattie, A. (2008, November 14). *Good question, Ma'am*. Financial Times.

Bennett, M. (2013). The financial industry business ontology: Best practice for big data. *Journal of Banking Regulation, 14*(3–4), 255–268.

Bird, A. (2007). *Nature's metaphysics: Laws and properties*. Oxford University Press.

Bodie, Z., Kane, A., & Marcus, A. J. (2018). *Investments* (11th ed.). McGraw-Hill Education.

Bunnell, L., Osei-Bryson, K., & Yoon, V. (2020). Development of a consumer financial goals ontology for use with FinTech applications for improving financial capability. *Expert Systems with Applications, 165*(3), 1–14.

Butler, T., Abi-Lahoud, E., & Espinoza, A. (2015). Designing semantic technologies for regulatory change management in the financial industry. *ICIS 2015 Proceedings, 8*, 1–11.

Ceusters, W., Elkin, P., & Smith, B. (2007). Negative findings in electronic health records and biomedical ontologies: A realist approach. *International Journal of Medical Informatics, 76*, 326–333.

Ceusters, W., & Smith, B. (2015). Aboutness: Towards foundations for the information artifact ontology. *Proceedings of the Sixth International Conference on Biomedical Ontology (ICBO*) *CEUR, 1515*, 1–5.

Chadha, P. (2016). What caused the failure of Lehman brothers? Could it have been prevented? How? Recommendations for going forward. *International Journal of Accounting Research, s1*, 1–5.

Chen, J. (2015). *Evaluation of application of ontology and semantic technology for improving data transparency and regulatory compliance in the global financial industry* (Master's thesis). Massachusetts Institute of Technology.

Chmielewski, M., & Stapor, P. (2018). Hidden information retrieval and evaluation method and tools utilising ontology reasoning applied for financial fraud analysis. *MATEC Web of Conferences*, 210.

Cregan, A. (2008). *Weaving the semantic web: Contributions and insights* (Doctoral dissertation).
University of New South Wales Sydney].

Dombret, A. (2013). Lehman lessons. *International Economy, Fall 2013*, 34–36.

Donohue, B. (2020). *Social ontology and social normativity* (Doctoral Dissertation). University at Buffalo.

Epstein, B. (2015). *The ant trap: Rebuilding the foundations of the social sciences*. Oxford University Press.

Fouque, J., & Langsam, J. A. (2013). Introduction. In J. Fouque & J. A. Langsam (Eds.), *Handbook on systemic risk* (pp. xx–xxv). Cambridge University Press.

Fürber, C. (2016). *Data quality management with semantic technologies*. Springer.

Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., & Schneider, L. (2002). Sweetening ontolo- gies with DOLCE. In A. Gómez-Pérez & V. R. Benjamins (Eds.), *Knowledge engineering and knowledge management: Ontologies and the semantic web. EKAW 2002* (*Lecture Notes in Computer Science*) (pp. 166–181). Springer.

Gruber, T. R. (1995). A translation approach to portable ontologies. *Knowledge Acquisition, 5*(2), 199–220.

Guala, F. (2020). Money as an institution and money as an object. *Journal of Social Ontology, 6*(2), 265–279.

Guarino, N. (1998). Formal ontology and information systems. In N. Guarino (Ed.), *Formal ontology in information systems* (pp. 3–15). IOS Press.

Guarino, N. (2017). BFO and DOLCE: So far, so close. *Cosmos+Taxis, 4*(4), 10–18.

Guarino, N., Oberle, D., & Staab, S. (2009). What is an ontology? In S. Staab & R. Studer (Eds.),
*Handbook on ontologies* (pp. 1–17). Springer.

Harré, R. (1970). Powers. *The British Journal for the Philosophy of Science, 21*(1), 81–101.

Helleiner, E. (2011). Understanding the 2007–2008 global financial crisis: Lessons for scholars of international political economy. *Annual Review of Political Science, 14*, 67–87.

Hughes, E., Rosenthal, A., & Worrell, C. (2013). Enabling data analysis for addressing systemic risk. In J. Fouque & J. A. Langsam (Eds.), *Handbook on systemic risk* (pp. 852–868). Cambridge University Press.

Jansen, L., & Brochhausen, M. (2022, September 28). *Begriffe in Schichten: Die 'Basic Formal Ontology' erbringt Dienstleistungen für die Medizin*. Das wurde jetzt mit der Anerkennung als Industriestandard gewürdigt. Frankfurter Allgemeine Zeitung.

Jasper, R., & Uschold, M. (1999). A framework for understanding and classifying ontology applications. In *Proceedings 12th international workshop on knowledge acquisition, modelling, and management* (pp. 16–21).

Joshi, K., Pande, K., & Mittal, S. (2019). A semantic approach for automatic knowledge in policies of cyber insurance services. In 2019 *IEEE international conferences on web services* (pp. 33– 40).

Kingston, J., Schafer, B., & Vandenberghe, W. (2004). Towards a financial fraud ontology: A legal modelling approach. *Artificial Intelligence and Law, 12*(4), 419–446.

