

TOWARDS A GENERAL CONCEPTUAL MODEL FOR BIBLIOGRAPHIC AGGREGATES:
FOUR CASE STUDIES FROM OUR BIBLIOGRAPHIC STANDARDS

BY

JACOB JETT

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Library and Information Science
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2019

Urbana, Illinois

Doctoral Committee:

Professor Allen H. Renear, Chair
Research Associate Professor David Dubin, Director of Research
Assistant Professor Karen M. Wickett
Professor Timothy W. Cole
Professor J. Stephen Downie

Abstract

Bibliographic aggregates such as anthologies, collections, journal issues, and media series are increasingly becoming the focus of bibliographic description. Bibliographic description, typically in the form of bibliographic metadata records, forms the cornerstone of information retrieval systems. Library users rely on bibliographic metadata records to find, identify, select, and obtain information resources of interest to them. In turn, library catalogers and metadata librarians rely on high-level conceptual standards to inform them regarding what metadata is central to each kind of bibliographic entity's description, including bibliographic aggregates like those mentioned above. However, not all of our high-level conceptual standards agree on how bibliographic aggregates should be modeled and what metadata is significant enough to be recorded in their bibliographic descriptions.

This dissertation analyzes conceptual models for bibliographic aggregates central to metadata descriptions for bibliographic description in library settings. More specifically, this dissertation focuses on the variations in conceptual models for bibliographic aggregates in four high-level library-centric conceptual models: Dublin Core Collections Application Profile (DC-CAP), Functional Requirements for Bibliographic Records (FRBR), Object-Oriented FRBR, and Library Reference Model (LRM).

The first three standards take an approach to modeling bibliographic aggregates that is based on concepts of parts and wholes. The more recent LRM standard takes a different approach by closely linking its bibliographic aggregate model to its central model for bibliographic entities in general—Work-Expression-Manifestation-Item (WEMI). This

dissertation makes a conceptual analysis of all four approaches in order to compare, contrast, and reconcile their conceptual models for bibliographic aggregates.

Acknowledgements

I would like to thank the members of my committee, especially Professors David Dubin and Karen Wickett, whose thoughtful feedback made this dissertation possible to write. I would also like to thank the members of the iSchool's Conceptual Foundations Group, whose feedback and open discussion of conceptual issues helped hone the analyses and arguments set out in this document. Finally, I would like to thank all of my friends who supported me through this journey.

Contents

Overview & Summary	1
1. Introduction & Method	6
1.1. Introduction.....	6
1.1.1. Bibliographic Metadata Standards	6
1.1.2. Bibliographic Aggregates	7
1.1.3. Research Problem	10
1.2. Method.....	13
1.2.1. Conceptual Analysis as Method.....	13
1.2.2. Conceptual Analysis in the Library and Information Sciences	13
1.2.3. Characterizing Conceptual Analysis.....	15
1.3. Dissertation Layout.....	18
2. Bibliographic Aggregates in Context.....	19
2.1. Chapter Overview	19
2.2. Bibliographic Metadata & the Library Catalog	19
2.3. Bibliographic Entities & Their Models.....	26
2.3.1. Competing and Complementary Models for Bibliographic Entities	26
2.3.2. Moving Beyond WEMI	38
2.4. Bibliographic Aggregates	40
2.4.1. Bibliographic Aggregates as Bibliographic Entities.....	40
2.4.2. The Problem of Paratext	46
2.4.3. An Initial Analysis of Bibliographic Aggregates.....	48
2.5. Chapter Summary	56
3. Part/Whole Conceptual Frameworks	58
3.1. Chapter Overview	58
3.2. Parts, Wholes, and Bibliographic Aggregates	58
3.3. Winston, Chaffin, and Herrmann’s Meronymous Relationships.....	60
3.3.1. Establishing an Initial Formal Framework	60
3.3.2. Initial Distinctions.....	66
3.3.3. General Descriptions of Metaproperties for Part/Whole Relationships	72
3.3.4. Candidate Part/Whole Relationships for Bibliographic Aggregates	76
3.3.5. Further Extensions to EM	79

3.3.6. Formal Analysis of Containment	91
3.3.7. Formal Analysis of Componenthood	95
3.3.8. Formal Analysis of the Membership Relationship	99
3.4. Chapter Summary	104
4. Analysis of Mereological Aggregate Models	107
4.1. Chapter Overview	107
4.2. Mereological Aggregate Models.....	107
4.3. Aggregates in the Dublin Core Collection Application Profile.....	112
4.3.1. DC-CAP Aggregate Model.....	112
4.3.2. Initial Analysis of “ <i>is-gathered-into</i> ” Relationship.....	115
4.3.3. Homogeneity Among Collection Members	121
4.3.4. Transitivity with Regard to Collection Membership.	127
4.3.5. Separable Collection Membership.....	130
4.3.6. Discrete and Dense Collection Membership.....	133
4.3.7. Overlapping Collection Members.....	140
4.3.8. Immediate Predecessors and Immediate Successors	141
4.3.9. Initial Findings	145
4.4. Aggregates in the Functional Requirements for Bibliographic Entities	147
4.4.1. FRBR Aggregate Model	147
4.4.2. Formal Account of FRBR WEMI Model	148
4.4.3. Formal Analysis of FRBR Aggregates	153
4.5. Aggregates in Object-Oriented FRBR [FRBR _{OO}]	162
4.5.1. FRBR _{OO} Aggregate Model	162
4.5.2. Formal Analysis	164
4.5.3. Initial Findings	180
4.6. Findings on Mereological Aggregate Models.....	182
4.7. Chapter Summary	184
5. Analysis of the Content-Artifact Aggregate Model.....	186
5.1. Chapter Overview	186
5.2. IFLA’s Library Reference Model	187
5.2.1. The Content-Artifact Aggregate Model.....	187
5.2.2. Bibliographic Entities and FRBR	188

5.2.3. FRBR, LRM, and Bibliographic Aggregates.....	190
5.2.4. Initial Analysis	194
5.3. Formal Analysis of the LRM WEMI Model—WEMI2	195
5.4. Discussion of the LRM Aggregate Model.....	209
5.5. The WEMI2 Content-Artifact Model	213
5.5.1. Developing an Alternate Illustration.....	213
5.5.2. Discussion of the “conveys” Relationship.....	214
5.5.3. Discussion of the “appears-in” Relationship.....	218
5.5.4. Discussion of the “carries” Relationship	219
5.5.5. Discussion of “Manifestations” and Their Particulars	220
5.5.6. Simplified WEMI2 Illustration.....	224
5.6. Initial Findings.....	225
5.6.1. Trivial Implications.....	225
5.6.2. Non-trivial Implications.....	227
5.7. LRM’s Other Aggregate Model.....	232
5.8. Preliminary Conclusions about the Content-Artifact Aggregate Model.....	234
5.9. Chapter Summary	238
6. Findings & Next Steps	240
6.1. Reprise of Research Goals.....	240
6.2. Summary of Findings.....	241
6.2.1. Finding 1 – Comparability of <i>Mereological Aggregate Models</i>	241
6.2.2. Finding 2 – Gaps in FRBR ₀₀ ’s Aggregate Model.....	243
6.2.3. Finding 3 – Failings in the <i>Content-Artifact Aggregate Model</i>	243
6.2.4. Finding 4 – The WEMI2 <i>manifestation</i> Entity is Semantically Overloaded.....	244
6.3. Excising WEMI2 <i>Manifestation</i> from the LRM Standard.....	245
6.4. The Paratext Issue Revisited.....	249
6.5. Avenues for Future Research.....	254
6.6. Conclusion	257
References.....	260
Appendix A: Discussion of Functionality.....	275

Overview & Summary

This dissertation analyzes conceptual models for bibliographic aggregates such as anthologies, collections, journal issues, and series. These conceptual models are central to metadata descriptions for bibliographic control in library settings.

Models for bibliographic entities are intended to inform the shape of metadata schemas supporting their management and access. Well-known examples of these models that this dissertation examines include the *Functional Requirements for Bibliographic Records* (FRBR [IFLA 1998]), the *Dublin Core Collection Application Profile* (DC-CAP [DCCDTG 2007]) and CIDOC's *Object-Oriented FRBR* (FRBR_{OO} [Bekiari et al. 2016]).¹ These models all take a *mereological* approach to the description of bibliographic aggregates; thereby, we will group them together by calling them *mereological aggregate models*. That is, they use the traditional terminology of parts and wholes, and the relationships of parts and wholes, as the basis for conceptualizing bibliographic aggregation.

