SWIFT INSTITUTE

SWIFT INSTITUTE WORKING PAPER NO. 2013-006

A CRITICAL AND EMPIRICAL EXAMINATION OF CURRENTLY-USED FINANCIAL DATA COLLECTION PROCESSES AND STANDARDS

DR. SUZANNE MORSFIELD

DR. STEVE YANG

SUSAN YOUNT

PUBLICATION DATE: 2 NOVEMBER 2016

 The authors would like to convey special thanks to Dr. James Hendler, and Michael Atkin for contributing on-the-record insights reflected within this study. We also would like to express our appreciation to those who participated anonymously as part of our interviews and surveys. This study would not have been possible without this support. All errors are our own.

Abstract

This paper provides a framework for analysing and assessing financial data standards. We test the proposed framework by examining common financial data standards or contexts in use today—FIX, FpML, ISO 15022/20022, XBRL, and the US CFTC's Swaps Data Repositories' data. The empirical information for each standard was gathered by conducting what we believe to be the first 360-degree survey of all the stakeholders to each standard—ie their respective designers, regulators, data input personnel, and data end users. This approach led us to the discovery of a new construct for evaluating existing and making recommendations for future financial standards; we characterise this construct as stakeholder harmony vs stakeholder dissonance. Our findings indicate that extreme instances of dissonance exist not only between types of stakeholders, but among the same category of stakeholder. From these survey data, we also designed a new empirical score of standard quality, and note that the ISO standards outperform the other standards, per our measure, with XBRL achieving the next highest score. We conclude with detailed list of practical recommendations which are organised to reflect the structure of the framework and stakeholder survey; they build on the theme of identifying and preventing stakeholder dissonance. Some key recommendations include: keep a full representation of all types of stakeholders at the design and implementation table; start small and with actual use cases provided by the end users of the data; use formal design and maintenance tools wherever possible; make the technical specification as invisible as possible; and listen to contrarian voices, especially if coming from the end user community.

1 Introduction

Data standards, and standards in general are an often under-appreciated, if not misunderstood element of everyday life. There are thousands of them in action at any given moment around the world. From electrical sockets to DVD formats to the world's currencies, they are what make our lives work. However, they can confound the simplest task like turning on a lamp if the plug doesn't match the socket. When one begins to examine this topic, one quickly realizes that not all standards are created equal. That is, not all of them serve their initial or stated purposes or their end users equally well. Why is that? What differentiates efficacy among and between standards, and how could one assess a standard?

The goal of this paper is to provide a framework for such an assessment within one subsector of data standards—i.e., within the world of financial data standards. We have chosen five financial data standards to examine in detail within this framework—FIX, FpML, ISO 15022/20022, XBRL, and the U.S. CFTC's Swaps Data Repositories (SDRs). Finally, we gather and analyze empirical data by surveying standards designers, data input personnel, and data end users.

These standards and settings are of interest because they cover a lot of ground, both technically and practically. For example, they range from being:

- Completely open source to completely proprietary with use fees
- Focused on the granular transaction to focused on the aggregate final regulatory reports
- The only game in town to being only one of several possible solutions
- Primarily a business and data model that could be indifferent to the communication language employed to primarily a communication language with little to no data model
- Highly rigid with respect to the execution to a loosely-defined application

Data standards are created and exist in an environment of policy, technology, and business-related decisions.

Where is the London Whale?

"Unfortunately, I must report that the Commission's progress in understanding and utilizing the data in its current form and with its current technology is not going well.

Specifically, the data submitted to SDRs and, in turn, to the Commission is not usable in its current form. The problem is so bad that staff have indicated that they currently cannot find the London Whale in the current data files. Why is that?

In a rush to promulgate the reporting rules, the Commission failed to specify the data format reporting parties must use when sending their swaps to SDRs. In other words, the Commission told the industry what information to report, but didn't specify which language to use. This has become a serious problem. As it turned out, each reporting party has its own internal nomenclature that is used to compile its swap data.

The end result is that even when market participants submit the correct data to SDRs, the language received from each reporting party is different. In addition, data is being recorded inconsistently from one dealer to another. It means that for each category of swap identified by the 70+ reporting swap dealers, those swaps will be reported in 70+ different data formats because each swap dealer has its own proprietary data format it uses in its internal systems. Now multiply that number by the number of different fields the rules require market participants to report."

Scott O'Malia, Former Commissioner, U.S. Commodities and Futures Trading Commission in the Keynote Address to SIFMA Compliance and Legal Society Annual Seminar, March 19, 2013.

Our focus in this paper is on the business-level decisions and issues, and we examine the following:

- The standard's design, purpose, and evolution
- Filers'/Preparers' ability to create high quality data using the standard
- Users' ability to use the data
- Regulators' ability to use the data to monitor and oversee their respective markets.

The impetus for our study was a call for research by the SWIFT Institute that highlighted statements made by Scott O'Malia, formerly of the U.S. Commodities and Futures Trade Commission (CFTC)[1] (see textbox) about the CFTC's inability to find the infamous London Whale trades.[2] The CFTC had been collecting a huge volume of swaps trading data under Dodd-Frank rules which were designed to help regulators understand and limit systemic risks. The role of "financial data standards" as the cure for data that cannot be deciphered was championed by Former Commissioner O'Malia, while others pointed to perhaps subtler solutions like data quality audits and semantics clarifications. The SWIFT Institute responded by asking for academic research that asks the more general question: what is the future of financial data standards?

Hence, our study seeks to answer questions about how current standards are working, and what the lessons learned are from these specific implementations. Other important questions include: when and where exactly is a financial data standard necessary and when might data science expertise be an adequate, and sometimes better solution? While important, these are

[1]Mr. O'Malia is currently the CEO of the International Swaps and Derivatives Association

[2] The London Whale trades refer to an aggressive trading strategy in credit default swaps, which ultimately resulted in over $6 billion in losses. [source: Wikipedia and Bloomberg(for loss amount)]

beyond our current scope, although they should be addressed by future research in this area.

This paper contributes to the current discussion in three ways:

1. We propose the first framework (to our knowledge) for assessing financial data standards; this framework also can serve as a comprehensive set of considerations for designing new standards as well;
2. We conduct the first 360-degree survey (to our knowledge) of the complete set of stakeholders for each standard we examine; this included in-depth interviews with key designers and standardisers of standards;
3. We provide the first empirical model (to our knowledge) for evaluating and comparing the quality of financial data standards.

By conducting all three steps in this study, we hope to have identified areas for regulators and standardisers in particular as best practices or as areas for improvement when developing and implementing standards.

2 Background

2.1 Types of data standards

Financial data standards can be organised as follows:

1. Whether they are a horizontal vs. a vertical standard, or
2. What exactly are they standardizing.

McKenna, Nichols, and Northey (2014)[3] define a horizontal standard as one that potentially works across all industries and sectors. A common example of this is HTML, used to create websites regardless of the underlying topic, purpose, or industry. A vertical standard, on the other hand, is generally designed for a specific setting within a specific industry or setting, and thus doesn't have a broad use or appeal outside that arena. Applying this categorisation, we would argue that all of the standards we are studying are vertical standards. That is, their application and effectiveness is limited to some aspect of business reporting or messaging, and in some cases that application is only to a very narrow business transaction, such as an individual trade.

