White paper:
Interconnectivity enabled through XBRL

Authors: Marc Houllier, Bodo Kesselmeyer, Karl Magnus Westerberg, Gilles Maguet, Patricia Myles, René van der Meij, Tord Andersson – Members of the XBRL Europe Sustainability/ESG Working Group

Date: 14 January 2022

Call for feedback

This paper has been prepared by members of XBRL Europe's ESG/Sustainability working group.

It aims to **provide guidance** on how to **develop XBRL taxonomies relating to Sustainability/ESG topics**, while creating **interconnectivity** with financial taxonomies thereby enabling optimal usability of the data.

We would appreciate feedback on this document with additions and comments, which will contribute to making this paper as valuable as possible.

We would be grateful to receive your feedback at info@xbrl-eu.org before the 15th of February 2022.

Executive Summary

This is the first paper in a series of four papers discussing, from the viewpoint of various types of stakeholders, the use and benefits of XBRL for the materialisation of relevant interconnections between domains (such as Sustainability/ESG reporting and financial reporting).

We highlight the need for such materialisation by demonstrating the loss of clarity and usability of digital data that happens when no connections are drawn between digital standards.

The XBRL standard provides a wide array of alternatives for the modelisation of such interconnections. We present those alternatives from a conceptual point of view while providing technical references and examples.

We also provide some guidance on when each alternative is most appropriate by evaluating its genericity, the value it creates, and the amount of work it requires from the standard-setter and software providers.

Table of contents

1. Introduction

With the legacy and experience of the last 20 years from XBRL enabled reporting and recently on ESEF reporting, our XBRL community at large and XBRL Europe in particular has assessed that with the emergence of the sustainability/ESG reporting, it could be of importance to underline :

- that considering the digital representation of ESG reporting requirements, concepts and their relationships with each other as early as possible is necessary
- that the necessary interconnectivity between ESG data and financial data is enabled through XBRL.

This document prepared by the XBRL Europe Sustainability/ESG working group commences a series of papers about interconnectivity between ESG data and financial data enabled through XBRL.

Multiple standard-setters are currently working on developing XBRL taxonomies for exchanging ESG and Sustainability reports. Examples include:

- The Corporate Sustainability Reporting Directive (CSRD) in the European Union prescribes the use of XBRL, for which the European Financial Reporting Advisory Group (EFRAG) is delegated to design the content of the taxonomy by June 2022.
- The International Sustainability Standards Board (ISSB) is consolidating with the Value Reporting Foundation (VRF) and the Climate Disclosure Standards Board (CDSB) which will likely result in a renewed XBRL taxonomy by June 2022 also.

This paper aims to support these and other standard-setters in their work to define technical standards for sustainability data based on XBRL and underlines the importance of interconnectivity for all stakeholders.

This series of four papers will establish the following:

1. Document 1: Generally speaking, that interconnectivity must be included in the technical models, and that XBRL taxonomies provide flexible options to do so
2. Document 2: Targeting standard-setters, in which interconnectivity should be considered in architecture choices
3. Document 3: Targeting data consumers, in which technically modelled interconnectivity makes integration of ESG data into decision-making models easier
4. Document 4: Targeting issuers, in which data collection and reporting can be made more efficient by considering interconnectivity. The papers are divided according to target audiences so that their message can be delivered more efficiently.

Consequently, there will be some overlap in the topics discussed in the four papers. Standard-setters should of course still be interested in reading the third and fourth paper to understand in more detail the effect of their taxonomy design choices.

No technical experience with XBRL is required to read this series of papers, but readers are encouraged to familiarise themselves with the issues of structured digitisation and/or with the language at the conceptual level[1] to best enjoy the following content.

This document is thus the document N°1 of the series.

1.1. Interconnectivity as an objective

[1] Introductions to the standard can be found on the website of XBRL International, Inc. . Although we voluntarily limit our use of technical vocabulary in these papers, readers are encouraged to consult the glossary to familiarize themselves with terms commonly used in publications on the subject.

Regulated periodic information issued by companies contains data from diverse categories. Even when that data is gathered in large documents such as annual reports, these documents are often divided into sections that make it clear which regulation each section corresponds to.

Consequently, data that is closely related from a stakeholder's point of view but pertains to different domains (governance, environmental and social responsibility, financial information) might be scattered throughout the document, making it much harder to find. Even more importantly, given that different sections of the document correspond to different regulations, it is not always clear whether similar concepts have consistent definitions across these sections.

A pair of domains that have come under a lot of attention lately are the combination of financial data and ESG data.

Data consumers seek to use both kinds of information concurrently, regardless of whether their primary objective is finance-oriented or sustainability-oriented.

- Financial analysts use ESG data to assess risks and/or forecast costs directly linked to that data, as they seek to forecast the future cash flows of a company.
- Investors need financial data to compute hybrid metrics[2] that allow the comparison of similar initiatives, as they aim to support sustainable projects and avoid greenwashing.

Many standard setters have stated the interconnection of these two domains as important, and its improvement as an objective in the design or evolution of their standards[345].

The examples we use in this paper will be drawn from the connectivity between these two domains, but all considerations are applicable to any pair of domains.

[2] Meaning metrics calculated from a financial measure and a non-financial measure. For instance, **energy consumption divided by revenue** can be used to compare the energy consumption efficiency of similar activities of differing scale.