IFLA's recently issued bibliographic conceptual model—Library Reference Model (LRM [Riva, Le Bœf, and Žumer 2017]) abandons this mereological approach. LRM provides an alternate conceptual account based on the concept of *embodiment*, as that notion has been defined and used in various IFLA documents over the years (IFLA 1998; FRBR-WGA 2011; Žumer and O'Neill 2012; Riva, Le Bœf, and Žumer 2017). We will call this non-mereological approach the *content-artifact aggregate model*, as it is primarily concerned with leveraging particular aspects of

¹ The referenced version of FRBR_{OO} is the most recent at the time of writing. The 1st version of the FRBR_{OO} model emerged in 2006.

intellectual abstraction along a continuum that spans from content (informative, artistic, etc.) to the real-world (concrete) artifacts upon which it is stored and through which it is transmitted. Specifically, the *content-artifact aggregate model* defines bibliographic aggregates as *manifestations* in the Work-Expression-Manifestation-Item (WEMI) model,² and implies that they do not have parts in a mereological sense (Riva, Le Bœf, and Žumer 2017, pp 93-4). The analysis presented here compares these two approaches.

To demonstrate where core concepts of the mereological models (*parthood* relationships such as *isGatheredInto*, *[bibliographic]PartOf*, etc.) fall among the family of standard mereological relationships, we first examine a set of 20 metaproperties. This set of metaproperties supports a clarification of the ontological status of the aggregating relationships used in the mereological models. We then show that all of the aggregating relationships in these standards are specialized versions of Winston, Chaffin, and Herrmann's *collection-member* relationship (1987) and should therefore not be interpreted as specialized versions of the *[set-]member-of* relationship found in standard mathematical set theory (Fraenkel and Bar-Hillel 1958). This determination is based on the fact that some form of the Homogeneous metaproperty (e.g., WEMI Homogenous, etc.), in addition to the Transitivity metaproperty, is necessarily possessed by these aggregating relationships but such is not the case for the *[set-] member-of* relationship defined in set theory.

The axiomatic approach to defining and analyzing mereological models is familiar and well-developed (Simons 1987, Varzi 1996, and Varzi 2016), but there is no comparably familiar formal approach to *embodiment*. In this dissertation, we develop an account of *embodiment* that helps us

² Also known as the FRBR Group 1 Entities.

understand its role in aggregate models, or as Riva, Le Boëf, and Žumer imply, its role as an *aggregating relationship*. We do this by analyzing the WEMI model set forth in LRM. As this is a very different model from the one set forth in the original FRBR document, we refer to it as WEMI2.

We conduct an analysis of WEMI2's core entities and relationships much as Wickett and Renear (2009) have already done for the original WEMI model. Through this analysis, we show that while LRM presents a much-improved sense of *works*, *expressions*, and *items*, its concept of *manifestation* is semantically overloaded. This is due in part to WEMI2 *manifestations* being defined as curated sets. This leads to the WEMI2 model employing *manifestations* for three distinct use cases:

1. As bibliographic aggregates in which *expressions* “appear.”
2. As sets of WEMI2 *items* related to one another through shared characteristics.
3. As something that reflects the characteristics shared by the *items* that are elements of them, e.g., a production plan or a metadata record.

Additionally, the analysis of these models finds that:

- All of the *mereological aggregate models* examined are closely related to one another through sub-property relationships.
- All of the *mereological aggregate models* are agnostic with respect to specifics of kind of bibliographic object (e.g., collection, anthology, series, etc.), and thereby:

- Provide practitioners with flexibility of choice as all can be expected to perform similarly and,
- Are general enough to record the particulars of the part/whole relationships of the bibliographic object, without impinging upon other metadata facets needed to describe bibliographic aggregates as first-class bibliographic objects.
- The *content-artifact aggregate model* directly limits certain kinds of metadata from being recorded, e.g., topicality.
- The *content-artifact aggregate model* causes some confusing situations for information retrieval systems:
 - Since *exemplifies* is a one-to-many relationship, we become uncertain as to which set a WEMI2 *item* belongs when it is both mass-produced and part of a series.
 - This problem compounds when nesting series, such as trilogies that are part of a larger overarching series, need to be accounted for.
- The *manifestation* entity in WEMI2 is semantically overloaded.
- Ignoring the WEMI2 *manifestation* entity altogether results in some unique benefits:
 - An existing *mereological aggregate model* deployed in LRM can be used as a superior model for *bibliographic aggregates*.
 - With some extra-refinement, the trio of *work-expression-item* entities can be leveraged into a superior item-deduplicating model, as it provides clear links from content to signs conveying content (e.g., a particular version of some text), to objects carrying signs conveying content (e.g., a book on a shelf, a particular copy of a text file, etc.).

While the findings described above ultimately indicate LRM's approach to *bibliographic aggregates* lacking, LRM nonetheless provides many valuable additions to overall descriptive enterprise for bibliographic entities in general. In particular, the WEMI model described by LRM is both different from and much more clearly stated than the one set forth in FRBR. This new WEMI model indicates that progress in the description of both the physical and abstract aspects of bibliographic entities is being made as our high-level conceptual standards continue to evolve. But it also implies that practitioners need to be wary of idiosyncrasies in the standard document's text and be made aware that the ER-diagram previously used to illustrate the WEMI model is not necessarily an accurate representation of the WEMI model set forth by LRM.

1. Introduction & Method

1.1. Introduction

1.1.1. Bibliographic Metadata Standards

Library users rely on bibliographic metadata, typically communicated to them in the form of metadata records, to, as the International Federation of Library Associations (IFLA) tells us in their Functional Requirements for Bibliographic Records (FRBR), find, identify, select, and obtain information resources of particular interest to them (1998, p 8). Library catalogers and other metadata creators rely on bibliographic metadata standards to identify what information particular to an artifact (e.g., a book, etc.) or its content (e.g., a “work”) should be included in a metadata record to best facilitate library users meeting the goals set out by the aforementioned user activities.

Prior to the 20th century, these metadata standards could be relatively simple, since libraries were primarily concerned with the collection and organization of three variations of text media—monographs (i.e., books), serialized publications (e.g., journals, newspapers, etc.), and manuscripts. The 20th century saw an explosion in the kinds of media that libraries collected and organized. Serialized publications were collected and bound into pseudo-monographic formats or were photographed and copied onto storage formats like microfilm and microfiche. Sound recordings, video recordings, and vast collections of digitized materials (ranging from government documents to photograph collections) have all become objects that libraries collect and organize.

The means by which we create, store, and interact with the metadata records representing media objects have also evolved. Data models, document models, and similar information design guides require clear ideas both about what thing the information being stored describes, and what portions of that information are available to which users. The relatively recent advent of the Semantic Web (Burners-Lee, Hendler, and Miller 2001; Shadbolt, Hall, and Burners-Lee 2006) requires even

more stringent descriptive rules to avoid blending different entities together through our metadata descriptions (Jett et al. 2016b). The needs of library users, especially scholarly users, are also becoming increasingly sophisticated. Providing enough metadata to simply meet FRBR's user goals is no longer enough. Increasingly, library users need metadata robust enough to aid them with their analyses (Fenlon et al. 2014).

The standards by which metadata records for these objects are crafted have also evolved throughout the 20th century. Increasingly, divisions become increasingly fine-grained to describe these objects in a variety of ways, in order to better manage the organization of materials and the information they bear.

This dissertation examines one particular distinction made by bibliographic metadata standards—the division between bibliographic entities that are singular things (e.g., a book, a film, a game, etc.) and bibliographic entities that are groups of other things (e.g., a collection of items, an anthology of short stories, a series of publications, etc.). More precisely, the focus of this dissertation is on the latter of these two broad categories—groups of things or *bibliographic aggregates*. The ultimate goals of this research are to assess the manner in which four high-level conceptual standards model *bibliographic aggregates* in order to clarify similarities and differences among them, and to assess the implications of those similarities and differences for bibliographic control (i.e., the creation of metadata) for *bibliographic aggregates*.

