

Harmony: Semantics for FinTech

Stijn Verstichel¹ ^a, Thomas Blommaert², Stijn Coppens², Thomas Van Maele², Wouter Haerick^{1,2},
and Femke Ongenaë¹ ^b

¹*IDLab, Department of Information Technology at Ghent University - imec, iGent Tower - Tech Lane Science Park 126, B-9052 Ghent, Belgium*

²*Harmony, Tech Lane Science Park 122, B-9052 Ghent, Belgium*

Keywords: FinTech, Information Integration and Interoperability, Financial Education.

Abstract: As a result of legislation imposed by the European Parliament, in order to protect inhabitants from being exposed to a too high financial risk when investing in a variety of financial markets and products, Financial Service Providers (FSPs) are obliged to test the knowledge and experience of potential investors. This is oftentimes done by means of questionnaires. However, these questionnaires differ in style and structure from one FSP to the other. The goal of this research is to manage in a more cost-effective manner (aligned with the needs and competencies of the individual financial investor in terms of products and services) the management of the private equity and to facilitate the fine-tuned personalised financial advisory services needed. This is achieved by means of a knowledge-based approach, integrating the available information of the investor (e.g. personal profile in terms of financial knowledge and experience) and for an extendable amount of financial service providers with their financial products and demonstrated by a number of exemplary use case scenarios.

1 INTRODUCTION


The best decisions on financial investments are those supported by a personalised, holistic understanding of the investor. Information such as financial knowledge, experience and risk willingness are key topics in this profile. Typical ways to collect this information are by means of questionnaires – be it online or on paper – to be filled out by potential investors. Questions found in such questionnaires are a.o. (i) “Do you have an investment goal or are you looking for a specific product?”, (ii) “How old are you? What is your marital status?”, (iii) “Are bonds a no-risk investment?”, (iv) “Have you invested previously on the stock market?”, or (v) “What is your current professional status?”.


myHarmony (myHarmony, 2020), a fast growing Belgian FinTech startup, offers unique digital tools to support information collection and aggregation and to ease the burden on all stakeholders involved. All financial professionals who need to comply with the MiFID II (European Parliament, 2014a) and MiFIR (European Parliament, 2014b) regulation can benefit from the myHarmony platform.

One of the possible approaches to reach this personalised holistic understanding and support of the financial investor to be, is through the creation of a personalised profile, i.e. a one-stop-shop for all information related to that investor. Financial service providers have to comply with all rules and legislation and as such are obliged to have their clients fill out specific questionnaires, defining their level of financial knowledge and experience. This in turn ensures that appropriate products and services can be offered.

However, this can be a repetitive and boring task if a client wants to invest with multiple financial service providers. The goal of our research has been to validate semantic technologies to extract information from these information sources and how to meaningfully pre-populate such questionnaires. Additionally, we have investigated how gaps in financial knowledge can be discovered and appropriate informative material can be suggested to the client. In such a way, the financial literacy of the stakeholders can be improved.

The remainder of this paper is structured as follows. Section 2 presents related work for the semantic model that is detailed in Section 3. The database and architecture of the platform is described in Section 4, including the evaluation of a few semantic processing mechanisms that are adopted to successfully ac-

^a  <https://orcid.org/0000-0001-8007-2984>

^b  <https://orcid.org/0000-0003-2529-5477>

comply with the tasks. Section 5 discusses the details of how data quality is maintained. Summarised, this paper aims to give the reader a hint of how the starting point of facilitating a more holistic approach to financial service provisioning can be achieved through the adoption of semantic technologies and how it might be integrated with existing systems, both front-end as well as back-end.

2 RELATED WORK

A number of potential data integration mechanisms exist. Firstly, one could integrate individual applications by means of re-implementing them in one domain-wide application. Secondly, Application Programming Interfaces (API) could be exposed as well offering application developers the possibility to use one another's applications and business information. Lastly, the integration could also be done at data level. This means that the individual applications continue to be developed independently, but a commonly agreed domain model is established to exchange information between the collaborating stakeholders. For wider specifications of system integration, as is the case in the financial domain, the preferred solution for the integration of data is one that avoids any major alteration to existing system design. Further to this, there is a requirement to enable systems to provide data to other stakeholders that can be extended at any time without major redesign. This interaction can be achieved through the implementation of a common vocabulary that forms the foundation for communication between applications.