[3] EFRAG in Proposals for a relevant and dynamic EU sustainability reporting standard setting :
[24]"*[the European Standard Setter] should define through appropriate guidelines methodologies and processes enabling connectivity (direct and indirect) and reconciliations between financial reporting, under IFRS or local GAAP, and sustainability reporting.*"
[103]"*[anchoring of sustainability information into the management report] acknowledges and conveys that financial and sustainability information are interconnected and jointly required to understand an entity's full 'story'*"

[4] Global Reporting Initiative (GRI) in Response to IFRS Consultation Paper on Sustainability Reporting : *"The interconnection between financial reporting and sustainability reporting described here deserves particular attention in the future standard-setting activities of the IFRS. It is essential to limit the burden on reporting entities while enhancing the effectiveness of disclosure."*

[5] EFRAG in EFRAG Consultation on IASB Agenda : "IASB identifies a separate area of its activity to address the connectivity between financial reporting and sustainability reporting and increases the resources devoted to digital reporting"

1.2. Technical representations of interconnectivity

Though interconnectivity has been widely recognised as an important issue, it is still often vaguely defined, and its implementation is a principle-based approach. With no rules to rely on, it is difficult for software providers design and develop tools that assist issuers in creating that interconnectivity (for instance, referencing other parts within a document or other documents). Consequently, issuers and their auditors must do this task manually and often have trouble figuring out how to properly interconnect different parts of their reports in practice.

This is in contrast to the recent approach of designing the technical modelisation of reporting standards concurrently with the standards themselves, an approach where the standard-setter defines a report in very concrete detail, providing issuers with a lot of useful guidance and setting up a process where usability and comparability of the resulting data is greatly improved.

We assert that interconnectivity should be considered early in the design process of standards and that it should be made concrete and usable through the use of technical standards. In particular, XBRL taxonomies and Inline XBRL have successfully enabled the digitisation of several types of business and institutional reporting. We will show that interconnectivity can be easily modelled through the use of the XBRL standards.

XBRL uses **taxonomies** as a means of providing a shared set of concepts (e.g. Revenue, Goodwill, Employee Turnover) in order to ensure comparability between reported data. Such taxonomies are commonly defined by standard-setters and made publicly available. They provide the issuer with guidance on how to create the report; and they provide data consumers with rich metadata thereby clarifying the meaning of consumed data.

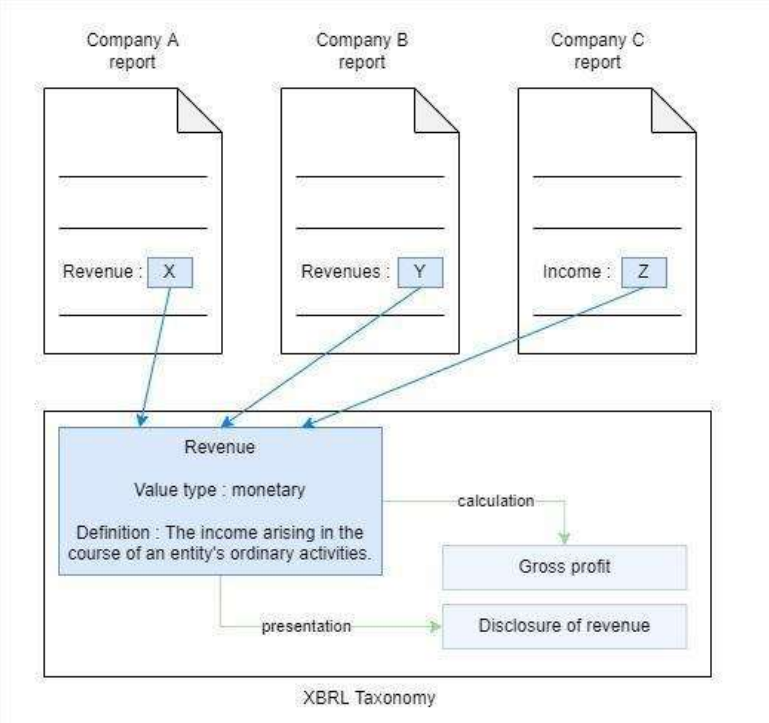

Figure 1 : Three companies report their revenue using a varying vocabulary, but since the figures all point to the same concept defined in a XBRL Taxonomy, a data user will know they are comparable. The taxonomy also provides useful additional information about the concept and its relations to other concepts.

From a standard setting point of view, it makes sense to develop financial and non-financial standards mainly separately, and therefore it also makes sense that their technical representations (XBRL taxonomies) should also be separate.

There is no technical issue with the use of two different base taxonomies within the same report, and there are several ways this can be achieved using XBRL.

However, when data marked with two taxonomies is expected to be used together (as is the case for financial and ESG data), several issues may arise.

- **Where the domains covered by the two taxonomies overlap, are the definitions consistent?**

For instance, ESG intensity metrics might involve calculations based on « revenue » or « enterprise value ». Research and development expenses into sustainable solutions may be mentioned in the report. It is of the utmost importance to the analyst to know whether these share the same definitions and accounting policies as the related metrics in the financial statements.

- **Are the taxonomy structures complete enough to match all users' intuition?**

In addition to providing common definitions of concepts, taxonomies also define standard relationships between those concepts. These relationships help clarify the meaning of concepts, but most importantly they provide a means to **discover and navigate** through the taxonomy concepts. Taxonomies commonly include several thousands of concepts, and a user should not be expected to have perfect knowledge about which concepts were included in the taxonomy and how they were named and/or defined in it. A taxonomy should therefore define a limited list of **entry points**, and relationships that allow any user to navigate to the intended concepts from there.