In our interviews with Michael Atkin, Managing Director of the Enterprise Data Management (EDM) Council, he lays out the following three categories for what a particular financial data standard may be attempting to standardise, along with examples of each:

1. Identification standards - unique and precise identification of the 'thing':
    a. 'Entity' standards – to identify the financial entity and unravel relationships to other entities (example: LEI);
    b. 'Instrument' standards – to uniquely identify the financial instrument (examples: CUSIP, ISIN, SEDOL, FIGI);

2. Messaging standards –efficiently communicate from one point to another (examples: XML, JSONN-LD). XBRL (financial reporting), MISMO

[3] K. McKenna, J. Northey, and B. Nichols. *Handbook of Financial Data and Risk Information, Volume II, Software and Data*; Cambridge University Press, United Kingdom; Chapter 19. Financial data interchange standards.

(loans), FpML (derivatives), FIXZ (pre-trade), ISO 20022 (post-trade) are all types XML schemas. Complexity comes in because these XML standards also contain syntax, dictionaries, and schemas, all of which are (were) needed to communicate with precision.

3. Content standards –precisely define terms in the context of how they operate in the real world. Think of this as "what do I mean by the word 'balance' or the word 'duration' or the concept of 'accrued interest'?" Here there are both data elements (interest rates, dates, schedules, etc.) and rules on how they are derived (balance can mean "net balance" or "gross balance" or "principle only balance" (example: Financial Industry Business Ontology/FIBO).

2.2 Selection of standards for study

The standards we examine are messaging standards, because they are at the heart of the concerns raised by Former Commissioner O'Malia. In particular we focus on: FIX, FpML, ISO 15022/20022, and XBRL (and SDR, where possible). We study these particular messaging standards because:

- They are currently in widespread (or growing) global use, and are the financial data standards most in play at present; and
- They were constructed based on the same underlying communication language—i.e., XML (or, eXtensible Machine Language), with the exception of the SDR setting (for which no technical specification has been required at this point, see Table 1 below).

Data standards consist of:

1. A technical specification,
2. The semantics (ie taxonomies and relationships), and
3. Their implementation and maintenance.

In Table 1 below, we lay out the differences and similarities across those three categories for each of the five standards. The key similarity is that, as noted above, all these standards are primarily XML-based specifications. The primary differences lie in what the standards are used for and by whom. For example, XBRL is used mainly in a financial or tax reporting context, while FIX and FpML both support trading information. The ISO standards are

unique in that they serve a variety of financial services processes, from payments and trades, to foreign exchange transactions.[4]

Table 1. The Standards

The Standards or Data Settings	Specification	Semantics	Implementations
ISO15022/20022 (Umbrella for payments, securities, trade services, cards, Fx)	XML	Consortium-defined data dictionaries & relationships	Inter-institution, regulatory
FIX (Trade capture, post-trade, STP)	XML	Consortium-defined data dictionaries & relationships	Inter-institution, regulatory
FpML (Trade capture, post-trade, STP)	XML	Consortium-defined data dictionaries & relationships	Inter-institution, regulatory
XBRL (Financial reports, FDIC Call reports, Mutual Fund Risk/Return summaries, Tax filings, Corporate Actions, Mutual Funds)	XML	Financial reports, call reports, mutual fund reports, and tax filings: Regulator-defined data dictionaries and relationships Corporate Actions: Consortium-defined data dictionaries and relationships	*U.S. SEC financial reports in annual and quarterly reports, Mutual fund risk/return summaries *U.S. FDIC "Call reports" *Various global tax and financial authorities *Corporate Actions (Global)
SDRs/TRs (Swaps, cleared and uncleared)	Not imposed by regulator	*250 column spreadsheet of data fields required by regulator *Little to no data definitions or relationships structure was specified	U.S. CFTC

[4] For additional details on the standards, see McKenna, Nichols, and Northey (2014) as referenced above, and Houstoun, Milne and Parboteeah (2015) *Preliminary Report on Standards in Global Financial Markets*; Working Paper, SWIFT Institute.

[RM]: "[Relational databases] could be the biggest impediment to progress. It's not that we don't need them, it's just not where the real world Big Data problems are coming from. We know how to do that and it's all in the past. Today's data doesn't fit into relational databases."

[VG]: Does it make sense to structure unstructured data (using tags, NLP, taxonomies)?

[RM]: "I do not think so. I think you should avoid doing violence to the data. Companies need to find ways to deal with the data as-is. They can do this by seeking out partners and vendors that specialize in analyzing these data types. We spend a lot of time forcing old methods on to new things, but that's not useful with data. You need to find new ways to manage the variety of today's data."

Dr. Roy Marsten, in an interview with Vincent Granville, on Data Science Central, 14 February 2014.

The standards we examine were created, for the most part, at a time when XML was the future of data standards, and as such rely on XML as the framework for communication. People we spoke with were divided on whether XML was here to stay and still useful, or whether it was overdue to be replaced by either newer specifications such as RDF/OWL or JSON, or better ways of mining existing information. At the extreme end of the spectrum, some of the proponents of other standards felt it was a waste of time to examine XML-based standards in such detail. Since we are studying what is, with an eye to what also can or should be, a key question we asked all survey participants was whether the specification/technology in use by a particular standard was, in their opinion, state of the art. If they thought it was not, we asked them to tell us what they thought was, or should be. We asked the stakeholders in the financial data markets to reflect on the here and now, and to provide input on the future.

The main future-oriented input from stakeholders was as follows:

- Ensure that taxonomies/ontologies are not tied to only one particular technical specification so that they can be 'future proofed'—the only specific alternatives mentioned were JSON and RDF/OWL;
- Move forward from a relational database only view of the data being standardised and collected (see textbox); and
- Develop something like the Legal Entity Identifiers (LEIs) for ease of correct and consistent identification of data with its appropriate legal entity.

3 Data standards development and analysis model

To develop and empirically test a framework to assess the quality of a data standard that would work across a variety of standards, we specified the following three key areas to study:

1. Core Elements—those areas that are likely to be present in any data standardisation or collection setting;
2. Big Data Issues—those issues that arise because of the vast quantities of data now available for analysis or that are being collected in many settings; and
3. 360-degree Stakeholder View—every data setting has stakeholders whose data and analysis needs will most likely not be the same, and possibly even incompatible with each other.

3.1 Core elements of the model

The core elements consist of the areas that are likely to be present in any data standardisation or collection setting. We define them in Figure 1 below.

Figure 1 – Core elements and their definitions

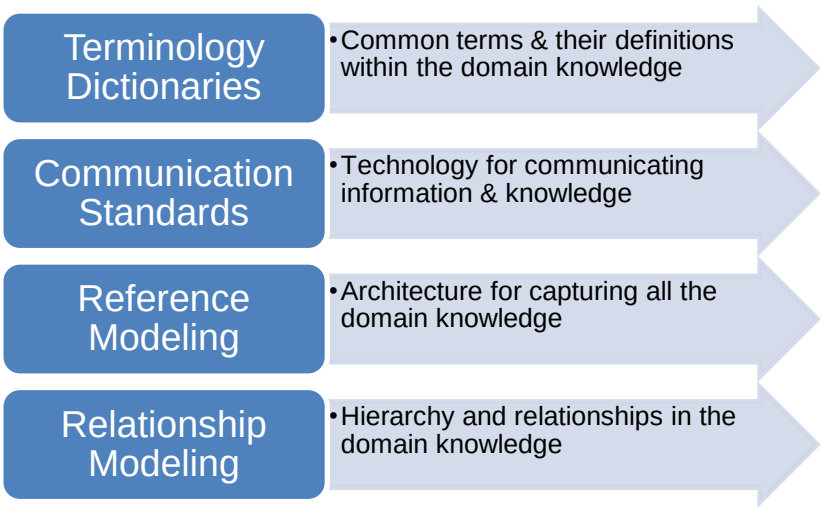

3.2 Big data elements of the model

We also draw attention to three broad issues relevant to considering the use of a data standard in a big data setting:

1. How precisely can or should one leverage the given standard in such a context?

2. What are the common obstacles to implementing the standard in a big data arena?
3. How should ongoing development and evolution of a given standard be provided for within a big data setting?