One such initiative is the Financial Industry Business Ontology (FIBO) (Bennett, 2013). It defines the entities and interactions/relations of interest to financial business applications. FIBO is hosted and sponsored by the Enterprise Data Management Council (EDM) and is a trademark of EDM Council, Inc. It is also standardised by the Object Management Group (OMG). As elaborated by EDM, FIBO is developed as an ontology in the Web Ontology Language (OWL). The use of logic ensures that each FIBO concept is framed in a way that is unambiguous and that is readable both by humans and machines.

Apart from developing wider semantic technology-based platforms and solutions, Ontotext (Ontotext, 2020) has previously made the case for the widespread adoption of these technologies in the FinTech domain. Moreover, as part of the 'Six Markers of Success' in their McKinsey report 'Cutting through the fintech noise: Markers of success, imperatives for banks' (Dietz et al.,

2016), the authors proclaim *Innovative uses of data*. One such initiative is, e.g. presented through the FintechBrowser (Abi-Lahoud et al., 2017) – a demonstrator of an integrated set of dashboards and a graph explorative-browsing tool. Rather than claiming a contribution to the general development of semantic platforms, not only specifically for the FinTech industry, but even for wider information integration purposes, this research focuses precisely on a specific, highly pertinent issue, as demonstrated by the MiFID II (European Parliament, 2014a) and MiFIR (European Parliament, 2014b) regulation, and wants to underline the appropriateness of semantic technologies for information integration and analysis.

3 SEMANTIC MODEL

This section presents the actual, in a scalable manner engineered, ontology that is suitable to interpret investment-related information. The architecture, the ontology and, in case appropriate, the mapping language will be detailed. As scoped in Section 1 the ultimate goal is to generate and pre-populate questionnaires to define the knowledge and expertise of potential investors.

3.1 Ontology Model

Initially, two separate sub-ontologies were designed, namely one representing the questionnaires (both the actual questionnaire with the questions as well as the filled-out instance by an investor) and the other ontology containing all information concerning the investment services and products, its characteristics and rules and regulations. The domain knowledge for the first ontology model was manually extracted from PDF documents, while the input for the second ontology model is located in a relational database.

Later on, these models were further extended with a more detailed structure of the questionnaires, so that automatic questionnaire generation can be supported, as well as the partitioning between individual financial service providers. Questionnaires have been modelled as a combination of an empty questionnaire template on the one hand and a completed version by an investor on the other hand. Questions in the questionnaires have two important relationships for this research, namely the topic concerned and the product which the question asks information about.

Questions in the questionnaires have two important relationships for this research, namely the topic concerned and the product which the question asks information about.

In a similar fashion all (potential) answers are instantiated into the ontology model. Apart from modelling the decomposition of a questionnaire into individual questions and answers, linked to the products and topics, one of course needs to be aware of correct and incorrect answers. This to be able, in a later stage, to suggest questions in an automated manner and contextualised to the situation of a given person. As such, the model as presented earlier can be fully summarised as in Figure 1.

Surveying through all available questionnaire PDF documents and using the expertise of the partners in the project, the taxonomy for the ‘Profile-Topics’ has been defined. Just over 200 questions have been analysed, and this has resulted in multiple topics being attributed to individual questions. The main modelling principle to facilitate capturing the knowledge and experience of a person allowing to suggest other questions is presented in Figure 2. Here, ‘Experience’ and ‘Knowledge’ represent placeholders for specific ‘Topics’ as discussed earlier in this section.

Depending on the (set of) answers given by a particular person/investor, an investment profile can be affirmed or inferred for that person. An example definition of a ‘ConservativeProfile’ could be that a ‘Person’ has answered a question of a ‘Questionnaire’ with the name ‘INGQuestionnaireSummer2017_Q1_A1’, should be realised as an individual of the concept ‘ConservativeProfile’.

The actual ‘Profiles’ to be modelled for all banks/financial service providers to be considered can be extended, as can the rules and axioms that define the automatic realisation. Ontologies allow to describe anything about anything. Thus, to enforce some constraints, two main methods are generally accepted:

- Usage of DOMAIN and RANGE specification for properties, e.g. ‘Entity’ ‘hasProfile’ ‘Profile’, and
- LOGIC AXIOMS, e.g. ‘Person’ ‘hasProfessionalActivity’ some ‘ProfessionalActivity’.