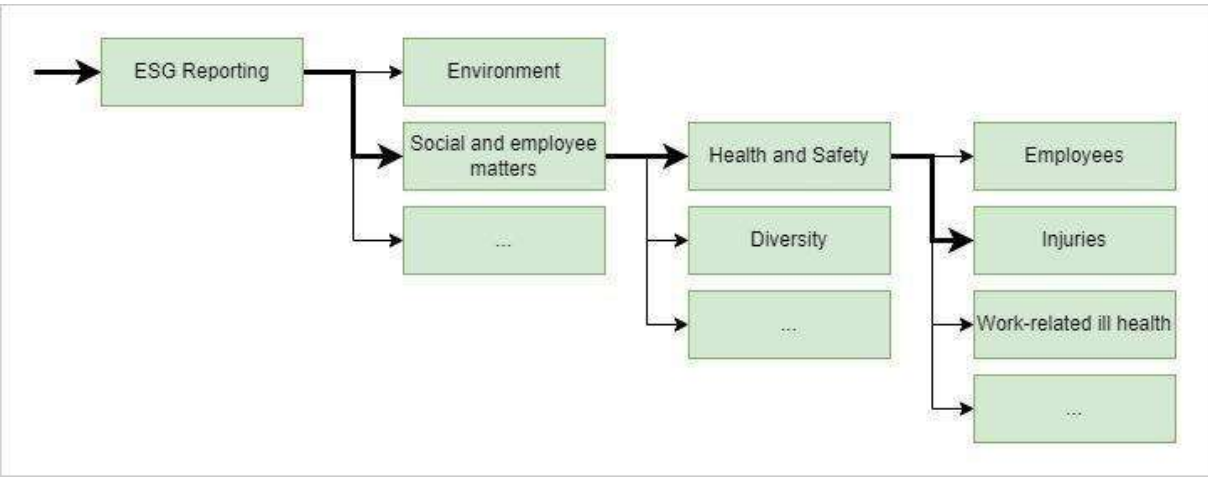

Figure 2 : A user is looking for data on work-related physical trauma. A well-built taxonomy contains the relations that allow them, starting from an intuitive entry point, to navigate to the corresponding concept in the taxonomy.

When two or more taxonomies are used, the issue is still the same: do taxonomy relationships modelise the logical connections a user is likely to follow? If two taxonomies cover related domains, then it is likely that a user should wish to transition from one domain to another at some point in their data discovery, and it is only possible if a relevant connection was drawn between the two taxonomies.

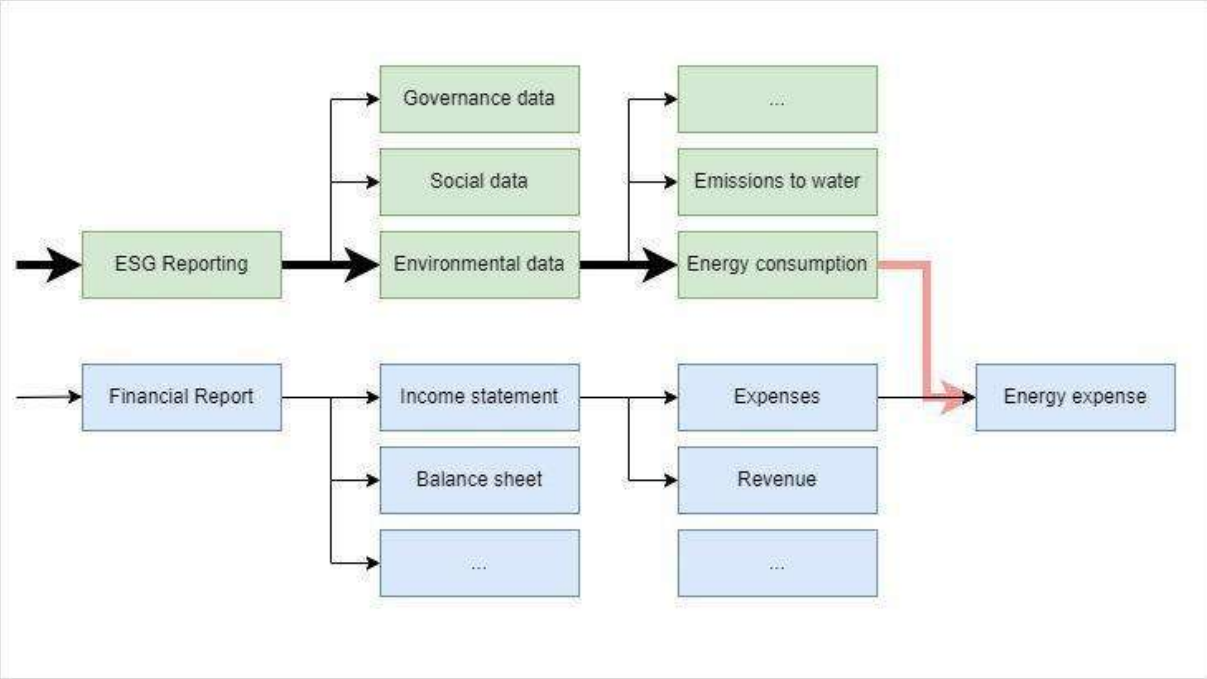

Figure 3 : Consider the above two non-financial (green) and financial (blue) taxonomy structures. To avoid duplicate definitions, the standard-setter chose not to redefine energy expenses in the ESG taxonomy. A user exploring the data looking for energy expenses and expecting it to be part of "energy consumption" might follow the bold path, but will fail to find the intended concept if the red arrow that links both taxonomies has not been materialised.

2. Options for the representation of interconnectivity in XBRL

XBRL is **extensible** in nature. This means that the language is designed to make it easy to create new concepts and new relationships between concepts, without any need to revise previously published documents.

In the context of modelling interconnectivity between two domains (such as financial reporting and ESG reporting), this means in particular that it is easy in a given taxonomy to create relationships between concepts defined in that taxonomy and concepts defined in any other taxonomy, **without having to involve the creator of that other taxonomy.**

A relationship in XBRL is usually modelled as an **arc,** an oriented arrow between two related concepts.