1.1.2. Bibliographic Aggregates

Bibliographic aggregates, like digital cultural heritage collections, anthologies, serialized publications, and similar objects, play an important role in many kinds of scholarship. They play

a particularly important role in humanities scholarship. Collections (or corpora as many humanists call them) are so important that there is even a TEI extension designed to capture relevant metadata about them—the Metamodel for Corpus Metadata (MCM).³ We can see there is clear user need for the capture and creation of metadata that focuses on aggregates as distinct bibliographic entities.

Whether they are objects of research themselves, or the objects from which scholars select parts for research, bibliographic aggregates are first-class bibliographic entities. As such, bibliographic aggregates receive special treatment in library standards regarding the kind of metadata to record for their best representation when situated among other kinds of bibliographic entities.

The representation of bibliographic entities is itself a tricky matter. Librarians and other information professionals have for centuries known that there is a distinction between the information content of an artifact and an artifact itself. More contemporary library cataloging scholars such as Verona (1959), Wilson (1968), Lubetzky (1969), Svenonius (2000), and Smiraglia (2001)⁴ have all addressed this other distinction between “text” (Wilson 1968, Smiraglia 2001), “documents” (Svenonius 2000, Smiraglia 2001), “bibliographic units” (Verona 1959), or “books” (Lubetzky 1969) on one hand, and “works” (Wilson 1968, Lubetzky 1969, Svenonius 2000, Smiraglia 2001) or “literary units” (Verona 1959) on the other. Conceptually, this distinction boils down to the difference between the artifacts that store (or transmit) inscriptions and the meanings that human beings ascribe to those inscriptions.

³ [see <https://zenodo.org/record/267999#.XDNsaPIKiM9>]

⁴ To name only a very few of the very many.

One might think that this distinction is orthogonal to our distinction between singular entities and aggregate entities. However, a new standard (Riva, Le Bœf, and Žumer 2017) recently set forth by IFLA, argues that matters of aggregation and singularity directly align with matters of metaphysical distinctions between artifacts and information content.

One group of metadata standards takes a *mereological* approach. That is, they propose models using parts and wholes to describe aggregates. Among these metadata standards are IFLA's *FRBR* (IFLA 1998), the *Dublin Core Collection Application Profile* (DC-CAP [DCCDTG 2007]) and *CIDOC's Object-Oriented FRBR* (FRBR_{OO} [Bekiari et al. 2016]).⁵ For the purposes of the analysis we carry out in this dissertation, we call this group of conceptual models using mereological approaches *mereological aggregate models*.

In contrast, new conceptual standards for metadata records employ a distinction between abstract content and physical artifacts (i.e., “works” and “text”) to describe aggregates. One of these new standards is IFLA's *Library Reference Model* (LRM [Riva, Le Bœf, and Žumer 2017]). LRM's aggregate model is focused on exploiting features of the Work-Expression-Manifestation-Item (WEMI) model first introduced by FRBR. The WEMI model is primarily concerned with illustrating the differences between content and artifacts. We call aggregate models built around this distinction *content-artifact aggregate models*.

Our goal in this dissertation is to compare and contrast these two descriptive agendas. By doing so we aim to assess how comparable the aggregate models used by our four cases—DC-CAP, FRBR,

⁵ The referenced version of FRBR_{OO} is the most recent at the time of writing. The 1st version of the FRBR_{OO} model emerged in 2006.

FRBR_{oo}, and LRM—are to one another. From the beginning, LRM stands out because it uses a different kind of aggregate model than the other three cases, but the implications of using this different model are not yet fully understood. The analysis herein aims to illuminate some of these implications.

1.1.3. Research Problem

The central focus of this dissertation is to clarify the differences between these two models and thereby better identify the features each approach has for the creation and implementation of metadata and the organization of knowledge. A superficial study of the *content-aggregate model* reveals some potentially troubling gaps in features it provides for the creation of high-quality metadata describing bibliographic aggregates (Jett, Fenlon, and Downie 2018). These gaps raise questions regarding how the aggregate model in LRM is to be implemented and how it should be expected to perform.

Lurking unresolved issues with the WEMI model itself may cause obstacles for the *content-aggregate model*. For instance, a key assumption that is often made with regard to the WEMI model is that properties possessed by an entity at one level of intellectual abstraction, e.g., *work*, are inherited by the entity at the next level “down” the model, e.g., an *expression* of a *work*. However, Renear and Choi (2006) argue that nothing in the FRBR document’s text asserts or even implies that an inheritance relationship exists among the WEMI entities. As they point out, when inheritance of attributes and properties from one entity to another is an expected feature of our bibliographic conceptual models, then the relationships accomplishing that task need to be specified in the standards documents. Simply stating that the entities in the models represent sets and sub-sets is not enough as inheritance is not entailed by ordinary set-subset relationships. As

Jett, Fenlon, and Downie (2018) point out, when there is no inheritance from one WEMI entity to the other, then bibliographic aggregates modeled using the *content-artifact model* lack properties, like topicality, that are vital for accomplishing FRBR user tasks.

Further, since the *content-artifact aggregate model* is built around the many-to-many cardinality of the WEMI *embodies* relationship, it raises questions about dependent works (Jett & Dubin 2018) and authorial intentions. Is it possible that FRBR *expressions* may realize more than one FRBR *work*, or that FRBR *items* might exemplify more than one FRBR *manifestation*? How can we know with certainty that authors intend their *works* to first be published and best be understood on their own? What are the implications of choosing to model aggregates as a specific WEMI entity, i.e., as *manifestations*?

Standards like FRBR and LRM are typically produced with some clear pragmatic benefit in mind. In the cases of FRBR and LRM, the goal is to provide some clarity in how various editions (e.g., mass-produced consumer artifacts like large-print books or trade paperback books) relate to one another through shared intellectual content. These standards thereby inform how metadata records are to be crafted and what representational roles they play in end user information needs satisfaction workflows.

However, as Dubin, Senseney, & Jett (2013) point out, these standards also provide detailed domain models that represent a particular joint understanding of an abstract phenomenon. In this

case the aggregate models used in FRBR and LRM⁶ represent particular views of *bibliographic aggregates*. From a researcher's point of view, the dichotomy of these disparate views is of interest because each model has distinct implications for how well metadata records crafted according to them can be expected to satisfy end user information needs. The implications of each model can indicate gaps in the fulfillment of end user information needs.

This dissertation uses formal methods to compare these two analytical approaches to *bibliographic aggregates*. Through our formal analysis of these two approaches we develop answers to the following research questions:

1. How comparable are the *mereological aggregate models* used by the standards in our case studies?
 - a. The texts of these standards often employ terms like, “container,” “component,” “member,” “part,” etc.; as these terms all refer to distinct part/whole relationships, are we to take them at face value or are they being employed synonymously to indicate one specific part/whole relationship?
 - b. If they are being used to indicate one specific part/whole relationship, which part/whole relationship is it?
2. How does the *content-artifact aggregate model* used in the LRM case differ from the *mereological aggregate models* employed by the other cases?

⁶ We should note that not all of the potential models that could fall into the first category are analyzed. For instance, since it is not primarily intended as library-centric standard, NISO's Z39.29 (Bibliographic References) standard is not examined. The still under-development Bibframe 2.0 standard is also not addressed by this dissertation, but some preliminary consideration of it appears in the further work section of the conclusion.

1.2. Method

1.2.1. Conceptual Analysis as Method

This dissertation examines both mereological aggregate models and content-artifact aggregate models through the critical lens of *conceptual analysis*, (Glock 2008, Beaney 2016). This is appropriate as these models *are* conceptual analyses, although not clearly documented as such.

Our objective is to better understand these models, to determine their differences and relative advantages, and to discover whether they succeed in modeling *bibliographic aggregates*—and if they fail to model *bibliographic aggregates*, the method of conceptual analysis will help us understand exactly why they fail and perhaps how that failure might be remedied.

1.2.2. Conceptual Analysis in the Library and Information Sciences

Discussion and analysis of core concepts is routine in LIS research. Certainly, much of the scholarship of influential LIS researchers such as Shera (1966) and Buckland (1991, 1997) is analyzing important LIS concepts. In the conceptual space of interest for this dissertation, we find important works that are quite clearly engaged in the analysis of concepts, including the aforementioned Verona (1959), Wilson (1968), Lubetzky (1969), Svenonius (2000), and Smiraglia (2001). Although they provide examples of conceptual analysis, these writers do not necessarily say that is what they are doing, nor do they typically deploy the specialized terminology or devices we might see in philosophical writing—there are exceptions: in Wilson (1968) and Svenonius (2000), the philosophical lineage of their approaches is evident.