To illustrate that not all information in the knowledge model should be entered manually, the hierarchical modelling of the known products and services in the relational database provided by Harmony has been imported using the OnTop (Calvanese et al., 2017) mechanism. This mechanism is further detailed in Subsection 4.1. So in detail, one implies that the individual (*i*) is of a certain type, i.e. specified by the ‘code’ column, (*ii*) is a product of a certain service provider, (*iii*) has a required minimal knowledge test score threshold, and (*iv*) has a minimal experience test score threshold; the latter two being a datatype relationship to values of the type ‘decimal’.

4 DATABASE AND APPLICATION ARCHITECTURE

4.1 Importing Information from a Relational Database

This section details the adopted technology stack and corresponding application architecture.

Obviously, one needs to create a semantic model first, i.e. the ontology. The approach adopted for this research is first to create a clean-field proprietary ontology model, according to the OWL2 standard. In a follow-up stage, alignment with existing models, such as e.g. FIBO (Bennett, 2013) can be provided, due to the inherent nature of ontology models and the OWL modelling language.

Having defined the ontology model in terms of its concepts, relationships and axioms, the raw data (i.e. individuals) need to be imported. This can be done manually, but of course much data already exist in proprietary systems. The example illustrated here is one where the data is stored in a relational database, namely an instance of MySQL. Mappings can be defined, e.g. using the OnTop engine, that specify how the cells in the tables of the relational database can be projected to become individuals of concepts in the ontology, together with their properties and relationships with other cells. As such a virtual RDF graph is created, allowing semantic querying in the same manner as if it were executed on a native RDF graph. In this context, RML (Dimou et al., 2014) needs mentioning as well, as this broadens the same mechanism for a variety of underlying persistency mechanisms.

4.1.1 Augmenting Imported Data with Additional Knowledge through the Use of JSON-LD

Apart from importing all data from a relational database, it should be considered that in certain situations, either not all information is available in the database, or manual enhancements/improvements on the imported data need to be performed. Several well established technologies exist to achieve this goal, such as by using spreadsheets, text files, XML documents, editors, etc. One particular technology of interest is based on the well-known serialisation format JSON, namely JavaScript Object Notation for Linked Data (JSON-LD). As its name suggests, it is an extension on the JSON format, and thus also backwards compatible with it. Information from existing systems or repositories that can produce JSON documents, can in a straightforward manner be extended to produce JSON-LD. Indeed, converting JSON into JSON-LD

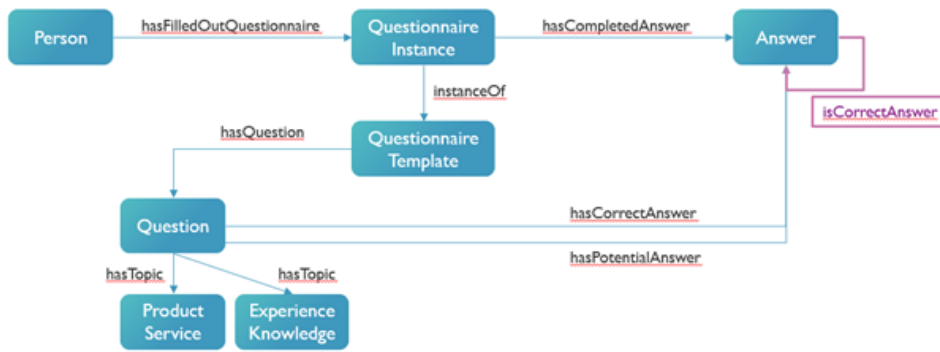


Figure 1: Decomposition of Questionnaires.



Figure 2: Three-way relation between ‘Topic’, ‘Product’ and ‘Question’.

only requires the addition of a concept reference. The example in Subsection 4.1.2 presents this approach in more detail and arguments why in a semantically rich environment it is of great additional benefit over the standard benefits when using plain JSON.