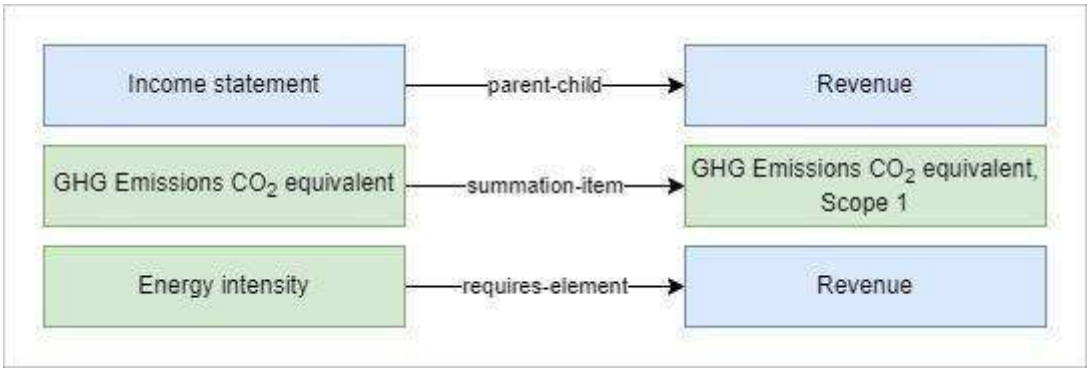

Figure 4: Examples of arcs drawn between concepts. Arcs can be drawn between concepts of the same taxonomy or concepts of separate taxonomies. Each arc also carries a "role" that specifies its meaning.

Regardless of the meaning of the arc, having a connection between two concepts creates a way for users to technically navigate from one to the other, which makes a considerable difference from having no connection at all. The direction of the arrow does not matter much for that purpose either.

Concepts can be considered as "connected" even when an arrow does not create a direct link between them. What matters most is that there exists some path, linking one to the other, that agrees with business logic.

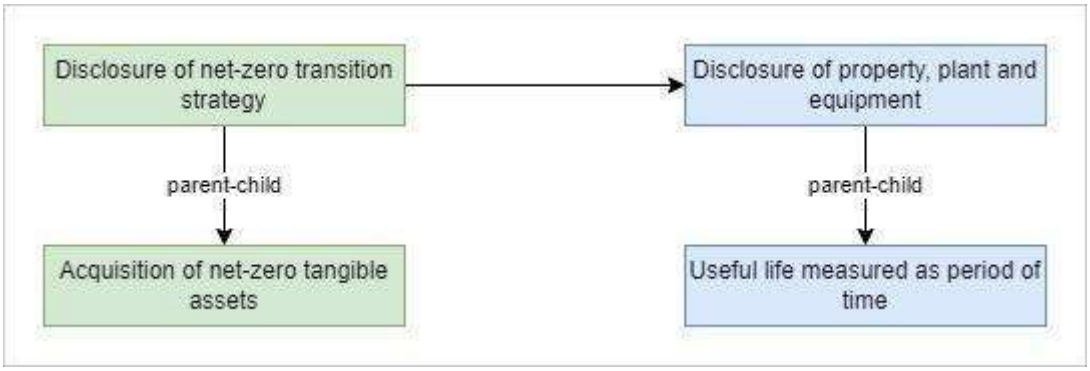

*Figure 5: An indirect connection between **Acquisition of net-zero tangible assets** and **Useful life measured as period of time** has been established by a direct connection between their respective parent concepts. Although it is not explicitly stated what the relation between these two concepts is, this indirect connection still allows software to show a concept as of potential interest to users looking at the other concept.*

Of course, the longer the path between two concepts is, the less clear its meaning and relevance will be for a user. Many factors should be considered when deciding the right level of granularity for creating direct relationships between two taxonomies. That will be discussed in a later paper dedicated to architecture choices.

The meaning of a relationship and the level of detail with which a standard-setter wishes to describe an interconnectivity relationship is expected to vary, depending on the nature of the relationship and on the resources available.

We will describe hereafter a few alternative ways to describe interconnectivity relationships using XBRL.

2.1. Overview

XBRL offers a wide array of options to represent an interconnection between two taxonomies. Each option is the most appropriate in a specific scenario:

- **Presentation arcs** create a simple means to navigate between concepts useful to both human and automated users, and only require the identification of related concepts.
- Other **standard arcs** can be used to define more precise technical relationships that can be used in automated validation and are already well processed by software.
- **Non-standard arcs** can be used to define other recurring technical relationships or to define new domain-specific relationships, but require careful specification by the standard-setter.
- **Arc titles** can be used to add free-form text to provide human users with documentation about a specific relation between two concepts.
- **Formulas** provide a standard way to describe other logical or mathematical relationships, but require a more complex technical design.

	How many relationships can the relation be applied to ?	What information is conveyed ?	How much work is required to define and maintain it ?	How much work is required from software providers to work with it ?
Presentation arcs	Can always be used.	No information beyond the existence of a connection.	Simplest.	None (standard).
Standard arcs	Requires a close relationship between both taxonomy concepts.	Specific technical information.	Requires careful consideration of consequences on the validity of documents.	None (standard).
Non-standard arcs	(Depends on the meaning behind the relationship)	Expected to be particularly relevant.	Requires definition of relationship, documentation and support.	Depends on the relationship and its specification, generally non-negligible.
Arc titles	Can always be used, but there is not always relevant information to add in a title.	Since it documents a specific relationship, expected to be quite specific.	Potentially large amount of work if many connections are documented.	Little (standard but seldom used).
Formulas	Requires a close relationship between both taxonomy concepts.	Specific technical information.	Requires careful consideration of consequences on the validity of documents, even more so with complex formulas.	None (standard).

2.1.1. Relationships specifically designed for exploration (presentation arcs)

It may happen that there is a consensus that the availability of certain data about ESG concept A is useful when dealing with financial concept B, but that there is no consensus on how exactly that data contributes. It would then be natural for a standard-setter to take a more principle-based approach and ask issuers to disclose data about concept A and explanations of its effects on B.

Figure 6: There would be no formula linking directly the number of incidents to the corresponding provision. In fact, an increase in incidents reported could be the result of either an actual increase in incidents or an improved monitoring.

Even though it cannot be precisely defined, a relationship still exists between the two concepts that should be expressed in the taxonomy. It would ensure that software can highlight the availability of relevant data about concept A (and help navigate to it) when a user is looking at data about concept B.

