



NINE

SOME ISSUES THAT ARISE IN FRBR

There are some fundamental problem areas in FRBR as I read the document. There have been hints some of these in the preceding description of the FRBR entities: marked differences between the underlying structure of the three groups and a lack of cohesion within the groups themselves (e.g., there are no shared qualities among group members). Some of these issues just don't feel quite right; others pose problems in implementation. Among the truly difficult problems are those of inheritance and hierarchy, of the disjoint nature of the Group 1 entities, and how bibliographic aggregates fit (or do not fit) into the FRBR model.

INHERITANCE AND HIERARCHY

It is a common assumption that the four entities in Group 1 represent a hierarchy with inheritance, even among FRBR Study Group members: “Expressions inherit properties from works and manifestations inherit properties from expressions, not vice versa” (O’Neill et al., “Final Report on the Working Group on

Aggregates”). In fact, it might be best if they did, but they don’t, as Renear and Choi demonstrated in their paper “Modeling Our Understanding, Understanding Our Models” (2006). In fact, the E-R modeling technique used in FRBR *cannot* express inheritance. The four entities of Group 1 are linked through relationships in a kind of bibliographic daisy chain; therefore, for example, a FRBR manifestation has access to the attributes of a FRBR work through the link to the FRBR expression. But as described, a manifestation cannot “have” an author or subjects because those are only attributes of a work.


In addition, the inclusion of many-to-many relationships between expression and manifestation means that the structure of Group 1 in data modeling is by definition a network structure. However, FRBR treats the structure as “a continuous chain,” as described in section 5.2.1:

It should be noted that although the relationships between *work*, *expression*, *manifestation*, and *item* are depicted in the entity-relationship diagram in a segmented way, they operate logically as a continuous chain. That is to say that the relationship from *work* to *expression* carries through to the relationship from *expression* to *manifestation*, and those two relationships subsequently carry through to the relationship from *manifestation* to *item*. Thus when a relationship is made between an *expression* and a *manifestation* that embodies the *expression*, the *manifestation* is at the same time logically linked to the *work* that is realized through the *expression*, given that the *expression* has been linked to the *work* it realizes. (FRBR Final Report, 58–59)

It’s easy to understand why people assume that there is inheritance. The statement that relationships “carry through” does sound like an inheritance model. Ignoring the many-to-many relationships (which are illustrated in the FRBR diagrams only through the use of multiple arrow-heads, and are easy to overlook in that form), you can produce an instance of FRBR that looks like that shown in figure 9.1.

FIGURE 9.1

A hierarchical model




This is visually the same as a taxonomy. In a taxonomy, the lower levels do inherit from the upper levels, because each lower level is a “type of” the defined upper level. For example, a dog is a type of mammal, so dog inherits qualities from the super-class mammal. However, an expression is not a type of work, and a manifestation is not a type of expression. The structure of FRBR does not represent a taxonomy and the entities are not in a class/sub-class relationship to each other.

The shape of the diagram is deceptive. The same diagram can, for example, represent an organization chart, as in figure 9.2. An organization chart does not have inheritance—lower-level employees are not “types of” the levels above them, and do not inherit the tasks or salary of their superiors. Each employee is a separate entity with a defined relationship to the employees “above” and “below.” Like the FRBR entities, these relationships are not of the “type of x” or “is an x” nature; an office worker is not a type of manager even though linked to the manager in the organization chart.

FIGURE 9.2

An organization chart



The four levels of Group 1 are not hierarchical in the common sense of that word. They do have a defined order, and the order goes from concrete (item) to increasingly abstract (manifestation, expression, work). But the entities as diagrammed are independent of each other, in the same way that a person entity is independent of a work entity, with which it can have certain relationships.


In the formal modeling of FRBR as a Semantic Web vocabulary, the Group 1 entities are defined as “disjoint.” What disjoint means in that technology is that the two entities cannot share any attributes or relationships. It is for this

reason that each FRBR entity must have its own distinct set of attributes, like “work title” for the work entity and “title proper” for the manifestation entity, even though in some sense “title is a title is a title.” This disjointness also means, though, that the entities cannot inherit attributes from each other because an expression or manifestation is not allowed to be described with the attributes of the work, and vice versa.

More importantly, Group 1 is not a tree structure like a taxonomy or an organization chart. There are many-to-many relationships between the Group 1 entities, which means that Group 1 is a network of relationships. A more accurate diagram would look like figure 9.3.