One particular important use case in the context of this research trajectory is the augmentation of existing information about topics and subtopics of the questions in available questionnaires with manual extensions, defined by domain experts. This can be done in parallel with the ongoing database development, without interfering with it. Currently, the Harmony RDB only contains information about four high-level topics, namely: (i) Financial Knowledge, (ii) Financial Experience, (iii) Financial Situation, and (iv) Investment Goals. However as indicated earlier, more detailed classification about the topics concerned by certain questions in the questionnaire is needed. This supports the use case of more intelligent and personalised testing of products, topics and knowledge of an investor.

One way to support the domain expert in specifying this more fine-grained modelling is through the supply of independent information documents. Those can be generated by a proprietary tool, or manually by the domain expert itself. Positive experience has been gained in the past with using JSON-LD. Its natural symbiosis with semantics allows the import of the JSON-LD documents in a straightforward and generic manner. This in contrast to e.g. other formats such as CSV, which require more intervention to get it in line with the semantics of the existing model. Therefore, the JSON-LD serialisation format is used to demon-

strate the process of extending existing information with more detail. It should be kept in mind, of course, that JSON-LD is only a means to an end and not a goal in itself.

4.1.2 High-level Information Imported by Means of the OnTop Procedure

Indeed, this is what the OnTop processor does, using the mapping specification that is provided by the system engineer. If we now want to include more detailed information into the knowledge model about the fact that this question does not only consider financial knowledge, but also the topic ‘Liquidity’, a separate JSON-LD document could be provided with this information. In this way, the original mapping can be maintained, but in parallel can be augmented with the information that it is not only about the high level notion of ‘Financial Knowledge’, but more detail about the ‘Liquidity’ of the product. This additional document could then be formatted as follows:

```

{
  {1} "@context":
    "http://idlab.technology/harmony/",
  {2} "@type": "Question",
  {3} "@id": "question_170",
  {4} "hasTopic": "Liquidity",
  {5} "hasQuestionText": "Money market funds
  always offer a capital guarantee."@EN
  {5} "hasQuestionText": "Geldmarktfondsen
  bieden altijd een kapitaalgarantie."@NL
}
  
```

What we can see in this sample is that first a static reference is made towards the ontology URI 1 that is being referenced, namely that of the developed Harmony ontology: <http://idlab.technology/harmony/>. Next, 2 defines that the subsequent information and relations define knowledge about an instance of the type ‘Question’, as defined in the ontology referenced in 1. Rather than using the




Figure 3: Dataflow, augmenting the asserted data out of the inherent information, using reasoning, and making it available for querying and processing through an API.

plain numeric mapping from the relational database (`id_question : 170`), a more textual description is used in 3 to create a URI specifying that this question is actually a ‘Question’ with URI: `http://idlab.technology/harmony/question_170` in the resulting ontology. Line 4 in its turn specifies, in addition to the ‘`id_subsegment`’ specification earlier, which actual low level topic this question is about, namely ‘Liquidity’, again as specified in 1. Finally, 5 adds the information into the model that the description specified is actually corresponding to the value of the ‘`hasQuestionText`’ property, with language specification ‘EN’ and ‘NL’.

4.2 Analysing the Model, Its Data and Reasoning about the Contents

Once the ontology model and its corresponding data are aligned and available in their respective repositories, optional reasoning can be triggered, using e.g. the HermiT reasoner (Glimm et al., 2014). This way, one or more inferred models are generated which can then in parallel be queried by the application through some API. Those queries are oftentimes specified using the SPARQL query language. This is presented in Figure 3.

The memory and CPU usage of this pre-processing reasoning stage have been evaluated in a test datacentre environment. The results are presented in the following subsection. It has to be noted, however, that this process is only a periodic phenomenon, i.e. when the contents of the underlying data (currently in the Harmony DB and called the A-Box) change. Additionally, when the T-Box or R-Box changes, this process will have to be executed as well. The likelihood of this happening is even less than that of a changing database content. Using this approach we can more easily guarantee a responsive deployment. One can compare the adaptation of the T-Box and R-Box with the change of a database schema and corresponding triggers.

4.3 Evaluation of the Reasoning Performance

The configuration of the machine on which the experiment was executed is: 4 64-bit Intel(R) Xeon(R) CPU E3-1220 v3 @ 3.10GHz processors, a total memory of 16262552 kB, Ubuntu 14.04.1 LTS as operating system, running the 3.13.0-33-generic kernel, Java(TM) SE Runtime Environment (build 1.8.0_191-b12) and HermiT version 1.3.8.