XBRL proposes a standard relation to be used explicitly for exploration. Such arcs are called **presentation arcs.** They usually describe a parent-child relationship and are meant to be assembled into hierarchies such as the ones shown in figures 2 and 3.

When using non-specified relationships such as presentation arcs, it is likely that if a relationship is drawn between A and B, then there also exists an interesting relationship between concepts close to A and concepts close to B. Defining all possible interesting relationship scenarios directly would be a difficult task.

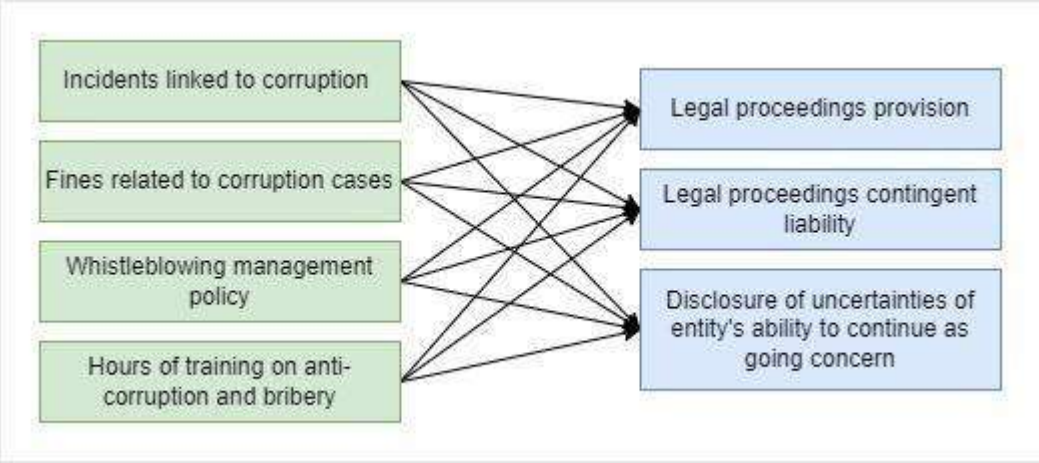

Figure 7: Most indicators linked to bribery or corruption cases could be argued to have an impact on future legal proceedings expenses. It would generally be impractical to try and define direct relationships exhaustively for all possibly relevant connections.

Defining a single relation between higher-level concepts is an interesting approach but is only possible if our concepts share common “parents”, and if the relation is still understandable at that high level.

A more consistent approach would be to use the XBRL notion of **abstract** concepts, designed especially to make the creation of relationships easier. Using an abstract concept, we can group together any concepts from both taxonomies for which we want an exploration connection to exist.

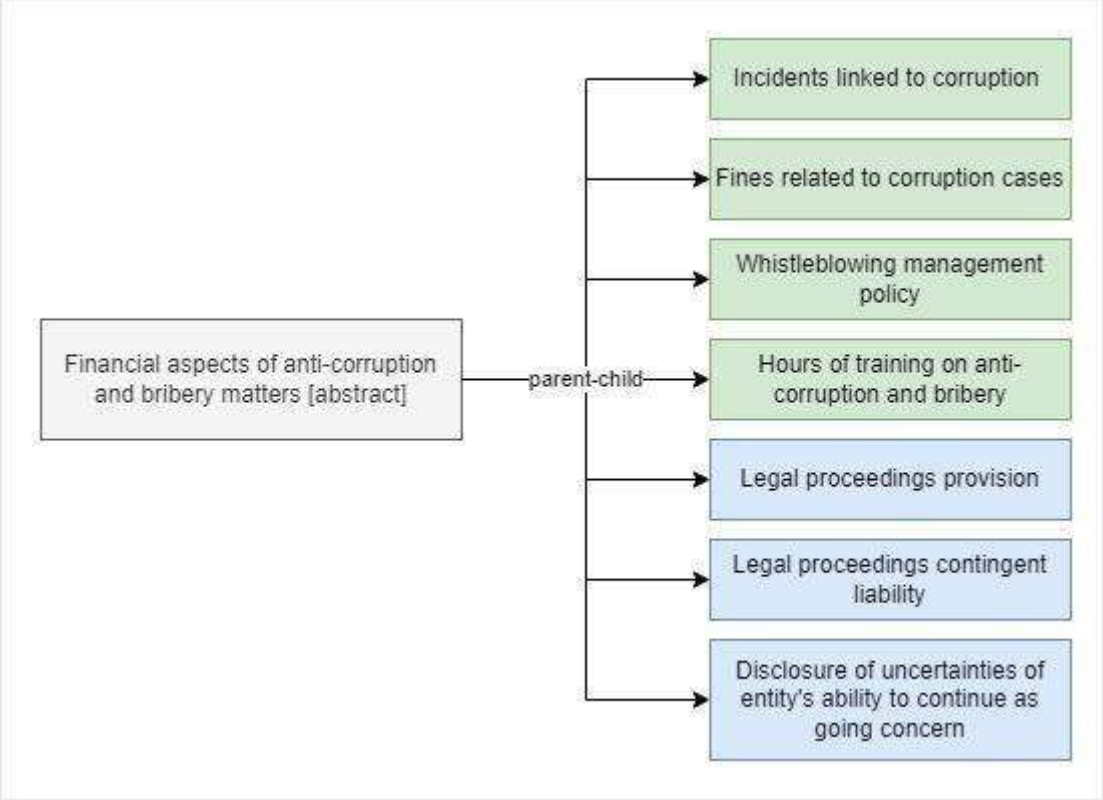

Figure 8: Related concepts can be gathered as children of a common abstract concept. The name of that abstract concept can optionally be used to document the common thread between the concepts if any.

2.1.2. Relationships that carry a simple or domain-specific meaning

If more is known about how concepts **regularly** relate to one another, then there is an opportunity to include that information within the taxonomy with a standard representation so that software providers can develop new functionalities.