Jonathan Furner⁷ has frequently employed conceptual analysis, for instance, in his works discussing *information* (Furner 2004), the place of philosophical methods in LIS (Furner 2010), and the nature of *data* (Furner 2016). Furner offers this description of conceptual analysis:

“Conceptual analysis is a technique that treats concepts as *classes* of objects, events, properties, and relationships. The technique involves precisely defining the meaning of a given concept by identifying and specifying the conditions under which any entity or phenomenon is (or could be) classified under the concept in question. The goal in using conceptual analysis as a method of inquiry into a given field of interest is to improve our understanding of the ways in which particular concepts are (or could be) used for communicating ideas about that field” -- Furner 2004, pp 233-4.

Furner links his application of the conceptual analysis method to the analytic traditions of Frege, Moore, and Russell, but like Glock (2008), he also notes that there is no one “method” that one can point to and definitively claim as **the** conceptual analysis method. Furner also employs conceptual analysis to respond to other authors (e.g., Buckland 1991, 1997; Duranti, Eastwood, and MacNeil 2002; Tournay 2003, etc.) who have provided meditations on LIS-centric conceptual problems (e.g., what is a document? what is information? etc.).

This dissertation makes use of previous conceptual analyses of the FRBR family of conceptual models (Renear & Choi 2006, Wickett & Renear 2009) and builds on previous meditations on

⁷ As well as Furner, a number of other LIS researchers, have been employing conceptual analysis as their primary research method. These include Dave Dubin, Don Fallis, Kay Mathiesen, Allen Renear, Karen Wickett, and Simone Sacchi.

bibliographic aggregates (Wickett, Renear, & Furner 2011, Wickett et al. 2013, and Wickett et al. 2014).

As Sacchi and Wickett (2012) point out, the ultimate goal of the kind of analytic process employed in this dissertation is to identify problems of consistency in conceptual and data models (i.e., points where the descriptive or narrative accounts they provide become confused) and to unpack the “black box” representations of information objects in various models (p. 2).

1.2.3. Characterizing Conceptual Analysis

Although there is no agreed upon formal definition of what conceptual analysis is, or how it should be deployed, we can provide a characterization that extends the account given by Furner.

A conceptual analysis “precisely defin[es] the meaning of a given concept by identifying and specifying the conditions under which any entity or phenomenon is (or could be) classified under the concept in question” – Furner, p 234. We extend this by saying more broadly that the conceptual analysis of a concept space also (i) *specifies relationships between concepts* and (ii) *gives an account of how phenomena of interest can be adequately described using these concepts and relationships*.

Generally, conceptual analyses are developed by identifying important concepts, conjecturing defining conditions for those concepts, and indicating relationships between concepts. In the case of relationships, one might also conjecture cardinality (whether a relationship is one to one, one to many, etc.) or relationship properties (transitive, reflexive, etc.), or classify entities as to ontological type (physical object, abstract object, event, class, etc.). Typically, in practical

applications of conceptual analysis, a number of concepts are being defined together. Clues as to what concepts are involved in a domain and need to be identified and analyzed are often found in the nouns, names, and other referring expressions used by practitioners in that domain.

Conceptual analyses should be clear, precise, and rigorously expressed. Even when presented in natural language, they should be easily translated into first order logic, or some well-defined extension of first order logic. ER diagrams, UML class diagrams, and RDF schemas are all considered useful for ensuring logical precision and clarity. As we will show in the subsequent chapters, our conceptual standards already employ a combination of natural-language definitions and ER diagrams to characterize bibliographic entities and various aspects of them in general. However, these natural-language definitions and their attendant ER-diagrams do not always agree with one another. Throughout this dissertation we also interpret the natural-language definitions using first-order modal predicate logic with identity as a means to more easily analyze their implications.

Difficult or specialized terms or locutions must either be themselves defined or, if primitive, identified as such and kept to a minimum. Indeed, much of Chapter 3 is spent clarifying what high-level conceptual standards mean when they use part/whole terms like container, component, part, member, etc., interchangeably. This analysis is carried out by identifying the specialized properties that each of the relationships suggested by the individual terms possess. These specialized

properties are known as *metaproperties*.⁸ Similarly, much of the analysis in Chapter 5 relies on notions of which terms employed in LRM's natural-language definitions should be taken as primitive and which are defined by them.

Individual analyses can be tested by counterexample, either by producing a case that falls under the concept but lacks a supposedly necessary condition, or by producing a case that meets all the conditions of the analysis but does not fall under the concept. We do this by closely examining the examples of bibliographic aggregates that are given in the four conceptual standards documents being examined and which are commonly occurring in library and similar settings. Examples and counterexamples are both interwoven throughout the text and represent a large portion of the analysis.

Finally, an analysis, or set of analyses, is considered inadequate if it does not provide the concepts and relationships necessary to describe the important features of the phenomena it is intended to be modelling. An important distinction to make here is that the models being examined are themselves prior analyses of particular phenomena. In the cases examined by this dissertation, the particular phenomenon of *bibliographic aggregates* in general is the central focus of both the prior analyses (i.e., the models employed in the standards documents) and our analysis. We have taken great strides to fully unpack terms in our analysis in an effort to avoid semantic overloading. Semantic overloading is symptomatic of inadequate analyses as it indicates insufficient

⁸ Since each *metaproperty* represents a property of a relationship which we can understand to be properties possessed by entities. Throughout the dissertation text though we use the term "relationship" to indicate properties that link two entities or more together and the term "metaproperty" to indicate properties possessed by relationships.

atomization of concepts important to the described phenomena. As such, where semantic overloading of terms is found during an analysis, it is strongly indicative that a model (i.e., the result of a prior conceptual analysis) is inadequate to fully describe the phenomena it sets out to.

1.3. Dissertation Layout

The core content of this dissertation is laid out across four central chapters and is followed by a concluding chapter. The first of these central chapters, Chapter 2, serves as a review of the literature that provides a general account of bibliographic metadata, bibliographic entities, bibliographic entity models, and bibliographic aggregates. A series of conceptual analyses are then carried out in the subsequent three chapters. Chapter 3 analyzes the part/whole terms employed in high-level conceptual standards for bibliographic entities in general by closely examining the varying metaproperties possessed by the part/whole relationship suggested by each part/whole term. Chapter 4 analyzes the three conceptual standards that employ *mereological aggregate models*—DC-CAP, FRBR, and FRBR₀₀, and Chapter 5 analyzes the *content-artifact aggregate model* in the LRM conceptual standard's context. Chapter 6 summarizes the dissertation's findings, considers some potentially beneficial modifications for LRM, and explores avenues for future research relating to the matters discussed in this dissertation.

2. Bibliographic Aggregates in Context

2.1. Chapter Overview

This chapter reviews the existing literature regarding *bibliographic aggregates* in library-centric contexts. We first review the central concepts of bibliographic metadata and the library catalog in order to situate the general context in which metadata describing bibliographic aggregates such as anthologies, digital collections, journal issues, monograph series, etc., is typically found. This is followed by a discussion of the concept of bibliographic entities and the general contention between the conceptual aspects, i.e., content, and their physical aspects, i.e., their nature as concrete artifacts. FRBR's WEMI model is deeply interwoven into this narrative as it represents a particular evolution in approaches to modeling the content-artifact aspects of bibliographic entities. Finally, we review the existing work discussing *bibliographic aggregates*, including the general confounding factors that paratext presents for adequate description through metadata and a previous conceptual analysis of particular aggregate entities—digital collections.

2.2. Bibliographic Metadata & the Library Catalog

As we noted in the introduction, competing accounts for describing *bibliographic aggregates* in general exist in the forms of high-level conceptual models of bibliographic entities. In one account, the one put forth in the *mereological aggregate models*, certain kinds of metadata, such as what parts a whole has or the reason all of the parts have been brought together in the first place, would be suggested as constituting significant general facts about a *bibliographic aggregate* that should be recorded through metadata. The other account, as described by the *content-artifact aggregate model*, suggests such metadata is already accounted for through existing metadata that describes something called a *manifestation*. Before we explore the specifics of these notions, we will first examine how this state of affairs has come into being in the first place.