3.3 Stakeholder elements of the model

Finally, with respect to the development of a common model, we introduce a 360-degree stakeholder view. We believe this is both a key contribution of our study to the discussion of the future of financial data standards, and an important component of any successful data standard design and implementation. Our hypothesis—to be tested by gathering survey data from all the stakeholders to each data standard—is that those implementations will be most successful that have full and engagement of all stakeholders across the entire lifecycle of the standard. We interviewed or surveyed:

1. Standardisers and designers (and subject matter experts, or SMEs)—a unique contribution of our study is that we interviewed the original and/or key creators/designers of the standards at length;
2. Regulators (can also be standardisers or users);
3. Preparers (those who put data into the standard); and
4. Users (those who consume the data for any purpose).

3.4 Putting the elements together—Our proposed analysis framework

Figure 2 illustrates our view is that it is important to address all three areas to design a meaningful, robust common framework for designing or assessing data standards.

Figure 2 – Proposed model for developing and analysing data standards

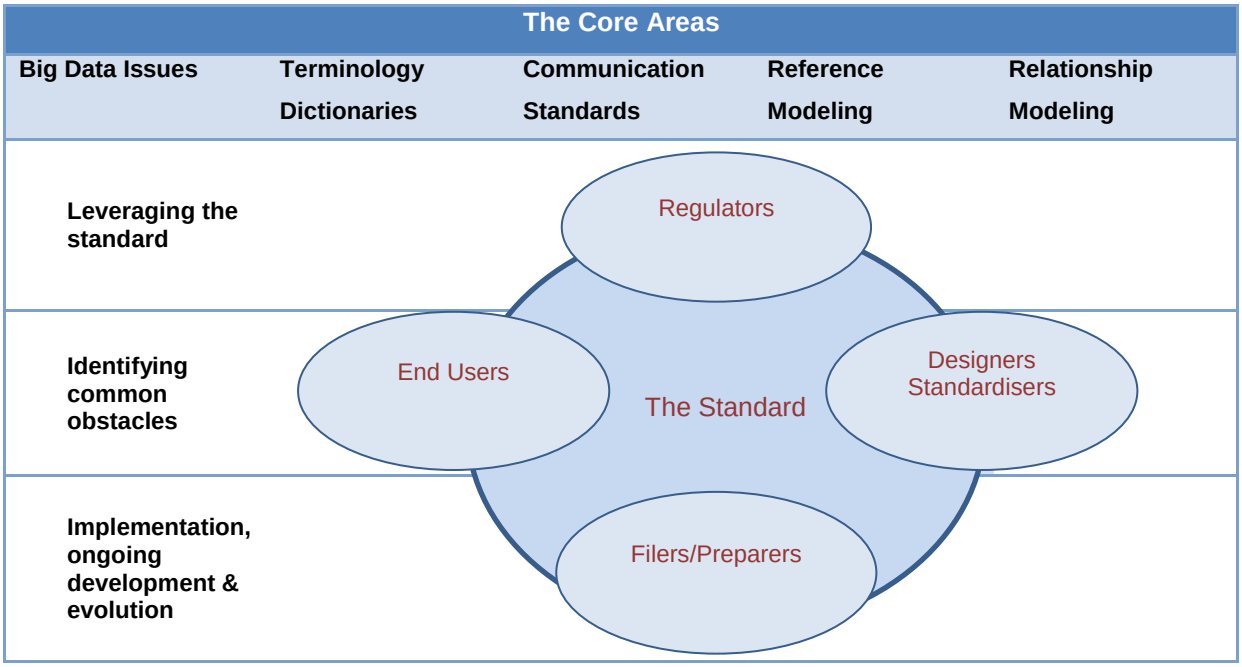

We identify sources of common issues across each of the core elements (see the following Table 2). A standard must have common terms and definitions to be usable in general, but especially in a big data setting. The technical design and technology by which the standard's communication is conducted must also be operational beyond an Excel spreadsheet or small set of financial messages. Reference modeling requirements should provide linkage from or enable feedback to outside resources or authoritative literature or legislative rulings. Finally, how or if the ontology or taxonomy is designed to function in a big data setting may make or break the usability of a given standard.

Table 2. Big Data Issues across the Core Elements

	The Core Areas			
Big Data Issues	**Terminology & Dictionaries**	**Communication Standards**	**Reference Modeling**	**Relationship Modeling**
Leveraging the standard	Common terms Definitions	Technical Design Technology	Application Presentation Session Transport Network Datalink Physical layer	Ontology Taxonomy
Identifying common obstacles	Disagreement Understandability Granularity	Costs Integration Usable	Costs Integration Usable	Complexity Accessible Useful
Implementing ongoing development & evolution	Widely used Updates Governance Feedback loop	Widely used Updates Obsolescence Governance Feedback loop	Widely used Updates Governance Feedback loop	Widely used Updates Governance Feedback loop

The common obstacles to creating a working dictionary revolve around overcoming disagreements about definitions and uses of the terms, ensuring that those same definitions are understandable to all stakeholders, and deciding on the level of granularity at which to identify the needed data items. For both communication standards and reference modeling the obstacles are similar—dealing with the monetary and human resource costs to learn, implement, and integrate both elements, as well as ensuring that each of them are truly usable for all stakeholders and not just the most technically proficient.

The final big data issue is how best to implement a program for the ongoing development and evolution of a given standard. The obstacles across all four elements are the same for this issue—i.e., to make sure that the standard is widely used, is carefully updated, is protected against obsolescence (or that there is a seamless transfer to something better), has a robust and effective governance mechanism, and has a healthy feedback loop for those who use the

standard as either a filer/"inputter" or an end user of its data (whether a regulator or a market participant).

4 Technical design features of each standard—desk research

The primary role of our desk research was to gather information about each standard or setting from publicly available sources such as the standard's website or other promotional/educational materials. In particular, we used this method to develop a high-level view of some of the technical design features relevant to our common model template.

From the summary in Table 3 below, a few things can be observed. All of the settings for which a formal standard is operational have formal dictionaries, and reference and relationship modeling to some degree, depending on the setting in which the standard is deployed. These settings vary across and within each other with respect to the exact nature of these formalities, but they exist, and appear to be quite strong in most cases. In contrast, the SDR setting permits filers to choose their own data standard and communication format. This may or may not be bad, but what immediately draws our attention is that common terminologies and reference or relationship modeling appears to not exist within this data collection process. This observation is made solely from our desk research.