FIGURE 9.3

FRBR Group 1 with many-to-many illustrated



Note that the diagram in figure 9.3 only includes the primary relationships of Group 1. You might say that they are the glue that holds Group 1 together; they are structural in nature. The bibliographic relationships, like “adaptation of,” “translation of,” or “supplements,” are also important, and presumably are part of the fulfillment of user needs in the library catalog. Adding only a few of these relationships provides a much more complex and graph-like image.

In data design, a network structure is quite common, and the many-to-many relationship is hardly unusual. It does, however, mean that the implementation of the conceptual model as a logical model must do some extra work to make the many-to-many relationships function as desired.

GROUP 1

ENTITIES AS “DISJOINT”

As mentioned above, the Group 1 entities defined in FRBR do not share any attributes. The attributes are each exclusively related to a single entity, and in the case of attributes that could be used for more than one FRBR entity, the attribute has been defined to describe only one entity, as in work identifier, expression identifier, manifestation identifier, or item identifier. Thus, each entity is a walled garden, distinct from the other Group 1 entities.

Yet this separation of entities is contradicted in the FRBR Final Report’s textual description of the entities. In the text of the document, an expression is a work as expressed, and therefore includes the content and meaning of the work. A manifestation is a real world realization of the expressed work, and therefore includes the content and meaning of the work as well as the expression of the work. This alone disproves the theory that the entities are entirely separate from each other. However, as we’ve seen, the E-R diagram and the text of the FRBR Final Report are not perfectly aligned. The E-R diagram shows a bibliographic description that has been broken into four dependent parts that must be reassembled to create a whole. One way to interpret this is that the text describes the concepts, but the E-R diagrams represent a further analysis that pertains to FRBR as a data model. However, this does not resolve the differences between those two views.

One of the first arguments against disjointness of the FRBR entities as a choice for bibliographic data is that, as the FRBR Final Report itself allows, not all bibliographic communities are expected to have the same precise definitions of the FRBR entities:

Because the notion of a *work* is abstract, it is difficult to define precise boundaries for the entity. The concept of what constitutes a *work* and where the line of demarcation lies between one *work* and another may in fact be viewed differently from one culture to another. Consequently the bibliographic conventions established by various cultures or national groups may differ in terms of the criteria they use for determining the boundaries between one *work* and another. (FRBR Final Report, 16)

Clearly, if the boundaries can be different, then the assignment of attributes to entities must be able to differ. However, nothing in the description of the attributes of each entity in FRBR takes this into account. Somehow, moving

from page 16 of the FRBR Final Report, which describes the work, to page 32, where the attributes are defined, all notion of possible variation is abandoned. As Robert Maxwell says in *FRBR: A Guide for the Perplexed*: “Given that FRBR emphasizes the fluidity of the concept behind the entity work, it is somewhat surprising that the document immediately gets down to the business of defining exactly where that line or boundary is (FRBR 3.2.1, 16–17)” (Maxwell 2008). The FRBR Final Report’s description of the work and its list of attributes as a single, unvarying view means that any deviation from that view is a deviation from the FRBR model itself. The declaration of disjointness between entities enforces the view that there is only one “right way” to model the bibliographic universe.

One must be cautious in defining entities as disjoint because disjointness results in considerable rigidity. That rigidity makes it difficult to change (“women cannot be priests” or “two men cannot marry”), and it also makes it difficult for data from different communities that may not have the same restrictions to interoperate. Disjointness means that all users of FRBR must have the exact same definitions of the entities of Group 1 (which is most likely the purpose of the strict definition). If a specialist community determines that a physical attribute like color defines a new expression instead of being an attribute of manifestation, as defined in FRBR, they will be violating the FRBR model. This rigidity, as well as the fact that FRBR is considered “true” in its current form, means that any application must be either FRBR or not-FRBR, thus splitting the bibliographic world into noncompatible factions.

AGGREGATES

FRBR depicts Group 1 as a network in which each entity is in a specific relationship with one other of the Group 1 entities. They therefore form a continuous line from work to item, and item back to work. Difficulties arise, however, in the relationship between expression and manifestation. What often interferes here is the complication that publishers and producers of creative works add to the picture. Although it may be quite accurate to say that an expression is manifested in a physical product, it is something else to say that the physical product is solely the manifestation of the expression. The reason is that the physical, publisher-produced package nearly always has content and qualities that are in addition to the expression. From the design of the package to liner notes, creator biographies and prefatory material, the expression is packaged as a manifestation with content provided by the publisher or producer.