Bibliographic metadata is one of the cornerstones of libraries and related institutions such as archives, historical societies, and museums. Libraries, in particular, were one of the earliest places in which metadata was a central feature to the proper functioning of the institution as a whole, through the innovation of the library catalog. Through the library catalog, a library's various users could search for information resources that meet some information need that they have. They could also discover whether or not the library has a suitable resource to meet that need, and in the event that library does possess such a resource, then they might obtain it for their own use.

Vernacular English dictionaries (e.g., the Oxford English Dictionary, Webster's Dictionary, etc.) tell us that catalogs themselves are systematically organized lists that are intended to be authoritative (i.e., complete) with regard to their topic. Prior to the late 19th Century, catalogs of all kinds were typically implemented using the print publishing technologies of the day. This was particularly the case for library catalogs prior to the late 19th Century (Panizzi 1841).

Library catalogs became standardized during the 19th century through efforts led by Antonio Panizzi (1841), Melville Dewey (1885), and Charles Cutter (1891). One of the outcomes of this process is that the metadata describing the objects in a library's catalog also became standardized. Another direct outcome of this process was the invention of the card-based catalog (Coyle 2016a), which was translated in machine readable format in the 1960s through the Machine-Readable Cataloging (MARC) document format (Avram 1968).

Great innovations in computing during the 20th Century wrought great changes in how catalogs can be implemented. Digital technologies such as databases are now used to implement catalogs

such as the online public access catalogs (OPACs) used in libraries today. These technologies are fundamentally more flexible than old print technologies, like card catalogs, and do not suffer from constraints such as limited space in which metadata can be recorded and communicated or singular ordering regimes. Hence, OPACs do not represent any singular catalog of a library's collection. Instead an OPAC's user is free to define queries that will build "catalogs" that are best suited to meet their information needs.

Published catalogs and card catalogs, then, are very specialized technologies designed expressly to meet very specific user needs. In a library's case, these needs include identifying what bibliographic entities a library owns and which of those they may access. As Coyle (2016b) points out, even when the publishing industry was young and primarily concerned with the printing and reprinting of relatively few books, this was a difficult task, primarily due to the economics of reprinting old texts, translating them into new languages, or rearranging them into new editions. As the demand for copies increased, publishers moved to meet it. However, each new translation and edition further complicated the bibliographic universe in which the library catalog was expected to operate.

When mass-production began to emerge in the 19th Century, the universe that library catalogs sought to articulate was complicated to such a degree that simply entering factual information regarding an edition or translation's title, author(s), publisher, etc. was no longer enough to aid the user in making sense of the dense and diverse bibliographic universe. By the mid-20th Century, the problem was beginning to come to a head as there were simply too many editions, reprints, and

copies of what appeared to be the same book, journal, manuscript, etc. in library catalogs for users to easily disambiguate one apparently same thing from another.

A method for disambiguating editions, reprints, and copies of bibliographic entities from one another was needed. Verona (1959) proposed one by decomposing the concept of bibliographic entity into a distinction between what she calls “bibliographic units” (i.e., the things on the shelf) and “literary units” (i.e., the content that authors create). As she points out, both of these things must be adequately represented through metadata in order for a library catalog to successfully meet the users’ needs to find:

- “1) the rapid location of a particular book;
- 2) the provision of information concerning all editions, translations etc. of a given work as far as they exist in the library; [and]
- 3) the provision of information concerning all works by a given author as far as they exist in the library” – Verona 1959, p 79.

The operations that Verona noted that the library catalog needed to meet were further codified through the Paris principles (IFLA 1961). However, these principles were not uncontroversial. Ranganathan (1962) in particular thought that they fell short in many areas and so produced a paper discussing the many shortcomings with respect to cataloging practices at the time. It is important to note that we can already see that the term “work,” wholly distinguishable from the term “book,” was used to draw distinctions by all three.

Unfortunately, the library card catalog, like the horse and buggy, has become a technology that is obsolete when compared to modern digital technologies like the database software used to implement today's library OPACs. Metadata too, once the principle domain of libraries worldwide, is now a ubiquitous feature of information technology. As many of today's experts on metadata have noted (Pomerantz 2015; Coyle 2016a, 2016b, Gartner 2016; Riley 2017), metadata has grown to become the foundation of and driving force behind much of today's information technologies. From the internet to the software applications on our smart phones, the world we live in today is completely awash with metadata describing where we are, what we are doing, what things we are interested in purchasing, the things we do purchase, and so on.

It is somewhat surprising then that the bibliographic metadata that is so vital to the proper functioning of libraries is behind the times. However, as Coyle notes regarding the development of the first online public access catalogs (OPAC) in the 1980s:

“We, and by ‘we’ I mean all of us in library technology during this time, created those first systems using the data we had, not the data we would have liked to have. The MARC records that we worked with were in essence the by-product of card production. And now, some thirty-five years later, we are still using much the same data even though the information technology has changed greatly during that time, potentially affording us many opportunities for innovation. Quite possibly the greatest mistake made in the last two to three decades is failing to create a new data

standard that would be more suited to modern technology and less an imitation of the library card in machine-readable form.” – Coyle 2016b, p 51.⁹

The above quotation is from Coyle’s book *FRBR, Before and After: A Look at our Bibliographic Models*. It is a meditation on the state of bibliographic metadata in libraries today. One of the primary points Coyle makes is that, despite efforts to innovate bibliographic metadata through new overarching standards, a great deal of chaos is erupting in communities central to the creation and maintenance of library metadata.

On one hand, new technologies are emerging at an unprecedented rate and are creating new expectations for metadata functionality in consumers. On the other hand, the library community has amassed an enormous collection of metadata that still seems relevant, even though its format is at least more than two generations of technology obsolete.¹⁰ It is not so surprising that possibly divergent approaches to articulating the minimal metadata sufficient for describing bibliographic entities in general are emerging.

⁹ Indeed, a key problem is that catalogs (like directories) are a technology from a bygone era and are completely unsuited to today’s highly interactive (meta)data intensive information technologies, as the failings of Yahoo demonstrated (Sullivan 2014).

¹⁰ Coyle specifically mentions relational databases, XML (and declarative markup languages in general), and object-oriented programming as technologies that libraries have missed. However, through object-oriented markup languages like JSON-LD, which are designed to integrate the object-oriented programming approach with Semantic Web technologies like RDF and OWL, it does not seem that librarians have missed the object-oriented programming boat yet. We might also point out that to some extent Coyle has misunderstood the nature of object-oriented programming, which is an alternative approach to writing computer code, as opposed to functional programming, for instance. While object-oriented programming has some value in the manner in which it handles data, especially in high-performance environments like web-browsers, it does so by removing control of the data from the data creator and giving it to the software developer. This often leads to either oversimplification of the data in order to make it easier to move from one application to another or over-specification of the data to make frequently used parts of the data (e.g., DOI identifiers, ISBN identifiers, etc.) easier for the developer to type when writing code.

Compounding this problem is the continuing focus on the metadata record, a specially-structured document designed to communicate assertions about an information resource to an end user. The metadata contained in a record has a number of potential uses for the end user ranging from the descriptive to the administrative and beyond. Just as how contemporary databases may be used to create any number of “catalogs” suitable for an end user’s information purposes, so too do they provide great flexibility regarding how various metadata assertions can be combined to form documents like metadata records. Nevertheless, the concept of the *metadata record* as the canonical holder of the metadata persists in library cataloging and metadata circles.

The needs of users are also continuing to evolve. In 2000, digital humanist John Unsworth noted that scholars’ research cycles employ seven primitive information manipulating activities: *discovering, annotating, comparing, referring, sampling, illustrating, and representing*. Through linking, social widgets like Discus, social media applications like Twitter, free online tools like Google sheets, etc., virtually every user of the Web can successfully carry out many of these activities with almost any information resource they can find online. Not so with libraries and library catalogs, as the catalog itself is only capable of empowering users to complete just the first of the seven primitive tasks that Unsworth lists.

At least, this is the case in traditional libraries. In digital libraries, more and more services are being added that allow library users to accomplish tasks like *annotating, comparing, referring, sampling, illustrating, and representing*. Digital libraries have been able to accomplish this primarily by abandoning the traditional library catalog and the notion of metadata records and embracing flexible models for representing and interacting with the bibliographic entities

contained in the digital library's collection. The goal is to provide the users with sufficient metadata (not metadata records) to accomplish these user tasks. Two key technologies make these innovations possible—models and standardized vocabularies (or ontologies in today's semantically-sensitive linked data environments).