Table 3. Technical Design Features by Standard

The Core Areas				
Big Data Issues	**Terminology & Dictionaries**	**Communication Standards**	**Reference Modeling**	**Relationship Modeling**
ISO20022 (Umbrella for payments, securities, trade services, cards, FX)	Yes	XML ASN.1	Yes Varies by setting	Yes Varies by setting
FIX (Trade capture, post-trade, STP)	Yes	XML ASN.1	Yes Varies by setting	Yes Varies by setting
XBRL (Financial reports, US FDIC Call reports, Corp Actions, Mutual Funds, GRC metrics)	Yes	XML	Varies by setting	Yes Varies by setting
SDRs/TRs (Swaps, cleared and uncleared)	No	None	No	No

5 360-degree stakeholder view—interviews and survey findings

The survey instrument was organised as follows:

First, we examined what we have identified as the four core elements relevant to a standard's design and implementation:

- Terminology and dictionary;
- Communication standard;
- Reference Model; and
- Relationship modelling.

Next, we examined the four core elements over the life cycle of a standard and across three aspects of what are considered "big data" methods:

- How has the development and maintenance (including updating for technological advancements) of the standard's technology, content, organisation, and governance of the data collection/data standard made the financial data more useful than prior to this standard's existence?
- How is the standard or data collection process actually leveraged when using various data analysis technologies such as the cloud, big data, etc.? and
- What are the common obstacles for the financial industry to use big data technology with respect to this standard or data collection process/setting?

For each stakeholder group we report on the more notable results in the sections below.

5.1 Standards designers/standardisers' views

A critical step in our project was to interview some of the key standards designers or standardisers for each of the settings. We began our project here in order to best understand the standard itself as envisioned by the designers; this also sets the stage for a 360 degree stakeholder view.

Table 4. Distribution of Standard Designer Responses

	A). XBRL	B). FIX	C). FpML	D). ISO15022	E). ISO20022	F). SDRs	Total
Responses	6	6	4	6	38	1	61
Responses (%)	9.84%	9.84%	6.56%	9.84%	62.30%	1.64%	100%

Designers tended to have a variety of skills and prior experience, for example:

- XBRL and SDR designers did not have data or information technology experience;
- The design teams for all standards had a significant proportion of its membership with 5+ years' experience in the general areas of financial data standards or filings; and
- The ISO20022 design team is the most experienced with 100% of it 38 responders reporting 10+ years.

5.1.1 Initial development and maintenance of the standard

Terminology and dictionaries

- Most standardisers had formal policies for creation and maintenance of the standard;
- In some instances, strong dictionaries already existed, driven by legal or regulatory frameworks;
- Major issues were encountered by all standardisers, except for ISO20022 and a one specific implementation of XBRL; some noted that when stakeholders disagreed about dictionaries a single stakeholder group tended to win out, sometimes to the detriment of others; and
- Tensions existed between those who envisioned using the standard to collect granular data vs. collecting information at a more commonly presented or reported, summarised level of granularity.

Communications standards

- Most had formal policies for creation, selection and maintenance with regards to the communication standard;
- Generally, the technology chosen was considered the best available at that time; and
- Most standards had a formal selection process; XBRL did not.

Reference modeling

- Standardisers were divided on whether formal policies and techniques were used during initial development. However, standardisers responded that formal approaches now exist for all standards; and
- Weak availability of modeling tools in most cases.

Relationship modeling

- Responses were divided on whether formal or informal relationship modeling policies existed during initial development;
- If modeling methods were used, they were often used only informally;
- Formal software tools for relationship modeling were not used except for FIX and FpML; and
- Formal ontology methods were considered by a minority of the designers for XBRL, FIX, and ISO.
- Across the board, standardisers expressed a strong recommendation that formal methods should be used now.

5.1.2 Leveraging the standard

Terminology and dictionaries

- XBRL and FpML standardisers believed users would utilise the dictionaries; FIX did not believe they would be used much, ISO standardisers had divided responses on whether they would be used;
- There were significant gaps reported between the terms and definitions used by users and preparers, although this was less pronounced for XBRL;
- Standardisers believed the dictionaries would be utilised often by users, except for FIX; and
- Standardisers believed it would require low effort by users to utilise the dictionaries.

Communication standards

- XBRL and FpML standardisers believed users would utilise the communication standard; FIX did not believe they would be used much, ISO standardisers had divided responses on whether they would be used; and
- Standardisers were divided on the level of effort that would be required by users to utilise the communication standard

Reference modeling

- XBRL and FpML standardisers believed users would utilise the reference modeling; FIX did not believe it would be used much, ISO standardisers had divided responses on whether it would be used;
- Standardisers were divided on how similar the design reference model was to the existing user reference model;
- Standardisers were divided on how similar the design reference model was to the existing preparer model;
- Standardisers believed the reference model would be used often by users, except for FIX; and
- Standardisers believed it would require low effort by users to utilise the reference modeling.

Relationship modeling

- XBRL and FpML assumed users would utilise the relationship modeling; FIX assumed users would not use it much; ISO results were mixed;
- There were mixed responses on whether all stakeholders had a common understanding of the ontology that was implemented; and
- All standardisers believed relationship modeling would require high effort by users to utilise

In a later section of the paper, we look at this again from the users' point of view to see how this assumption has worked out in practice.

5.1.3 Common obstacles

5.1.3.1 Using the data

In general the designers believed that users of their standard would find it easy to bridge whatever differences might exist between the users' own processes and the standard's requirements. Notable exceptions are XBRL and ISO20022:

- For ISO20022, we received explanations that many users of the standard are also interacting with internal legacy systems that are based on quite old or proprietary standards that, once mapped, would function well;
- Most XBRL designers thought the bridge would be difficult, but did not elaborate;
- In the case of at least two standards, some of the designers noted that end users' needs were not directly taken into account up front:
    - The solution was designed and implemented, with high level or no direct knowledge of impacts on the users of the standard or its data;
    - In these instances, some designers further report that there is an ongoing attempt to "work backwards" to attempt to solve users' needs, and in some cases, to discover and formally define users' business requirements;
- An interesting observation is that, except in the case of FIX, the designers assumed that data users would utilise or exploit the value of the components of the standards; and
- XBRL and FpML designers were particularly optimistic about this.

5.1.3.2 Dealing with complexity

Before discussing questions about infrastructure, we asked standardisers a broader question about how easy or not they thought it would be for users and preparers to bridge any gaps between the requirements of the standards and that of their own internal processes. In the same vein, we asked whether the designers expected that the users would need a separate technology just to consume the data;

- XBRL was viewed as generally the most difficult to bridge, and the one that uniformly was expected to require separate software to be able to use the data;
- FIX was envisioned as the easiest to bridge and the least likely to require separate technological consumption solutions;

We also asked for designers' perceptions about the actual effort-levels that designers expected data users to have to expend to use the output/data captured by the standards:

- The common terms are generally low on the effort scale;
- The ontology is high on the effort scale; and
- The metadata and reference model results are more mixed across the standards.

We then come at this issue from another angle. In a pair of questions, we asked whether or not internal frameworks for dealing with data complexity are either assumed to be in place already, or are even deemed necessary on the part of those who would want to consume the data:

- XBRL and FpML standardisers assumed quite strongly that these frameworks would already be present. This is important to note, because both of these standards also appear to require that they are in place in order to consume the data: and
- Conversely, FIX designers neither assumed nor does the standard appear to require such frameworks;
- The ISO results are mixed.

5.1.3.3 Errors

In this section we discuss the standardisers' views on errors in the data captured by the standard (other potential obstacles to using data standards are discussed in later sections of the paper).

For purposes of this study, we focused on the following error types:

- Data inconsistency
- Missing Data
- Incorrect Data.

We also allowed responders to add any other error types through the use of a text box in the survey.

5.1.3.3.1 Effect of errors on a data user's ability to utilise the data

- Only the XBRL standardisers appear strongly split on issue of whether errors would impact the usability of the data, whereas the other standards' designers tend to agree that there would be an effect.