This additional content is described in detail in Gerard Genette's *Paratexts: Thresholds of Interpretation* (1997). Paratexts are all of the contents of a publication that are not the primary text. Some are provided by the creator of the primary content, but others are added by the product creator, the publishing house. Paratexts are sometimes recorded in library catalog data, depending on their perceived importance (e.g., a preface by a well-known scholar), but not always. In cataloging, one of the more important paratexts supplied by the publisher is the title page, which is considered the primary source of information for the catalog record. Yet the publisher is recorded only as an element of that page, and is not credited with any specific creative role for included content except in some rare books or archival publications.

Not only do publishers include paratexts in publications, they also create aggregations of expressions of works, and not always with the participation of the original creator. Aggregates can contain works of a single author or those of multiple authors. Because of this, the difference between multiple expressions in a single manifestation and the presence of an expression along with paratexts is not nearly as clear as it could be. A version of Dante's *Inferno*, with a translation into English, a lengthy introduction, and copious explanatory footnotes is hardly a mere manifestation of the expression of that work, yet there is no physical separation between the works of Dante and his commenter that could easily lead to the definition of two separate works. Some collections of seemingly separate works may in fact have the intellectual characteristics of a single work, as in publications of a professor's lecture series. On the other hand, a recording of two or three musical pieces by a single composer may be little more than the publisher's determination of an optimum number of minutes given the physical medium of publication. A publication of separately authored papers may in some cases cover a wide range of topics, making "workness" nearly impossible to define for the whole. But in other cases the separate essays may have a coherence that is hardly different from a group of chapters written by a single author.

The question that one asks at this point is whether the FRBR manifestation refers to the expression that has been published, or whether it refers to the entire published package. The FRBR Final Report allows for either the treatment of an aggregate as a work in itself, or as a whole/part relationship, such as with individual chapters in a book or articles in a journal. The whole/part relationship has two types: the dependent part, which does not stand alone or is integral with the whole (such as illustrations that accompany and support a text) or the independent part, such as an article in a journal. The FRBR Final Report concludes

that although “there often will be no reason for a dependent part of a work to be separately identified or described . . . independent parts of a work are more apt to be identified and described in their own right.” (FRBR Final Report, 68)

If a FRBR manifestation does not refer to the entire publication, then it isn’t clear how the fullness of the publication is covered in that model. It also isn’t clear what the difference is between “is manifested as” and “is part of.” FRBR includes part/whole relationships in its model. This, however, is where things get complex.

Because the bibliographic description is separated into the four Group 1 entities, WEMI, it is implicit that all creative output has this structure. Thus, an introductory essay in a published book also should have the four-part nature of work, expression, manifestation, and item. One could take that further and consider each paratext and each design element of the finished publication to have “WEMI-ness.” A book cover design surely would be considered a manifested expression of a work to the extent that any other visual resource is, and the same could be said for photographs or illustrations within the book.

In this way, the layered bibliographic description that is WEMI complicates the part/whole relationship, because one has to decide where along the cascade of abstractions of WEMI the part/whole relationship exists. Is this part of a work? Part of an expression? Part of a manifestation? These are not easy questions, as the FRBR Study Group discovered.

The FRBR Study Group created a subgroup to study the complications brought on by aggregate publications. The Working Group on Aggregates narrowed down the potential solutions to aggregations to two competing views: aggregates as manifestations, which was authored by Ed O’Neil and Maja Žumer, and aggregates as works, by Barbara Tillett.

The argument by O’Neill and Žumer (undated) on aggregates as manifestations makes use of the allowed many-to-many relationship between expressions and manifestations in FRBR. Because any manifestation can manifest multiple expressions, an aggregate is a “combination of expressions” in a single manifestation. They also note that there are different types of aggregates: collections (which are aggregates of independent works), augmentations (which contain supplemental material around a primary work or works, like the aforementioned translation and commentary on Dante), and parallel aggregates, which are often texts that appear in more than one language in the same publication. An aggregate as manifestation has separate expressions of works that are manifested in a single manifestation.