Kulvatunyou, B., Wallace, E., Kiritsis, D., Smith, B., & Will, C. (2018). The industrial ontologies foundry proof-of-concept project. *APMS 2018: Advances in Production Management Systems.* Smart Manufacturing for Industry *4.0,* pp. 402–409.

Mäki, U. (2001). Economic ontology: What? Why? How? In U. Mäki (Ed.), *The economic world view: Studies in the ontology of economics* (pp. 3–14). Cambridge University Press.

Mäki, U. (2021). Reflections on the ontology of money. *Journal of Social Ontology, 6*(2), 245–263.

Marmodoro, A. (2017). Aristotelian powers at work: Reciprocity without symmetry in causation. In J. D. Jacobs (Ed.), *Causal powers* (pp. 57–76). Oxford University Press.

Martin, A., Manjula, M., & Venkatesan, P. (2011). A business intelligence model to predict bankruptcy using financial domain ontology with association rule mining algorithm. *International Journal of Computer Science Issues, 8*(3), 211–218.

Merrell, E., Kelly, R., Kasmier, D., Smith, B., Brittain, M., Ankner, R., Maki, E., Heisey, C. W., & Bush, K. (2021). Benefits of realist ontologies to systems engineering. In *8th international workshop on ontologies and conceptual modelling (OntoCom)* (pp. 1–10).

Molnar, G. (2003). *Powers: A study in metaphysics*. Oxford University Press.

Organ, J. (2018). The control of human factors in catastrophic financial systemic risk: A case study of the Irish banking crisis. *IFAC-PapersOnLine, 51*(30), 580–585.

Schoen, E. J. (2017). The 2007–2009 financial crisis: An erosion of ethics: A case study. *Journal of Business Ethics, 146*, 805–830.

Schuler, J. C., & Ceusters, W. M. (2018). The problems of realism-based ontology of design: A case study in creating definitions for an application ontology for diabetes camps. *AMIA Annual Symposium Proceedings, AMIA Symposium, 2017*, 1517–1526.

Searle, J. (1995). *The construction of the social reality*. Free Press.

Searle, J. (2010). *Making the social world*. Oxford University Press.

Shue, L., Chen, C., & Shiue, W. (2009). The development of an ontology-based expert system for corporate financial rating. *Expert Systems with Applications, 36*(2), 2130–2142.

SICoP. (2005). *Introducing Semantic Technologies and the Semantic Web Vision.*https:// staff.icar.cnr.it/pontieri/didattica/SIIR/lezioni/SemanticWeb.pdf

Smith, B. (2003). John Searle: From speech acts to social reality. In B. Smith (Ed.), *John Searle* (pp. 1–33). Cambridge University Press.

Smith, B. (2004). Beyond Concepts: Ontology as Reality Representation. In A. Varzi, L. Vieu (eds.), *Proceedings of Fois 2004: International Conference on Formal Ontology and Information systems* (pp. 1–12).

Smith, B. (2008). Searle and de Soto: The new ontology of the social world. In B. Smith, D. Mark, & I. Ehrlich (Eds.), *The mystery of capital and the construction of social reality* (pp. 35–51). Open Court.

Smith, B., Ashburner, M., Rosse, C., Bard, J., Bug, W., Ceusters, W., Goldberg, L. J., Eilbeck, K., Ireland, A., Mungall, C. J., OBI Consortium, Leontis, N., Rocca-Serra, P., Ruttenberg, A., Sansone, S. A., Scheuermann, R. H., Shah, N., Whetzel, P. L., & Lewis, S. (2007). The OBO foundry: Coordinated evolution of ontologies to support biomedical data integration. *Nature Biotechnology, 25*(11), 1251–1255.

Sylos Labini, F., & Caprara, S. (2017). On the description of financial markets: A physicist's viewpoint. In E. Ippoliti & P. Chen (Eds.), *Methods and finance: A unifying view on finance, mathematics and philosophy* (pp. 67–71). Springer.

Tett, G. (2009). *Fool's gold: How unrestrained greed corrupted a dream, shattered global markets and unleashed a catastrophe*. Little, Brown and Company.

Turner, C. J., Emmanouilidis, C., Tomiyama, T., Tiwari, A., & Roy, R. (2019). Intelligent decision support for maintenance: An overview and future trends. *International Journal of Computer Integrated Manufacturing, 32*(10), 936–959.

Wahlberg, T. H. (2018). Causal powers and social ontology. *Synthese, 197*(3), 1357–1377. Weglarz, G. (2004). Two worlds of data - unstructured and structured. *DM Review*.

Williams, N. (2019). *The powers metaphysic*. Oxford University Press.

Ye, K., Wang, S., Yan, J., Wang, H., & Miao, B. (2009). Ontologies for crisis contagion management in financial institutions. *Journal of Information Sciences, 35*(5), 548–562.

Zhao, G., Kingston, J., Kerremans, K., Coppens, F., Verlinden, R., Temmerman, R., & Meersman, R. (2004). *Engineering an ontology of financial securities fraud* (Lecture Notes in Computer Science) (pp. 605–620). Springer.

.