5.1.3.3.2 Common errors

- All standards designers report that they have formal methods in place to prevent the error types noted above. In one instance, an interviewee provided examples of some of the specific methods that their standard employs to prevent errors: no open text fields, and no fields that are redundant or too similar to another field. For standards that are transactions-based, automated validations ensure that the transaction won't be completed if certain errors exist.
  - FIX and FpML appear to have the same frequency of errors across all types;
  - For ISO, the most frequent errors appear to be incorrect data; and
  - For XBRL, incorrect and inconsistent data are the most frequent types of errors.

5.1.3.3.3 Severity of errors

- XBRL and FpML indicate that the errors in their data are the most serious relative to the other standards; and
- ISO reported the least serious errors.

5.1.3.3.4 Reasons for and Prevention of Errors

We provided an open-text box to learn more from designers about to what they attributed the amount and types of errors. The main responses were:

- Lack of robust validity checks built into vendor software at the input stage;
- Human error that cannot be prevented by software;
- Lack of use of the validation tools available (some of which are free to use, but are not used); and
- Bad or inadequate testing of data that are input.
- Across all standards, the designers reported that they had formal methods in place to prevent these common errors.

5.1.4 Some additional life cycle issues

5.1.4.1 Initial vision and development

In most cases, we note reasonable consensus among the designers within each standard on the initial vision for the standard. However, there were a few questions for which there was some surprising variation in responses:

5.1.4.1.1 For and by whom was the standard designed

By whom was the standard constructed—professional standardisers, preparer consortium, user consortium, or regulators:

- XBRL and ISO20022 had noticeable variation in the perceptions of who was do the designing;
- XBRL was viewed as mostly developed by regulators, and ISO20022 was viewed as mostly developed by professional standardisers; while
- FpML and SDR designers were unanimous in the view that it was developed by a preparer consortium.

For whom was the standard constructed

- XBRL and FpML both had noticeable variation in the perceptions of who the standard primarily was intended to serve;
- XBRL was split between regulators, preparers and users, but with the majority of responses for regulatory use; and
- FpML was split between regulatory and non-regulatory users.

5.1.4.1.2 Formality of design methods

- In general, the various standards designers had formal or informal methods for choosing or creating the various aspects of standard;
- The most commonly formalised area was that of dictionary development, whereas the least formalised was relationship modeling;
- An area worth noting that seemed to emerge more from the interviews than the survey responses had to do with how much exposure to the complexities of the standard the preparers or end users were confronted with during the initial development and implementation:
    - One designer described their setting as follows, “we designed the implementation process so that ALL the difficult aspects of the implementation and dealing with

[the standard] were completely hidden from the [data preparers] and the end users." This is the observation from an implementation of one of the Standards that is considered by many external impartial observers as very successful.

5.1.4.1.3 Governance

- Virtually all the standards designers considered governance issues from the beginning, and had formal processes in place for each of the big data issues;
- Some of the more interesting observations about the initial governance issues are:
  - Conflicting demands among the various stakeholders—e.g., regional differences, trade-offs between stability vs. improvements, semantic completeness vs. performance, and even what "data" really is; in some cases, certain stakeholders seemed to be able to monopolise the outcomes;
  - Various complexities related to being volunteer-driven organisations and efforts, especially initially; for example, time constraints, absence of deep enough technical skills; and
  - Global coordination across multinational constituents and governing bodies, and varied levels of understanding of the standard itself; some at the table did not understand the standard, but had input into it:
    - One standardiser noted that formal processes and governance mechanisms become even more important and relevant in light of the extensive and multi-faceted coordination needed for most of these standards.

5.1.4.1.4 Is it state-of-the-art technology?

- Standardisers and designers generally viewed their respective standard as state-of-the-art across the board; and
- We will learn later, however, that this perspective is often in contrast with that of the standard's end users or preparer/filers.
  - This disparity may be the result of a time lapse between the design and the implementation, which means that, at the time of design, the standard may have been using cutting edge technology.

5.2 360-degree stakeholder view across all standards

We believe one of the important contributions of this study is that we document how all the stakeholders to each standard view the same topics in our framework for analyzing a data standard. We can then compare the responses to assess the standard within the environment in which it operates. We divide our discussion of the findings into two parts:

- Descriptive

  Below we provide an overview of the notable aspects of the survey responses organised along the key dimensions of our model. We observe that both across and within standards there are varying degrees of harmony and dissonance among the stakeholders (which is discussed in more detail in the next section of the paper); and

- Analysis

  In the next section of the paper we discuss our anecdotal observations of the survey findings and provide results of our empirical examination of the survey responses. By combining statistical analysis with our analytical model, we believe we are able to provide a framework for objectively measuring the quality of each standard that is useful for future research as well.

The 360-degree survey questions covered similar ground to those poised to the standards designers during their interviews. However, some additional aspects were queried, and the structure of the survey was slightly different. The outline of the survey questions is noted in Figure 3 below.

Figure 3 – Outline of 360-degree Survey Questions

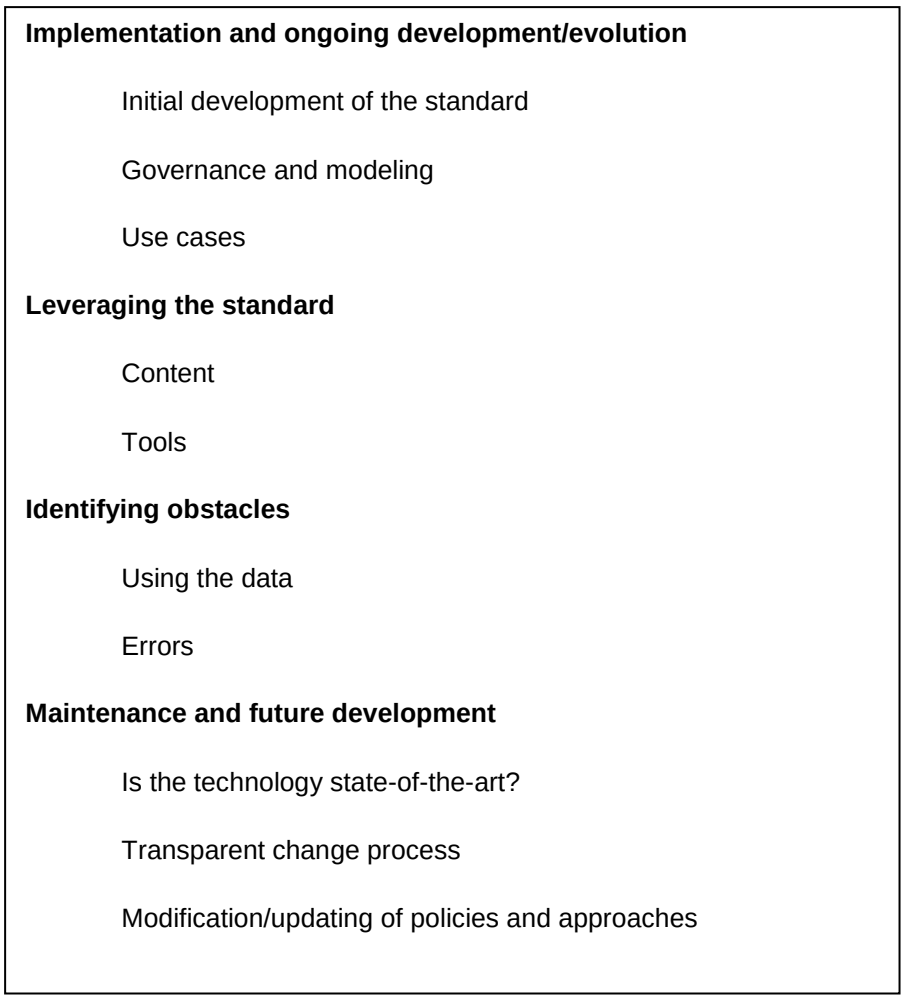

5.2.1 Implementation and ongoing development/evolution

5.2.1.1 Initial development of the standard

- There was little consensus among the stakeholders (and even among the same types of stakeholders in some cases) about by and for whom the standard was designed.
  - Except for FIX: strong agreement among and between all the stakeholders that data users constructed the standard;
- In one extreme case (FpML), the average data user stated they did not know for whom the standard was created; and