The XBRL specifications define a handful of **standard arcs** that carry technical meaning. Examples include::

- A relationship to define that two concepts are equivalent. For instance, that relation could be used to define financial concepts in an ESG taxonomy and identify the corresponding concepts in financial taxonomies.
- A relationship to define the arithmetic relation between several concepts.

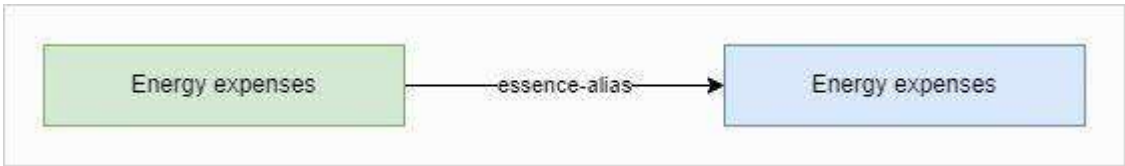

Figure 9: If our ESG taxonomy is likely to be used as standalone, it could be useful to include a "copy" of a financial concept, for instance energy expenses. In this case, users should be made aware that the two concepts are identical and interchangeable. The "essence-alias" XBRL relationship is a standard relation that can be used for this exact purpose.

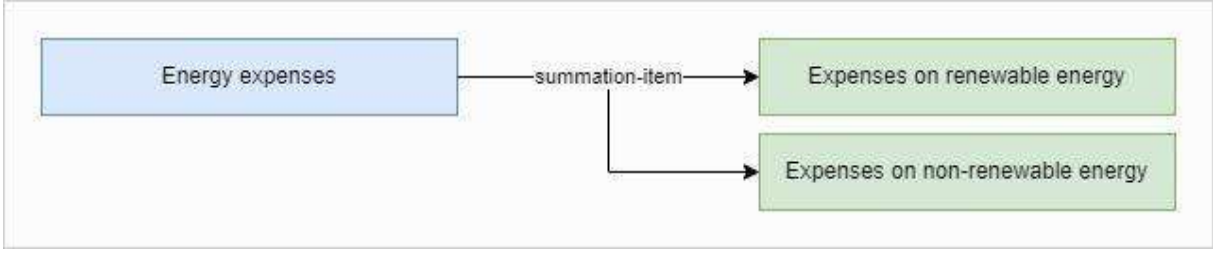

Figure 10: Alternatively we decide here that in the ESG report, energy expenses should always be disaggregated between those on renewable and non-renewable sources. An exact arithmetic relationship with the financial concept can be documented using the standard "summation-item" XBRL relationship.

Because they are standard relationships, they are well implemented by most available software available to all kinds of users.

If the identified recurring type of relationship does not match one of the standard relationships, XBRL has mechanisms in place that allow the creation of new types of relationships (**custom arcs**).

In essence, the only technical work necessary to create a new type of relationship is to define a "name" for it.

Figure 11: We want to create a new relationship that specifically links a numerical concept to the concept that should contain disclosure of hypotheses used in the calculation of the numerical concept. The only step is to choose a unique name for it; to ensure the name is unique, we prefix it with a domain name we own (here, the fictive www.xasb.org).

As an illustration, the code used to define the anchoring relationship used in ESEF can be found in Annex 1 of this paper.

Of course, for the new relationship to be understood and usable, its meaning needs to be well defined. This is usually best done through free-form specifications that involve domain-specific language. Although that documentation is free-form and is expected to be included in "paper" specifications, it can also be technically linked to the newly created relationship so that human users can access the documentation easily when using the relationship.

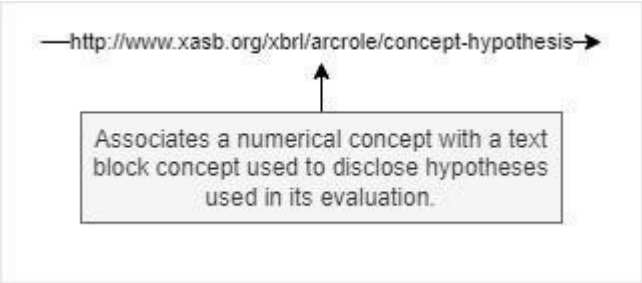

Figure 12: We can attach a documentation to the definition of our concept-hypothesis arcrole. This documentation will help human users differentiate between the different interconnections we create between our two taxonomies. It can also help software editors to better understand the relationship and implement specific functionalities for it.

As an illustration, the code corresponding to the definition for the anchoring relationship can be found in Annex I.

For available software to implement this new relationship in the most user-friendly ways, attention should be paid to defining in particular:

- The kind of concepts the relationship can be applied to (should they be expressed in similar units? Should they usually be defined over similar periods? Is it alright to define a relationship between an actual concept and an "abstract" one that represents a category of sub-concepts? etc.)
- What can and what cannot be inferred from the relationship, with examples if possible.

The more documented the relationship is, the more software is likely to implement the new relationship (and implement it *properly*).

Another way to improve the quality of implementation is to increase the exposure of the relationship, notably by sharing it with other standard-setters. The XBRL International organization maintains a registry of non-standard but common relationships (here) in order to attract the attention of the community to them and maximize their longevity.

For instance, the anchoring relationship defined by ESMA is hosted by XBRL International in its registry: http://www.xbrl.org/lrr/arcrole/esma-arcrole-2018-11-21.xsd#wider-narrower . Any standard-setter with a need similar to anchoring can now use this relationship in its requirements and take advantage of software having already implemented the relationship.

2.1.3. Documenting a relationship in particular (arc titles)

For relationships where the standard-setter wishes to include additional information regarding the meaning of a relationship between two concepts in particular, the XBRL standards allow the creator of a relationship to add free-form documentation to a specific instance of a relationship between several concepts.

The relationship can be any of the previously mentioned types of relationships, whether it is standard or not.