2.3. Bibliographic Entities & Their Models

2.3.1. Competing and Complementary Models for Bibliographic Entities

A core concern for the authors and maintainers of library metadata is what metadata is needed for.

In the era before the computer revolution of the 1950s and 60s, this was a fairly constrained space, consisting primarily of books, serialized publications (like journals and newspapers), and manuscripts. Since then, the numbers and kinds of information-bearing media have exploded, so much so that it is easier to refer to them more generally with the term *bibliographic entity*. With this term, we might be indicating a film, a play script, a novel, or some other media type.

Things have only become more complicated for libraries and similar institutions since the late 20th Century. Now there are electronic editions of books that need to be accounted for by the library catalog, and beyond those, libraries now collect a larger variety of media types than ever before, as the practice of maintaining large collections of films, graphic novels, and even board games and more esoteric media is becoming more and more common place.

Additionally, by the late 20th Century, databases and similar information technologies were becoming the norm in all sectors that employed metadata as a fundamental part of their business models (e.g., in banking, market exchanges, government agencies, and of course, in libraries, among other enterprises). Databases rely on a series of interlocking models to achieve full

functionality, and these are: *conceptual models* (sometimes called domain models), *query models* (sometimes called operational models), *data models*, and *document models*. And in this environment, three additional innovations to Verona, Wilson, and Lubetzky's conceptual analyses occurred.

As we mentioned in the previous section, a distinction between “books” on one hand and “works” on the other was already being drawn. Verona tried to clarify these notions by renaming them “bibliographic units” on one hand and “literary units” on the other. Similarly, Wilson (1968) and Lubetzky (1969) drew distinctions between “text” (Lubetzky actually still uses the existing term “books”) and “work” to meet the same user needs that Verona listed. Despite these meditations, a true conceptual model for bibliographic entities eluded the library sector until Barbara Tillett closely inspected the kinds of information library catalog records were trying to communicate (Tillett 1987, 1991a, 1991b, 1992a, 1992b). In her own words, the goal of her study was to “provide the groundwork for understanding the conceptual structure of the ideal library catalog in terms of bibliographic relationships” – p 1.

Bibliographic relationships, specifically, became the cornerstone of her analysis and of them she says, “A bibliographic relationship exists when we associate two or more bibliographic items or works” – p 1. In the pages of her dissertation, much of which was published in LIS-centric journals in 1991 and 1992, Tillett demonstrates that catalogs were already recording much finer-grained distinctions than just differences between “works” and “items,” as the Paris Principles called them.

In turn, both Tillett and the chair of her dissertation committee, Elaine Svenonius, participated in a series of meetings whose foundations stretch back as far as 1988 (Madison 2009) and which directly led to the development of the WEMI model. As Madison tells us, the focus of these initial meetings was to design an agenda “around cooperative cataloging that focused on standardization and cost-benefits vis-à-vis the current technological and economic environment.” – p 17.

One of the most important of these meetings was a 1990 symposium held in Stockholm just prior to IFLA’s annual conference that year. Madison tells us the scope of this symposium:

“The central issues supporting the framework of the papers and discussions were:

- The mounting costs of cataloging and corresponding interest in simplifying the bibliographic content of cataloging records;
- Interest in decreasing the cost of cataloging by increasing the sharing of bibliographic records and thereby reducing duplicate cataloging, both nationally and internationally;
- The explosion of the amount of published materials—regardless of format—throughout the world, thereby increasing interest in universal bibliographic control;
- The increasing awareness of the benefits to adapt cataloging practices and codes to electronic environments that support online library management systems and mega national and international bibliographic data systems;
- Increasing interest in examining the role of our bibliographic universes through the eyes and needs of their users.” – Madison 2009 p 18.

The FRBR study was a direct outcome of this seminar, and both Tillett and Svenonius were brought into the study as consultants in 1992. As Madison points out, Tillett, through her dissertation, was a pioneer in the use of entity-relationship analysis of bibliographic records and the things described by them. It is not so surprising then, that FRBR and especially the WEMI model it espouses, speaks of the bibliographic universe in terms of entities and relationships. Indeed, Madison tells us that Tillett's approach was successful enough when applied to the FRBR study that in 1996, when IFLA appointed a Working Group on Minimal Level Authority Record and ISADM, Barbara Tillett was appointed its chair (Madison 2009 p 23).

All along, though, a point Tillett had made through the entity-relationship analysis that she carried out in her dissertation was coming to the fore. Specifically, in existing metadata records of the time, there seemed to be sufficient information concerning editions or versions of them that they could be treated as entities in their own right for the purposes of the analysis. It seems likely, then, that the inclusion of the *expression* and *manifestation* entities in the WEMI standard was a result of Tillett's analysis and her subsequent participation in the development of the WEMI conceptual model through the mid and late 1990s.

Arguably, the first actual conceptual model for bibliographic entities in terms of entities and relationships is the one set forth through IFLA's FRBR bibliographic standard (1998), the WEMI model. The WEMI model, which invokes the ER-diagram method (Chen 1976) as an explanation for its appearance (Figure 2.1 below), "represents the different aspects of user interests in products of intellectual or artistic endeavor" (IFLA p 13), and these aspects are named to be—*works*, *expressions*, *manifestations*, and *items* (WEMI).

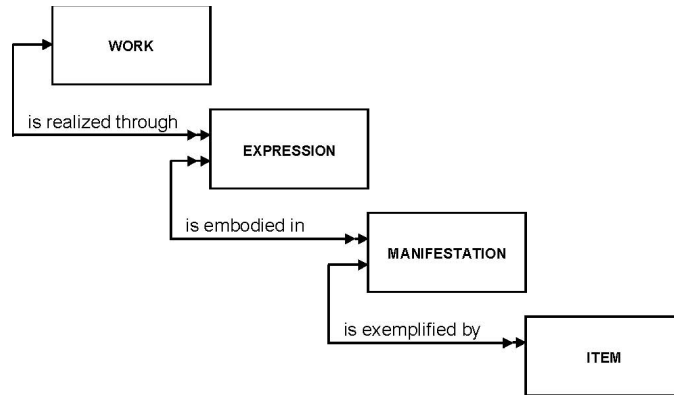


Figure 2.1: WEMI ER-Diagram

As we can see, this is a radical expansion from the binary decompositions of bibliographic entities that Verona, Wilson, and Lubetzky were arguing for (and which seemed to have been well-acknowledged by their contemporaries and peers, such as Ranganathan). However, this model was also very useful because this kind of entity-relationship conceptual design technique is highly valuable in the conceptual design of databases.

Figure 2.1 presents a view of the various aspects comprised by bibliographic entities. It implies that a successful, minimal metadata description for any bibliographic entity needs to provide characteristic information about each aspect that allows bibliographic entities to be grouped by the features of those aspects. Unfortunately, we are unable to provide much more detail because, as many have noticed (Taniguchi 2002, 2003; Renear and Choi 2006; Wickett and Renear 2009, among others), the manner in which these aspects are defined is not particularly illuminating.

From a philosophical perspective, such as the one used as a method in this dissertation, we can say that the conceptual approach taken by IFLA is one that is primarily concerned with describing the roles and participation constraints of entities within the model. We say this because, even though it is unclear precisely what *works* or *expressions* are, the model does tell us that they are the kinds

of things that are directly related to one another through the “is realized through” relationship. Wickett and Renear (2009) go on to point out that this relationship entails that for an *expression* to exist (whatever an *expression* is), there must exist a *work* for which the *expression* is in the role of *realizing*. It is possible, then, to analyze the WEMI model just by examining the participation constraints and roles that have been defined for it.

More pragmatically, we can say that the WEMI model was designed to help us make sense of situations where we have many physically different copies of the same text. We will examine a curious example discovered by Nurmikko-Fuller et al. (2015) during their process of assessing metadata schemas for the HathiTrust Research Center’s¹¹ various usages. In Figure 2.2 below we see what WEMI would call an *expression*, the book, [*The Game of Chess*].