Tillett’s paper (undated) takes the view that an aggregate is a work in itself, and that it has a whole/part relationship with any works that are included in the

aggregated work. Therefore, a “work can be a work of works,” which could be shown with a recursion symbol on the box for the work entity. She states that “The recursive symbol was omitted from the final text, but because this is causing confusion, I feel it should be reintroduced.”

The final report of the Working Group on Aggregates (O’Neill et al. 2011) concludes with an aggregation of its own, a combination of features from the two viewpoints described above. The group concluded that the FRBR approach to aggregates would use the many-to-many relationship between expressions and manifestations. However, because aggregates have some properties of their own (such as a creator of the aggregate), aggregates may also be seen as separate works. Any one instance can have multiple works as well as a single aggregating work, as shown in figure 9.4.

FIGURE 9.4

The solution to FRBR aggregates



The recursive element suggested by Tillett was not included in the FRBR Final Report. The Final Report also did not resolve the difference between a manifestation that manifests multiple expressions and a manifestation with parts. This seems to be a significant weakness in FRBR that needs resolution but may require more experience with the model.

Heidrun Wiesenmüller (2011) of Stuttgart Media University studied the FRBR aggregate model and questioned the logic of the sub-group’s conclusion. In her view, the treatment of both individual works and aggregating works as FRBR works at the same level did not accurately reflect the nature of those works. Her suggestion was that there should be two levels of work, one for the aggregate and one for the individual works. This, however, produces an even more complex

set of relationships, at least in terms of how they are diagrammed. However, no model that includes aggregations will be simple.

One solution not considered by the Study Group on Aggregates would be to consider that all manifestations are, essentially, aggregates, comprising at least a primary expression of a work and the paratexts accompany that content. Taking this view would lead to a consistent treat of all manifestations since a publication always goes beyond the mere manifestation of an expression. Even the choice of page design, the use of covers, employing a font and pagination, perhaps adding a table of contents and index—all of these make the manifestation more than a manifestation of an expression. The alignment of FRBR with the CIDOC resource model, FRBRoo, emphasizes the publication as an event in the chain from creator to audience, and therefore injects the act of publication as external to the creative workflow. However, FRBRoo very definitely defines aggregates as works and does not have an equivalent to FRBR's aggregation as manifestation. In the view of CIDOC's resource model, the FRBR manifestation is a compound of content and carrier that needs to be broken apart. The work on FRBRoo may be the best analysis of why the FRBR Study Group was unable to find a solid solution to the problem of aggregate publications.

Before concluding, it is important to acknowledge that decisions made for the purposes of library cataloging are not intended to solve philosophical questions about the nature of reality. The catalog has practical goals relating to inventorying the library collection and serving those who seek materials in the library. For that reason it is legitimate for library cataloging rules to make decisions that serve the library's need, even if one could argue that they are not somehow philosophically sound. When the members of the Working Group on Aggregates, a subgroup of the FRBR Study Group, were tasked with formulating a solution to the problem of how to deal with the part/whole and aggregates issue, their focus was rightly on solving this problem for library cataloging. Whether they did so in the best way possible is something I do not feel qualified to judge. It is unfortunate, however, that this problem did not lead to a reevaluation of any of the assumptions on the nature of the FRBR entities, in particular on the relationships between the expression and the manifestation. One possible adjustment could have been to redefine the manifestation separate from the package or container. The manifestation of the expression would then be something contained within the publication. This would have separated the publisher's package from the creative work it encompasses and made it possible to describe the two independently. Instead, the manifestation arises solely from the expression and has no individuality of its

own, leaving all of the effort of publishers to be treated as a kind of afterthought. I doubt if publishers would find that characterization flattering.

The study of aggregates is evidence of a significant difference between the cataloging view of library metadata and the view that a data designer might take. It is quite awkward (and not recommended) for data designs to have exceptions or “either/or” situations. The study of aggregates as manifestations cites statistics from OCLC showing that aggregates are only a small portion of the bibliographic universe. However, in data modeling, it does not matter if 2 percent or 98 percent of your instances will exhibit the characteristic in question; the model must solve the problem in a way that is valid for all of your data. For example, in a situation where you can have many-to-many relationships in your data, the data must be modeled that way even though a majority of your instances might be one-to-one. If there is the possibility to further develop the concepts first presented in FRBR, the treatment of aggregates is an obvious area that will need additional study.