- In another case (XBRL SEC implementation[5]), two groups constructed the taxonomy, with two different end users in mind. The financial statement portion of the taxonomy was developed from an investor model, while the footnote disclosures were developed by another group based on accounting standards, SEC regulations, and common reporting practice.

5.2.1.2 Governance and modeling

- Formal governance processes are reported to be in place for virtually all aspects of the initial design of the standard;
- However, the stakeholders are not in agreement about which components of the standard are subject to modeling, either formally or informally; and
- Formal ontological methods were employed according to the designers, but the other stakeholders did not necessarily support that view in all cases.

5.2.1.3 Use cases

- Both the standardisers and the preparers believed formal use cases (reflecting end users' needs) to have been employed during the design process;
- Most reported having pre-defined, specific use cases; and
- Data users conveyed that such use cases were not employed during the initial design.

5.2.2 Leveraging the standard

5.2.2.1 Content

- Both the filers and the end users reported that the content as well as the technical aspects of the standard met their needs;
- XBRL:
  - Users were split on whether they believed the content was sufficient;
  - Those who did not think it was sufficient cited disclosures outside of regulatory requirements for applying the standard – standard entity identifiers, and other information filed with the SEC; and
  - Users also expressed a desire for more interoperability between taxonomies that cover identical information.

5.2.2.2 Tools

- The perception of availability and cost of tools (whether in effort or monetarily) is quite mixed for this standard;

---

[5] XBRL was also implemented by the US Federal Deposit Insurance Corporation's (FDIC) for Call Reports in 2005. If relevant, we indicate which implementation—ie that of the US SEC or that of the US FDIC.

- ISO 20022:
  - The availability of tools appears to be strong, and the responses consistent across all respondent groups; and
  - The tools are generally reported as effective, except for those related to the ontology, for which a Neutral (or middle of the road, neither effective nor ineffective) rating is found

5.2.3 Identifying obstacles

5.2.3.1 Using the data

- ISO 15022: The designers overestimated how difficult it would be for preparers to understand the common terms, and underestimated how difficult it would be for both groups to deal with the ontology;
- ISO 20022: The designers overestimated how difficult it would be for preparers to deal with the reference model and the ontology, and underestimated how difficult it would be for both groups to understand the common terms;
- FIX: The designers were generally accurate that grappling with the ontology, metadata, and reference model would be a challenge, and that understanding the common terms would require little effort. However, data users reported a high level of effort required to understanding the common terms;
- FpML: There is little consensus among the stakeholders about the effort level to understand the various components of the standard; and
- XBRL: The majority of users report that they find it easy to understand and implement updates to taxonomies, technical design, reference models and ontology models.[6] Users were divided on the effort and man-hours needed to understand the data and make the data usable. Standardisers believed that it would not take users much effort to understand these four areas, and agreed that they would be used frequently.

5.2.3.2 Errors

- Generally, there was little agreement between and among stakeholders on the amount, types, and severity of data errors; and

[6] These findings are in contrast to those of a study carried out by Columbia Business School's Center for Excellence in Accounting and Security Analysis (CEASA), 'An evaluation of the current state and future of XBRL and interactive data for investors and analysts' (T. Harris and S. Morsfield, 2012). This can be explained by differences in the composition of the 'users' groups in each study—the 'users' in the Columbia study consisted of investors and analysts, while the 'users' in our current study consist primarily of data experts and intermediaries who provide software and data to investors and analysts.

- An exception is FpML: The stakeholders all agree that missing data errors occur frequently.

5.2.4 Maintenance and future development

5.2.4.1 Is the technology state-of-the-art?

- We find that across all the standards, the designers always believe that it is; and
- FpML: this is the only standard for which all stakeholders believe that it is.

5.2.4.2 Transparent change process

- Designers always state that the process is transparent and easy, while the other stakeholders generally do not agree; and
- The exception is ISO, where users agree with designers, and preparers are neutral.

5.2.4.3 Modification/updating of policies and approaches

- The final area that we focus on is how the users and preparers approach modifications and updates to technology within their own organisations; and
- Both stakeholder groups report that they have formal policies for internal technology and modifications related to the standard.

6 Analysis

6.1 Stakeholder harmony vs dissonance

Because our study does not empirically test actual data from each standard, we can refer to the extant research but cannot apply the same parameters and definitions directly. We can, however, look for ways to build on this literature to provide further understanding and new paths toward assessing the quality of a standard. Perhaps more importantly, we can make recommendations based on our findings that could enable the stakeholders for a given standard to design and implement standards that meet or improve the quality criteria demonstrated by Zhu and Fu (2009)[7] and Zhu and Wu (2011)[8].

[7] Zhu, H. and L. Fu (2009). Quality of Data Standards: Empirical Findings from XBRL. 30th International Conference on Information Systems. December 15-18. Phoenix, AZ, USA.

[8] Zhu, H. and H. Wu (2011). Quality of Data Standards: Framework and Illustration using XBRL Taxonomy and Instances. *Electronic Markets*, 21(2). 129-139.

With that in mind, we turn to our survey and interview results and note patterns as to the dissonance or harmony in the responses between and among stakeholders within and across the standards. Stakeholder responses are not always consistent within a particular stakeholder group. Dissonance and harmony emerge from our data as potential new constructs to consider in data standard development and evaluation. We do not test them empirically at this stage, but rather offer a descriptive view of where and when they occur over the life cycle and within the components of each standard that we have studied.
We summarise some of the more notable dissonance and harmony examples and trends. For this purpose, we define harmony and dissonance as follows:

- Harmony is defined as a context in which stakeholders generally agree with each other on a key concept examined in the survey.
- Inter-stakeholder dissonance is defined as a context where, for a given, standard, the different types of stakeholders do not agree on a key concept.
- Intra-stakeholder dissonance is defined as a context where individuals within a type of stakeholder exhibit significant disagreement amongst themselves about a key concept.

6.1.1 Areas of Stakeholder Harmony

- Formal processes and governance are common areas of stakeholder harmony.
- Content, syntax, technology generally appears to meet the needs of all parties.

6.1.2 Areas of Stakeholder Dissonance

- Although formal processes and governance are common areas of stakeholder harmony, dissonance among and between stakeholder groups is also very common
- There is confusion over by and for whom the standards were created
- The presence or awareness of detailed, specific use cases for all stakeholders and user experience testing is uneven
- There is inconsistent awareness of the availability of software and tools, and inconsistent reporting of costs and effort to learn
- Only designers regularly viewed the process for change as transparent and inclusive
- Only designers consistently viewed the standard as state-of-the-art.