As Nurmikko-Fuller et al. (2015) explain, the book in the figure ([*The Game of Chess*]) is owned by the British Library. In the 1960s, the British Library photographed the entirety of the book to produce a copy of it on microfilm. One of these copies was eventually sold to the University of California, which in turn produced a new print copy of the book from the microfilm. Later, both of these print copies were digitized. A digital copy of the University of California’s digitized version was given to the HathiTrust Digital Library (HT).¹²

¹¹ <https://analytics.hathitrust.org>

¹² <http://hathitrust.org>

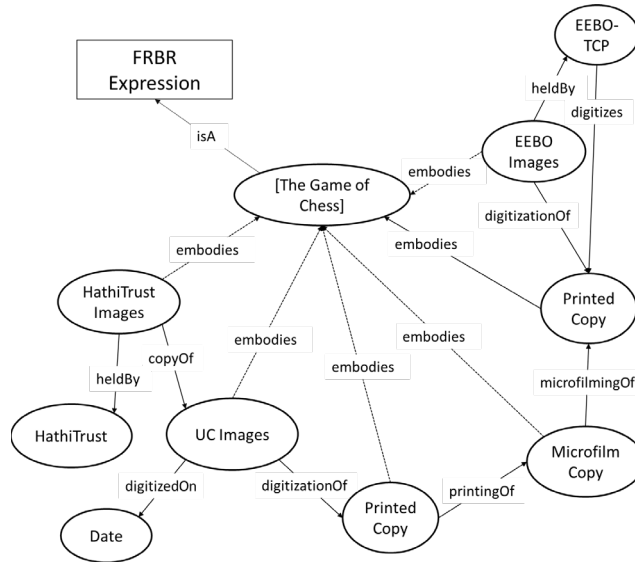


Figure 2.2: Different manifestations of [The Game of Chess] charting its reproduction cycle from the British Library through UC Berkeley to the HathiTrust Digital Library¹³

Nurmikko-Fuller et al. (2015) also tell us that there are separate cataloging records describing this book, one produced by the British Library, and one produced by the University of California. The catalogers who produced these records did not agree how the book should be cataloged and so the facts describing the book, even the title, are altogether different from one another. Despite this, Nurmikko-Fuller et al. (2015) have been able to link all of the copies together, because at some point before the first microfilm copy was produced, someone spilled coffee or tea on the book owned by the British Library. This created an item-level feature that has been preserved in all of the various copies, which according to WEMI, are putatively different manifestations.

In theory, WEMI's *manifestation* and *expression* entities should be able to draw these links exactly as Figure 2.2 illustrates. However, we do have some questions if what we are looking at are called

¹³ Nurmikko-Fuller et al. actually use the figure above as part of a figure contrasting descriptive metadata models with event-based metadata models.

manifestations or *items*. At this point though, we cannot completely resolve them, because as we mentioned above, the information about WEMI's entities is quite sparse.

Not long after FRBR was published, two additional accounts of bibliographic entities, possibly in response to all of the confusion that the WEMI model caused among the cataloging community, were published.

Svenonius (as we mentioned—also a consultant for the committee that initially developed the WEMI model and the chair of Tillett's dissertation committee) provides one of these accounts (2000). In Svenonius' case, she sets out to provide a much different, extremely positivist view of bibliographic entities. She employs set theoretics in an effort to reduce the overall number of entity types and relationships needed by any model of them. Thereby, she focuses on a small handful of things—*sets*, *documents*, and *constrained-set-member-of* relationships—all in the service of providing what some call a reductionist account for bibliographic entities.

This set theoretic view of bibliographic entities arranges “documents” into larger and larger sets, by first grouping all of the “documents” which are the same “versions” into sets, and then grouping all of the “version” sets onto bigger sets that are the same “edition.” In turn, all of the “edition” sets are grouped together into a “work” set which relates them all through having the same content. This imperfect subsumption hierarchy, where larger sets comprise smaller sets, provides a view of bibliographic entities that is quite similar to the one set forth by FRBR's WEMI model. Indeed, we suspect that she was trying to clarify what the WEMI model was trying to do.

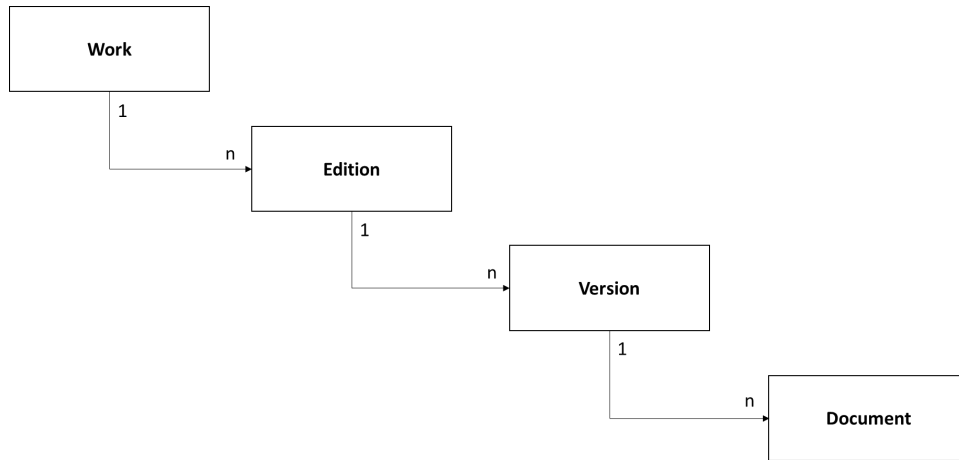


Figure 2.3: ER-Diagram Illustration of Svenonius’s Bibliographic Entities

Ironically, we can use the entity-relationship approach taken by Tillett to provide a partial illustration (Figure 2.3 above) of what Svenonius tells us are distinct bibliographic entities. There are many problems with Svenonius’s account, which we do not go into here. The important factor is that we have a WEMI-like picture of the bibliographic universe, but unlike the traditional entity-relationship approach taken by Tillett in her analysis, Svenonius takes the position that we can define each entity in set-theoretic mathematical terms.

While this approach can be a useful one when we focus on building actual IR systems, especially where we have relatively few works (i.e., our collection is small), it is not necessarily a scalable one because each work in a collection, as well as each edition, version, document, etc., would need its own unique specifications. As Svenonius remarks, “Specification is not easy.” – p 36. In part, this is because individual phenomena must be unpacked so that the language needed to describe them can be worked out. Set theoretics does not do this unpacking in and of itself, so one needs to apply other formal grammars.

For instance, a reductionist model of a phenomenon like ballistic range might propose that the distance a projectile travels can be reduced to a set of vectors. While this is true, this is not very helpful if one wants to model a particular vector. For that, one needs to develop a quadratic equation that can unpack all of the concepts necessary to model an individual ballistic range. This equation includes factors like the initial height of the projectile, the projectile's initial velocity, the time it takes to reach the ground, the angle of its trajectory, and acceleration effects from the force of gravity, etc. Similarly, Svenonius's reductionist account of bibliographic entities is not overly helpful because it is not sufficient to group bibliographic entities into sets. The languages needed to describe individuals in the sets still needs to be specified, and for this task, the conceptual approach will be necessary.

In contrast, Smiraglia's (2001) meditation on bibliographic entities is a much simpler one, but no less useful. However, it is more directly useful for conceptualizing bibliographic entities than for directly developing databases (and is probably also a response to the WEMI model). Like Tillet, IFLA, and Svenonius, Smiraglia is also considering how to unpack the subtle differences among different "documents" which all have a "family resemblance" in the Wilsonian (1968) sense. Instead of proposing a hierarchy of sets though, he sticks to higher-level conceptual spaces, and he proposes a view of bibliographic entities that decomposes them into three aspect entities—"documents," "text," and "works." Figure 2.4 illustrates his model for bibliographic entities.