The materialisation process has been executed 35 times. During this execution process, the CPU usage and memory consumption was recorded using the ‘top’ application.

Analysing the results, it can be derived that on average 165.15% of CPU is used (having 400% available in this setup), and 2.74% of available memory, which corresponds to 445419.3 kB. Standard Deviation are respectively: 50.53 and 1.86. The extreme values are: (i) Minimum: 6% CPU and 0.3% Memory (45787.66 kB) and (ii) Maximum: 264.6% CPU and 8% Memory (1301004 kB). Regarding the execution time, including some household keeping code, and combining all phases together, an average time of 9926 ms, with a maximum of 10251 ms and a minimum of 9766 ms is achieved. (Standard deviation of 134.9). The main conclusion that can be drawn from these measurements is that, given the fact that this materialisation process is not to be executed in every single incoming application request, these results are acceptable and can be accommodated with a batch-processing or pre-loading style of architecture.

A final evaluation is loading the materialised model into the query engine. Once this is completed, the query engine is available for querying in an operational environment. Results are on average 14424.97 ms, with extreme values of 14080 ms and 15005 ms. (Standard deviation of 271.75). To conclude, the DL expressivity of the model is SROIF(D).

Once the above reasoning process has been completed, the resulting model – i.e. the inferred model – is persisted into a separate instance. This is called materialisation. During operation, this materialised model is read at boot-time, without additional reasoning being performed. This means that queries can be executed on a semantically enhanced model, without the need for resource intensive reasoning every time, and results in good performance in regard to response times of those queries while maintaining the added value of using description logics annotated semantic models.

5 ENSURING DATA QUALITY AND USING ONTOLOGY-BASED DATA INTERPRETATION

This section focusses on the establishment of data-interpretation mechanisms (semantic modelling) to interpret the data presented in Section 4, with the aim to generate specific investor profiles. This will be achieved through the generation of questionnaires, tailored towards the needs of the investor concerned.

5.1 Exploiting the Semantics to Contextualise the Questions in the Questionnaire

It should be clear that the strength of using ontologies and semantics is that any relations and hidden knowledge in the model can be exploited to create innovative applications. Focussing on questionnaire generation, a number of potential scenarios are: (i) Check the answers given by a person on a previous questionnaire to find the topics and products in need of further training, (ii) Check which questions are frequently wrongly answered, (iii) Check which topics and/or products/services are suitable/unsuitable for a person, (iv) Find other questions on the same topic/product/service and (v) Find information documents or training material.

One thing to keep in mind is that any of these use cases can be supported in a variety of manners and with a number of very different technologies, indeed. However, some of those technologies will require more application programming, or more sophisticated query engineering with a tight coupling of domain knowledge inside the application, rather than the model. Therefore, the approach suggested in this work is to support the right balance between generally accepted domain knowledge (thus exploiting DL-Reasoning) and proprietary application specific knowledge (potentially supported by generic Rule Engines).

5.2 ‘Knowledge’ and ‘Experience’ in Relation to ‘Person’ and ‘Investment Group’

Based on the domain knowledge of the partners involved in the project, part of the model, as illustrated in Figure 4, has been defined. ‘Knowledge’ and ‘Experience’ of financial products is defined on an individual basis.




Figure 4: Relation Person / Investment Group / Knowledge & Experience / Balance / Investment Goal.

```
Person hasFinancialKnowledge
    some FinancialKnowledge
Person hasFinancialExperience
    some FinancialExperience
```

Furthermore, an ‘Investment Group’ is composed of one or more ‘Person’ individuals. And finally, the ‘Investment Goal’ is linked to the ‘Investment Group’.

```
Group hasMember
    some Person
InvestmentGroup = Group and
    hasInvestmentGoal some InvestmentGoal
```

5.3 Value Partitioning to Classify Persons

One of the important characteristics of ontologies in general, and of OWL DL in particular is the base Open World Assumption (OWA). In short, this means that no conclusions can be drawn based on the absence of information. Basically, it is not because it is not asserted that a ‘Person’ is knowledgeable about a certain ‘Product’, given a certain ‘Topic’, that it can automatically be inferred that they are NOT knowledgeable. There could indeed be some place, repository or database not currently considered that affirms that they are not knowledgeable indeed. In order to overcome this potential issue, a standard modelling principle, named ‘Value Partitioning’ has been introduced. Compare this to the ‘Design Patterns’ in software programming (Gangemi, 2005).