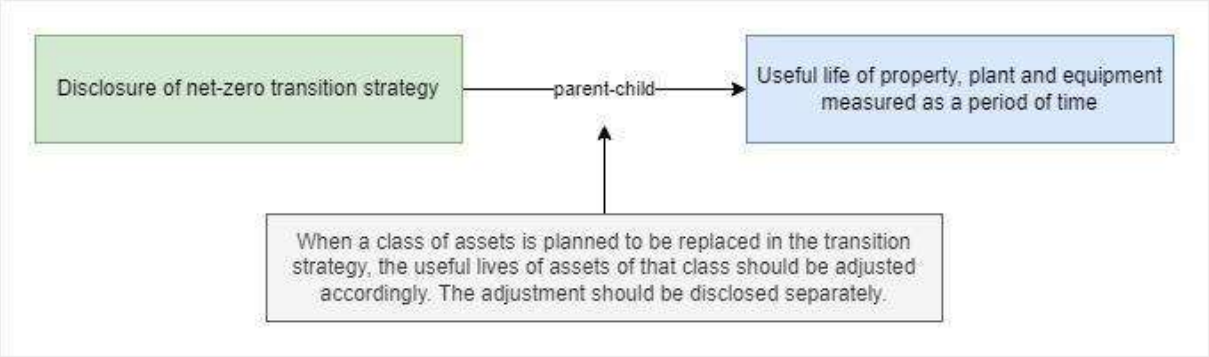

Figure 13: Here, we want to add a comment not to all "parent-child" relationships, but to this specific one between the strategy for net-zero transition strategy and useful life of assets assumptions used in financial reporting. The comment makes it clear why the two concepts should be considered together, and also includes requirements for issuers.

Software can then display that documentation when a user navigates from one concept to the other using that relationship.

Since the documentation is only applied to one specific instance of a relationship, documenting all interconnections this way provides very relevant information but requires additional work.

2.1.4. Relationships that carry a non-standard technical meaning (formulas)

If the relationship between two concepts can be expressed with logical or mathematical operators, then XBRL provides standard means ("validations" or "formulas") to express any such relation with its technical meaning within the taxonomy.

Figure 14: When the previously "essence-alias" relationship is not sufficient, more complex requirements can be described using formulas. Formulas can also be used to adjust levels of requirement ("should" vs. "must") if the standard relationship is deemed too strict.

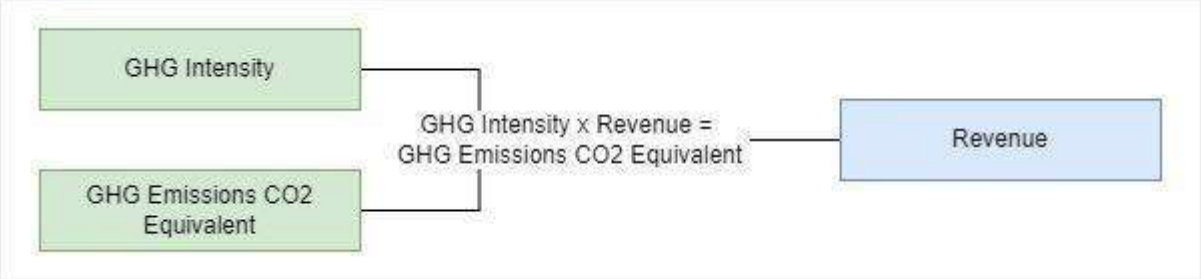

Figure 15: Such equations (other than weighted sums) and inequalities should be defined using XBRL formulas.

Since XBRL formulas are a well-used standard, all XBRL-compliant software would be able to use these relationships, with no delay needed for implementation and no opportunity for misinterpretation.

In practice, this kind of relationship is often used for inference and/or validation purposes. The standard provides ways to define messages that should be displayed to users when the expected logical or mathematical relationship doesn't hold in a document.