These observations are valuable because they highlight that involving a small stakeholder group in the process of designing and implementing a standard may be insufficient to capture input from the stakeholder group as a whole.

6.2 Empirical assessment of data standard quality

In this section, we propose a method to measure the standard quality using the survey responses, and we also devise a ranking mechanism to present the quality of the six standards in comparison. The objective is to define a quantitative approach to measure the relative quality of the standards based on our proposed analysis framework.

As described above, we collected survey data related to issues in the following three areas: levering the standard, identifying common obstacles, and implementing ongoing development & evolution. A set of survey questions has been designed in this study to extract the relevant quality-related properties of the various standards from the survey respondents. The set of properties of standards are defined in Table 5, and the overall quality is defined using the following formula:

$$\boldsymbol{Quality\ Score = w_1 * Relevancy + w_2 * Coverage + w_3 * Usability + w_4 * Cost\ Effectiveness + w_5 * Resusability + w_6 * Governance + w_7 * Policies;}$$

$$\boldsymbol{where} \sum_{i=1}^{7} \boldsymbol{w_i} = \boldsymbol{1}$$

Table 5. Properties for Quality Measurement

	Properties	# of Questions	Example Questions
Leveraging the standard	Relevancy	4	How effective are they? (Terminology, Technical design, Reference model, Ontology)
	Coverage	8	How much of each did you assume that data users would actually utilize, exploit, etc.?
	Usability	4	How often do you actually utilize the following aspects of the standard?
Identifying common obstacles	Cost of understanding	3	How much effort do you spend in understanding the following aspects of the standard?
	Cost of implementation	3	If so, what is the monetary cost to learn and use them?
	Cost of maintenance	3	If so, what is the effort in man-hours to learn and use them?
Implementing ongoing development and evolution	Reusability	4	As an end user of the data or standard, how easy is it in general for you or your company's I.T. department to understand and implement changes to various aspects of the standard when updates occur?
	Governance	4	Did you consider ongoing technological evolution and development as part of your maintenance processes?
	Policies	4	Did you have formal policies in place for the design and modification of the standard in creating and maintaining?
Total		41	

We used a Likert-type scale to measure respondents' answers to the relevant properties in this study. When responding to a survey question, respondents specified their level of agreement or disagreement on a symmetric agree-disagree scale for a series of statements. Thus, the range captures the intensity of their agreement or disagreement for a given item.

We designed a set of 41 different 5-Likert scale questions for questions related to the quality of the model, and we used a weighted average to represent the responses over the full range of the scale. The weight represents the percentage of respondents agreeing with the statements regarding the quality of the standards; where the response points to disagreement with the question, we invert the weighted average with a maximum score. In the end, we combined the different aspects of the model with equal weighting and calculated a single ranking score to represent the relative quality ranking between the standards we studied.

Due to the fact that survey respondents answered questions only for those standards for which they have knowledge, the number of responses for each of the 41 questions might have a variant emphasis depending on the standards. In most cases, they indicate the specific standard for which they intended to answer the questions, while in other cases we rely on our general survey gate questions to assess their confidence level toward each of the standards. These self-assessments combined with the number of responses for each of the questions provides a means for us to present the overall confidence level for the ranking of each of the 6 standards. For example, for XBRL we have responses in 5 Likert-scales: Very strong, Relatively strong, Neural, Relatively weak, No experience.

We then interpret the confidence level of a particular respondent for XBRL as 100%, 75%, 50%, 25%, and 0% respectively. As we pool all responses together, we created a weighted confidence based on population distribution. These weighted confidence cores therefore are presented along with the ranking scores to have a comprehensive qualitative representation of the standard qualities.

Figure 4 – Quality Ranking of Financial Standards

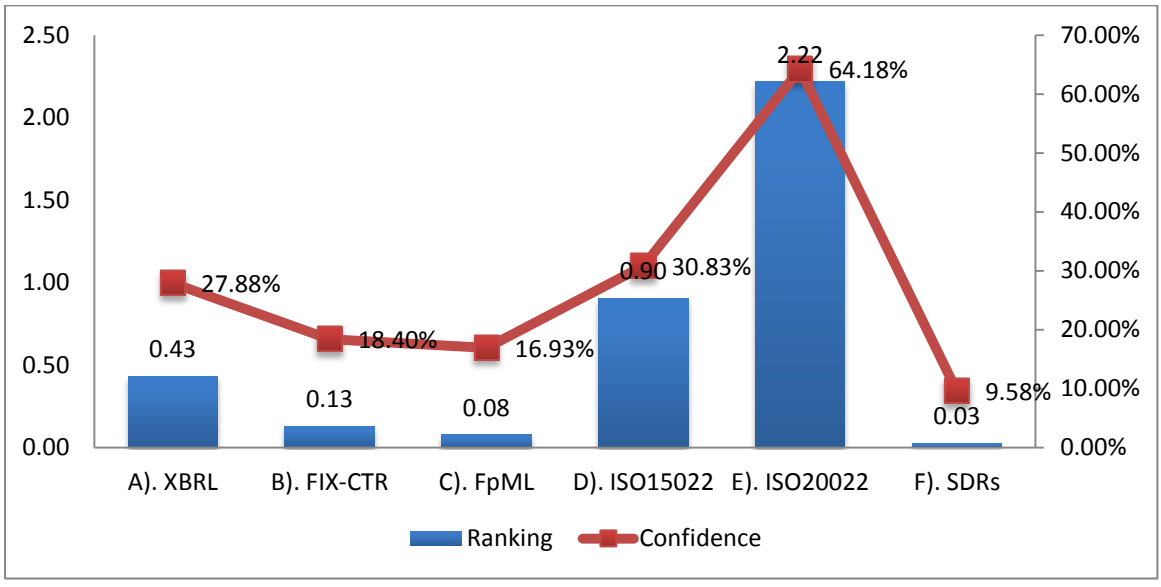

From this quality analysis of the six financial standards based on the 269 respondents, we conclude that ISO20022 has the highest standing among all 6 standards. There is also a relatively high overall confidence with this assessment. ISO15022 is placed the second only to ISO20022. This result is consistent with the fact that ISO20022 has become a replacement

for ISO15022 recognised by the industry.[9] This finding, in fact, further confirms the coherence of our proposed quality measurement approach. Moreover, among the other "semantically" complex standards, XBRL stands better than FIX, FpML and SDRs. But, the relatively low-weighted confidence scores for these four standards might be a reflection of the complexity of these standards. Unfortunately, our extremely small sample size for SDRs prevents for us to make any meaningful conclusion about this standardisation effort.

We readily acknowledge that small sample size in some cases, and assessing quality based on stakeholder views contains some inherent limitations. Nevertheless, we are confident that our assessment model does capture the essential elements of quality measure for these standards. The relative standing is informative for future improvements of these standards. The results also do provide initial support that there is merit in developing and assessing standards using our formalised model and 360-degree approach. We encourage additional research testing our model and approach.

[9] These quality findings represent an unbiased, independent empirical assessment, driven by the design of the empirical model. Alternative specifications of a model may find different results. However, we believe our model captures important, but hereto unstudied, aspects of data standards' quality.

7 Conclusions and Recommendations

We have provided and tested a proposal for a multifaceted model for developing, maintaining, and analyzing data financial data standards. The evidence gathered suggests that it would be a useful tool for the development of future standards, as well as for adjusting the approaches to those standards currently in use. The evidence also suggests that a new construct—stakeholder harmony vs dissonance—may play an important role in quality and outcome of the implementations of standards. The findings also support the importance of regular and effective interaction with all categories of stakeholders not only at the initial design, but throughout the entire life cycle of a standard.