The principle underlying value partitioning is rather straightforward, we state in the model that a person is either knowledgeable or not, and we affirm that those are the only two options. A covering axiom is another term referring to exactly this mechanism. This basically closes the world to the data-sets included in the run-time processing of the reasoner.

A concrete example in the scope of this project is given below. With the covering axiom, it is modelled

that a ‘Person’ is either knowledgeable about liquidity, or not, while without the covering axiom, this binary conclusion cannot be automatically drawn.

5.4 Analysing and Classifying the Person According to Their Topic Knowledge

One approach to support the automatic realisation of individuals according to the answers that were given on the questions of a particular questionnaire, and referring to the base modelling principle detailed in Subsection 5.3, is to define every category into persons that can be realised with a logical axiom, such as:

```
CAPITAL_LOSS_KNOWLEDGEABLE_PERSON =
  Person and
  (hasFilledOutQuestionnaire some
    (hasCompletedAnswer some
      ((isAnswerForQuestion some
        (hasTopic some CAPITAL_LOSS))
        and isCorrectAnswer value true))))
```

Combined with the covering axiom, this will represent in a ‘Person’ either be realised capital loss knowledgeable or not. However, this can quickly become cluttered and overly complicated. Additionally, reasoning with datatype properties is oftentimes not the most performant mechanism. So a more divide and conquer-based approach might be necessary. Thus, rather than using properties for classification, the use of fine-grained concepts with axioms is proposed. As such, easier querying is supported and it should be more performant, due to the underlying mechanisms of the tableau algorithms in the reasoner(s). As such, a new, additional, concept structure has been developed. This is presented in Figure 5.

Therefore, the corresponding query with the new structure can be formulated as:

```
Person and (hasFilledOutQuestionnaire
  some (hasCompletedAnswer some
    ((isAnswerForQuestion some
      (CORRECT_Answer
        and isAnswerForQuestion
          some CAPITAL_LOSS_Question))))
```

5.5 Automatic Questionnaire Generation

A number of typical base question categories can be defined to support the automatic questionnaire generation, and this for the three phases in the process.

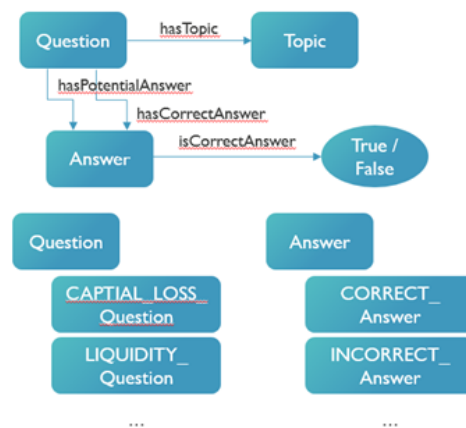


Figure 5: Additional (supporting) concepts in the model to facilitate distributed query answering and algorithmic pruning.

- *Learning*: find me all ‘CAPITAL_LOSS_ Question’ individuals
- *Testing*: find me a ‘Question’ individual for each ‘Topic’
- *ReTesting*: find me a ‘Question’ individual for every ‘INCORRECT_ Answer’

More detailed/fine-grained queries can be developed as well, e.g. (i) one could make combinations of (false) answers that lead to certain questions being asked, (ii) indicate compulsory questions, or (iii) indicate probabilities of questions or indicate a desired frequency of questions appearing in a generated questionnaire. Of course, it should be clear that potentially some service providers or banks will be accepting all, some or none of the questions present in the ontology model. Thus, a trust relationship will have to be introduced.

The above ideas are but a number of suggestions and can be further defined in further research later on. It should also be clear by now, that due to the adoption of the ontology modelling principle, additional information, stating anything about anything, can be added in a straightforward, and backwards compatible manner as long as the ontology model in itself is kept consistent.