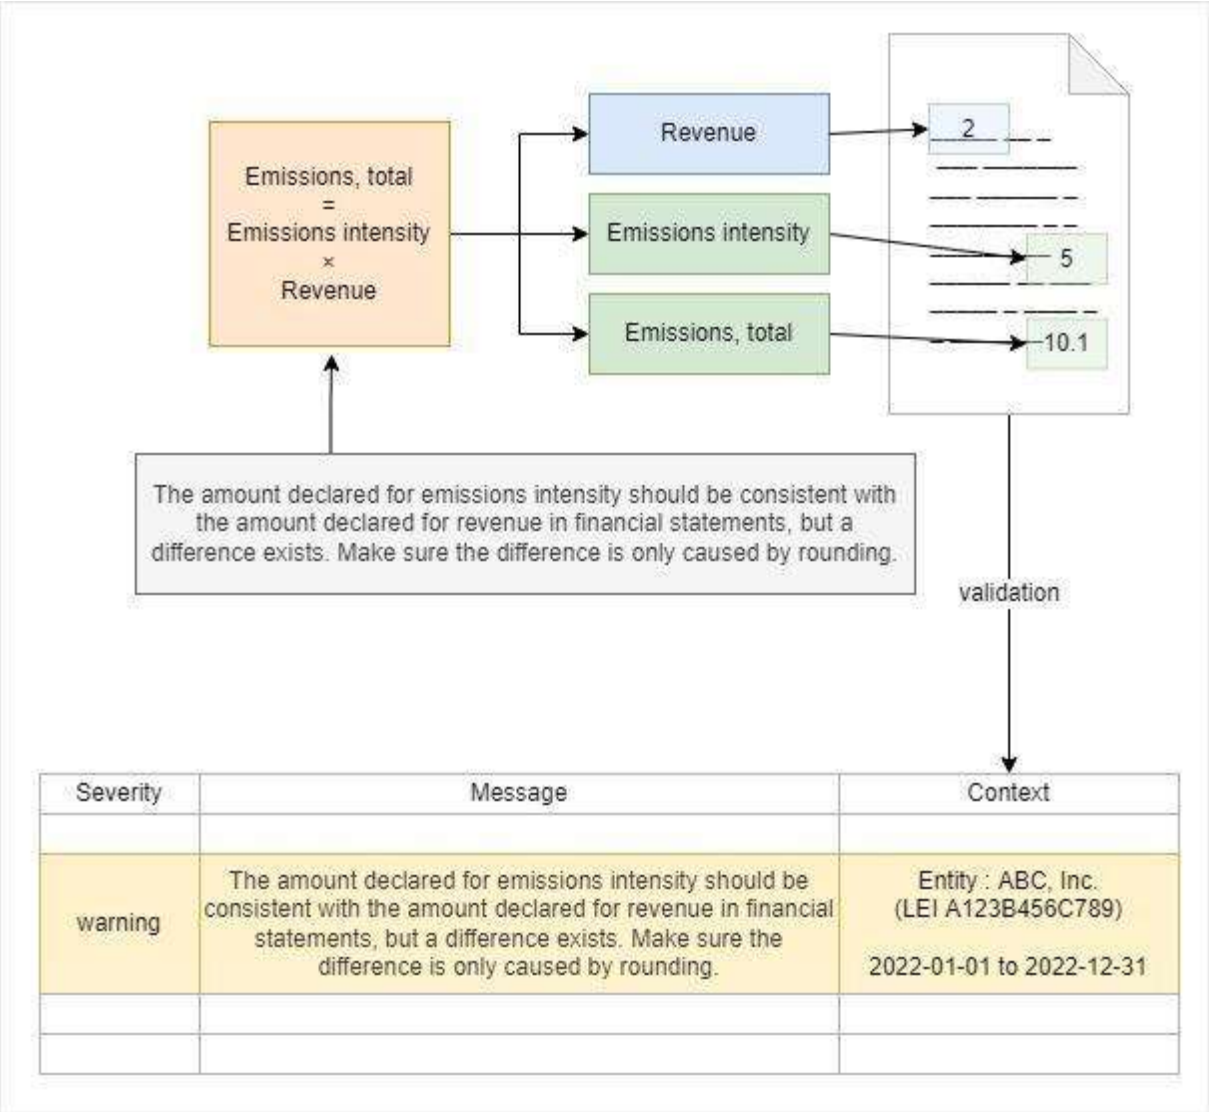

Figure 16: A message is attached to the formula defined in the previous example. When a user employs an XBRL processor to validate a report where all three concepts are present with the same context (same entity, same date or period), the formula will be checked with the corresponding figures, and if it is false, a line will be added to the validation logs for this report. A technical example can be found in Annex I.

Since it requires (in addition to the identification of the involved concepts) the definition of technical formulas, defining this kind of relationship usually requires additional work when compared to the previous solutions.

However, such definitions are the most suitable solutions for enabling automated processing of the documents and greatly improve the value of a taxonomy when they are present.

3. Annex I - Technical examples

3.1. Definition of a new type of relationship

```
<link:arcroleType
    id="wider-narrower"
    cyclesAllowed="undirected"
    arcroleURI="http://www.esma.europa.eu/xbrl/esef/arcrole/wider-narrower"
    >
...
</link:arcroleType>
```

The anchoring relationship in ESEF specifications is technically defined by the above snippet. The "name" that identifies anchoring is the "http://www.esma.europa.eu/xbrl/esef/arcrole/wider-narrower" string. An anchoring relationship between two concepts can be created by simply referring to that string.

3.2. Documenting a new type of relationship

```
<link:arcroleType
    id="wider-narrower"
    cyclesAllowed="undirected"
    arcroleURI="http://www.esma.europa.eu/xbrl/esef/arcrole/wider-narrower"
    >
    <link:definition>
        An anchoring mechanism that links source concepts of wider scope and meaning to target
        concepts with narrower scope and meaning where one end of the link is identified as an
        extension concept and the other end as a base taxonomy concept. This arcrole is intended
        for use on reporting concepts, and is not intended for use between domain members or
        other dimensional components.
    </link:definition>
    <link:usedOn>
        link:definitionArc
    </link:usedOn>
</link:arcroleType>
```

The human-readable definition included in the technical definition for the anchoring relationship.

3.3. Documentation of a specific instance of a relationship

```
<presentationArc
    xlink:type="arc"
    xlink:arcrole="http://www.xbrl.org/2003/arcrole/parent-child"
    xlink:from="esg_DisclosureNetZeroTransitionStrategy"
    xlink:to="fin_UsefulLifePPEMeasuredAsPeriodTime">
    <arcDocumentation xml:lang="en">
        When a class of assets is planned to be replaced in the transition strategy,
        the useful lives of assets of that class should be adjusted accordingly. The
        adjustment should be disclosed separately.
    </arcDocumentation>
    <arcDocumentation xml:lang="ga">
        ...
    </arcDocumentation>
    ...
</presentationArc>
```

Free-form text can be added to the relationship in several languages to help users understand what the relationship conveys and how it can be used.

3.4. Adding a message to an automated check

```
<va:valueAssertion
    xlink:label="formula"
    test="$ghg_total eq ($ghg_intensity * $revenue)"
    ...
    />
<gen:arc
    xlink:arcrole="http://xbrl.org/arcrole/2010/assertion-unsatisfied-message"
    xlink:from="formula"
    xlink:to="message"
    ...
/>
<msg:message
    xlink:label="message"
    xlink:role="http://www.xbrl.org/2010/role/message"
    xlink:type="resource"
    xml:lang="en">
        The amount declared for revenue in financial statements is expected to be consistent with
        the amount declared for GHG emissions and GHG intensity. Check whether the difference is
        due to rounding or due to other causes.
</msg:message>
```

If the formula that corresponds to the assertion is false, the message will be displayed to the user to inform them about the reason for the message, recommend further manual checks and, if the issue is confirmed, suggest possible fixes to it.