Finally, our survey asked about the future—we asked for recommendations from the respondents for "better" standards or technologies, if they felt that the current one was not state-of-the-art. The responses to that question were quite limited, with only one approach being mentioned more than once—i.e., RDF/OWL. The Financial Industry Business Ontology (FIBO) is an industry initiative to define financial industry terms using RDF/OWL, and is currently under development by the Enterprise Data Management Council. In James Hendler's (one of the creators of RDF/OWL) view, RDF/OWL, while an important part of the future, is not a stand-alone solution.

As with many disruptive technologies, the emerging and new technologies that spring from them go beyond the original vision of their early users. Hence, keeping this entire evolving ecosystem in view, not just an RDF/OWL ontology, is important. According to Dr. Hendler, what is also key (and not unrelated) is not just a new or disruptive technology or language, but a change of mindset. A change from thinking about data solely in terms of relational databases and being physically stored somewhere, to seeing a world of possibilities from the linking and accessing of widely-dispersed data across the web (and anywhere else it might be, including within relational databases). Finally, he recommends that with any major change, such as standardising vocabularies, those overseeing the standardisation should start small—eg standardise the first 10 items, and test that they work before proceeding to the next 10 and so on.

So, where does that leave us? From here, we find a fitting connection between the London Whale, Roy Marsten, and our study. The future may well be more about our mindsets than about a silver bullet technology or a life-saving structure that we can just superimpose on the data and information that we are creating, and accumulating at ever increasing speeds and amounts. Although we may not be able to identify a single next big technology disruption, history tells us that disruption is on its way. We suggest that financial data standards be flexible and adaptive wherever possible, and not be tied to a single technological approach or coding language. A change in mindset starts here.

Some market standards must-dos:

Market standards must meet the following challenges to promote and maintain industry-wide adoption.

They must:

• Fulfill the needs of **users**, not the perceived needs of 'professional standardizers'.

• Respond quickly to business and market **changes**.

• Stay **relevant** to the businesses and markets that they support.

• Be **inclusive** and allow the contributions and feedback of stakeholders and users to be included in the standards development and revision processes.

• Be **international** in scope and internationally accepted.

• Have continued access to, availability and participation of the necessary **technical and subject matter experts** as standards development is a voluntary effort. They must attract new participants and stakeholders to the standards development process.

Those promoting, developing and using standards must maintain and increase **collaboration**…

Karla McKenna, *representing ISITC in XBRLglobal Vol. 1 Issue 2 – Aug 2010*

What does it take to find the London Whale? The answer may lie somewhere in between data standardisation and data analytics, but robust data standards, where necessary, are a critical component. Standards should also consider use cases that go beyond a traditional data model, whether relational or ontology-driven.

For single transactions, when money needs to be moved from one place to another, how the data are stored and retrieved may not be relevant. But when someone is looking for that transaction in a sea of other transactions, high caliber analytical tools must be brought to bear that can quickly analyse vast quantities of data.

In all cases, we can safely say that standardisers should keep an eye and ear attuned not only to new technologies, but to the ecosystem and adaptations that arise around them. And, above all, they should be open to changing their mindset, whenever warranted.

With that, we supplement the advice of Karla McKenna (see textbox) with detailed, practical recommendations which result from our detailed interviews with designers, our 360-degree surveys, and our empirical tests. Figures 5a and 5b below build on the outline of the survey questions in Figure 3, and provide our recommendations by each category of our research questioning and analysis.

Figure 5a – Recommendations by Category of Survey Questions: Implementing and Leveraging the Standard

Implementation and ongoing development/evolution

Initial development of the standard

- Find and keep all categories of stakeholders at the table as the standard is being designed.
- Ensure agreement about who the standard is for.
- Be aware of and guard against the very strong tendency of standardisers to be overly optimistic about the usefulness and ease of use of the standard.
- Documentation does not equal shared vision and knowledge over the life cycle of the standard; supplement documentation with ongoing and detailed interaction with all categories of stakeholders that reexamines the vision and knowledge at regular intervals.
- Start small, and ensure that the standard works for the intended user from the onset, then test and retest that it continues to work for them.

Governance and modeling

- Pay attention to contrarian voices, looking for the kernel of truth in their concerns or criticisms.
- Use professional skepticism for those who are overly optimistic or who think they are the smartest in the room.
- Avoid having the expertise and decision-making about a standard's design and implementation reside with only one or two people.

Use cases

- Base the design on actual use cases from all types of stakeholders who will be consuming the data that the standard is transmitting. Do not substitute what one imagines a user will want with deep knowledge of what they do and how they do it, and actual conversations with them.

Leveraging the standard

Content

- Develop and maintain a set of stable content, commonly agreed upon, based primarily on input from the primary users/consumers of the data, rather than on abstractions and ideas formulated primarily by designers (unless they are intended users); this is linked to the notion of 'use cases' above.

Tools

- Use formal tools wherever possible in the design of the standard; conduct extensive formal user testing on both the tools and the resulting data.
- Make interaction with the standard itself as invisible as possible to the filer/preparer, as well as to the end user. If a non-technical person cannot use the standard without extensive training either to input or to consume the data, then you haven't met this goal.

Figure 5b – Recommendations by Category of Survey Questions: Implementing and Leveraging the Standards: Identifying obstacles, Maintenance and future development, and General

Identifying obstacles

Using the data

- Standardisers generally were completely wrong on how the users/preparers would find interacting with all aspects of the standards and the data it created; as a result, we would repeat again our recommendation in the Tools section above.

Errors

- Electronically-submitted data via a data standard should not be different than the related paper or HTML-submitted data, if relevant.
- Data quality issues will hinder the best of data standards; the extent of their presence implies that effective user testing was not conducted prior to release of the standard; as a result, we would repeat again our recommendation in the Tools section above.
- Hand-waving errors away as growing pains when the data are 'live' is not advisable; the underlying data are being used to support the global financial and securities markets, which means that errors can be extremely costly beyond the immediate data point(s) affected. Where external audits of the data are available and appropriate for ensuring data quality, they should be required.
- The context of the errors also matters and users rightly expect zero errors in the settings we studied; contending that commercial databases contain more errors than the actual data transmitted via data standards does not excuse or minimise the seriousness and cost of the latter errors.

Maintenance and future development

- Publicly document how the standard's design is 'future proof'; at a minimum, backwards-compatibility must be provided.
- Ensure that future-looking experts provide an unbiased, independent review of the standard and its functioning (and of other technologies available) at regular intervals over the life cycle; designers nearly unanimously viewed their standard more favorably than those who had to use it either to prepare or consume data, and as such, are not the best party to independently evaluate this issue.
- Ensure that user and preparer concerns are adequately addressed; although all standardisers would assert that this part of the process functions well, our survey findings suggest otherwise.

General

- Survey and 'resurvey' all stakeholders regularly, paying particular attention to areas where stakeholder dissonance is present and persistent.
- To facilitate the above recommendations we suggest using the questions from our survey as a checklist for enabling stakeholder harmony not just at the initial development stage, but throughout the entire life cycle of the standard; we noted that although all designers asserted that documentation of various decisions was a part of their process, the ensuing amount of dissonance provides evidence that something about the types or amount of documentation was not especially effective.
- Utilise the most formal methods and tools available for designing the core elements of standards.