5.6 Modelling the Structure of a Questionnaire

In Subsection 5.5, it was already indicated that a trust relationship is needed. This way banks/service providers can indicate whether questions on a certain topic for a certain product can be sourced in the collection of questions from a questionnaire from another trusted source. A key selling point of this




Figure 6: Modelling the structure of a questionnaire.

functionality is that banks and service providers can then automatically generate random or contextualised questionnaires for a given investor. However, special attention is to be given to the situation where one bank/service provider only accepts its own questions. In this situation, the questionnaire generator should generate an exact copy of the questionnaire of that particular bank/service provider. This requires the structure of a questionnaire to be modelled as well.

Having the structure of the questionnaire in the knowledge model allows a second functionality to be provided, i.e. that a given bank/service provider can accept the questionnaire structure of another service provider/bank, but can indicate that more than that bank's questions can be sourced from. This questionnaire template structure model is presented in Figure 6.

6 CONCLUSIONS

In the previous sections of this paper, we have presented our research into optimising the use of available financial consumer information in order to streamline personalised financial investment over the border of individual financial service providers. We have demonstrated that, using semantic technologies, an extendable, intelligent knowledge base can be created to support potential financial investors by lowering the administrative burden, through re-using previous information captures and to improve the investor's financial expertise in an intelligent manner.

Modelling questionnaires into the knowledge base in an innovative way by separating the structure of the questionnaire with the actual filled-out versions of the investors to be, has opened up a way of cleverly building new personalised questionnaires, either for different financial service providers or to enhance the financial literacy of the investors.

Using existing software libraries for most of the aspects of the overall application ensures a future-proof approach. Of course, the presented approach has to be approved by the necessary governing bodies in the financial sector. Although no official request has been made to the regulator, safeguarding this ac-

ceptance, was kept in mind throughout the research by the expertise of myHarmony.

ACKNOWLEDGEMENTS

This research or part of this research is conducted within the project entitled Harmony co-funded by VLAIO (Flanders Innovation & Entrepreneurship) and myHarmony.eu.

REFERENCES

- Abi-Lahoud, E., Raoul, P.-E., and Muckley, C. (2017). Linking data to understand the fintech ecosystem. In *SEMANTICS Posters&Demos*.
- Bennett, M. (2013). The financial industry business ontology: Best practice for big data. In *Journal of Banking Regulation*, volume 14, pages 255–268.
- Calvanese, D., Cogrel, B., Komla-Ebri, S., Kontchakov, R., Lanti, D., Rezk, M., Rodriguez-Muro, M., and Xiao, G. (2017). Ontop: Answering sparql queries over relational databases. In *Semantic Web*, volume 8, pages 471–487.
- Dietz, M., Khanna, S., Olanrewaju, T., and Rajgopal, K. (2016). Cutting through the fintech noise: Markers of success, imperatives for banks.
- Dimou, A., Vander Sande, M., Colpaert, P., Verborgh, R., Mannens, E., and Van de Walle, R. (2014). Rml: a generic language for integrated rdf mappings of heterogeneous data. In *7th Workshop on Linked Data on the Web, Proceedings*, page 5.
- European Parliament, C. o. t. E. U. (2014a). Directive 2014/65/eu of the european parliament and of the council of 15 may 2014 on markets in financial instruments and amending directive 2002/92/ec and directive 2011/61/eu.
- European Parliament, C. o. t. E. U. (2014b). Regulation (eu) no 600/2014 of the european parliament and of the council of 15 may 2014 on markets in financial instruments and amending regulation (eu) no 648/2012 text with eea relevance.
- Gangemi, A. (2005). Ontology design patterns for semantic web content. In Gil, Y., Motta, E., Benjamins, V. R., and Musen, M. A., editors, *Proceedings of the 4th International Semantic Web Conference (ISWC 2005)*, volume 3729 of *Lecture Notes in Computer Science*, pages 262–276, Galway, Ireland. Springer. doi: 10.1007/11574620_21.
- Glimm, B., Horrocks, I., Motik, B., Stoilos, G., and Wang, Z. (October 2014). Hermit: An owl 2 reasoner. In *J. Autom. Reason.*, volume 53, page 245–269, Berlin, Heidelberg. Springer-Verlag.
- myHarmony (2020). <https://www.myharmony.eu/>.
- Ontotext (2020). <https://www.ontotext.com/>.