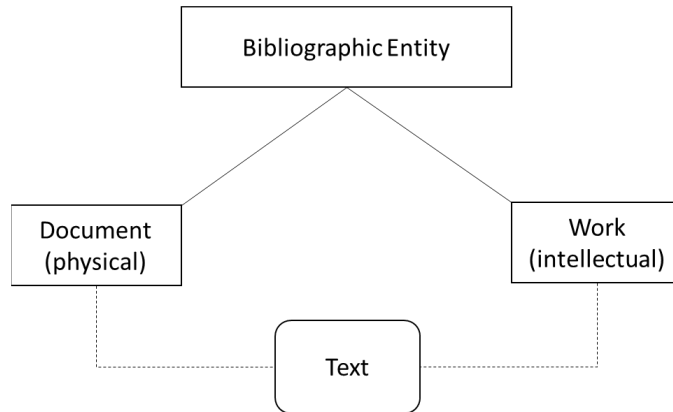


Figure 2.4: Smiraglia’s Bibliographic Entity Model¹⁴

Figure 2.4 is not an ER-diagram. As such, we should not interpret particular shapes as imparting any additional meaning. Indeed, Smiraglia is taking a semiotic approach which we might interpret as a kind of type-token analysis of works and documents (Smiraglia 2008; Wetzel 2009). While we do not employ this particular kind of semiotic analysis ourselves, preferring instead to examine roles and participation constraints on entities, Smiraglia’s diagram does possess two particularly important innovations that have major implications for what metadata is required at a minimum for the basic description of any bibliographic entity. For one thing, it is much clearer that the entity-like boxes in the model: *work*, *text*, and *document*, are not kinds of bibliographic entities in and of themselves, rather they are all aspects comprised by a bibliographic entity. The second important thing follows from this first one—namely, a complete description of a bibliographic entity possesses metadata for all three of these aspects.

Metadata describing each aspect is necessary so that any data infrastructure built around the metadata can group bibliographic entities in one of three ways:

¹⁴ Reproduced from Figure 1.1 in Smiraglia 2001, p 4.

1. By physical characteristics
2. By textual characteristics
3. By intellectual characteristics

Taken together, these three approaches leave us in an interesting space. This is because entity-relationship diagrams are intended to illustrate the conceptual space of data in databases. Thus, we are free to interpret the WEMI model illustrated in Figure 2.1 to actually be describing sets of tuples that correspond to things, i.e., *works*, *expressions*, *manifestations*, and *items*. However, the language on FRBR's p 13 (IFLA 1998) makes it clear that these four things are not distinct bibliographic entities, as Svenonius sets forth, but are instead distinct aspects comprised by a bibliographic entity. Additionally, their individual definitions on the same page of FRBR are vague enough that we might interpret the relationships among them as being the kind of semiotic relationships that Smiraglia discusses.

Complicating these interpretations are a series of attributes that are defined for each of the WEMI entities. Which, in turn, are further complicated by the fact that many in the cataloging community frequently speak of the values of these attributes as being able to be inherited (i.e., to propagate from one entity to another entity) across the disjoint boundaries in the model (Coyle 2016b). This is a trend despite the lack of any language in the FRBR document supporting the sharing of attributes by disjoint entity classes (Renear and Choi 2006). Regardless of these additional complications, we can more clearly say that the WEMI entities form a continuum of classes that describe bibliographic entities in a manner that spans from their physicality as artifacts (e.g.,

“documents,” “items”) to the abstract content that they are used to communicate (e.g., “works”), or in other words, they form what might be called a *content-artifact model*.

2.3.2. Moving Beyond WEMI

As we noted in the section above, information defining the WEMI entities is quite limited, and over the years there have been a number of attempts to either rehabilitate it as a model for bibliographic entities or move beyond it. Taniguchi (2002, 2003), in particular, tries to provide one conceptual mapping that proposes that by *expression*, WEMI’s authors really mean *text*, and by *manifestation*, they mean *medium*. However, with several examples of differences in *medium* (e.g., the film adaptation of *Moby Dick* vs the novel, *Moby Dick*) and Michael Gorman’s earlier (1998) point about describing physical objects (p 27), full rehabilitation of the *manifestation* notion specifically has proved elusive.

In particular, several mediums (e.g., video games, live performances, etc.) seem¹⁵ to evince features that suggest that there are *expressions* that are possibly directly related to other *expressions*, or *manifestations* possibly related to other *manifestations* through *embodiment*-like relationships. These dualistic expressions and manifestations caused problems when attempts were made to expand the general model espoused by WEMI to encompass specific kinds of bibliographic entities spanning from manuscripts to video games and on to live performances of various kinds (Jonsson 2005, Miller and Le Boeuf 2005, Nicolas 2005, Baca and Clarke 2007, McDonough et al. 2010, Lee et al. 2012).

¹⁵ And this still seems to be the case for many, if not all, of these mediums.

Video games (and digital media types in general) have been particularly troublesome. This is because there are frequently multiple languages involved simultaneously—one language for the human consumer and one language for the machine¹⁶ rendering the end-product to the human consumer. As the WEMI model is, itself, based on the decomposition of existing, mass-market monograph-centric catalog records, it is perhaps not too surprising that it begins to quickly break down when confronted by mediums that are well beyond the scope of what monographic-centric catalog records are designed to describe.

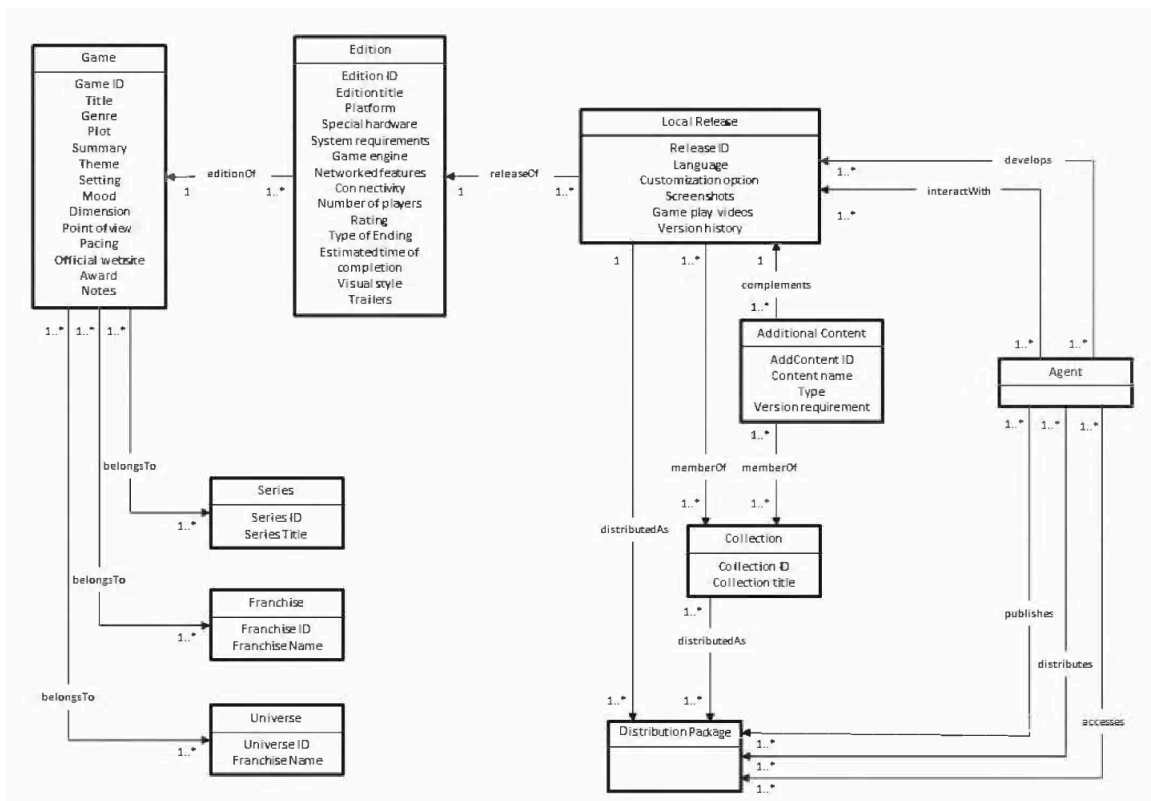


Figure 2.5: Jett et al.'s Conceptual Model for Video Games¹⁷

¹⁶ The actual situation is much more complex even than this.

¹⁷ Reproduced from Jett et al. 2015, p 507.

Various repairs for specific types of bibliographic entities have all been suggested. One example is the video game metadata standard (Figure 2.5 above) suggested by Jett et al. (2015) which more or less abandons the WEMI model in favor of an approach that expands WEMI's *expression* into "edition" and "local release" entities in order to articulate differences in both machine languages (via "edition") and human languages (via "local release").

Similar problems led to the formation of the FRBR working group on Aggregates which developed the *content-artifact aggregate model* (FRBR-WGA 2011; Žumer and O'Neill 2012) deployed in IFLA's new LRM standard (Riva, Le Boëf, and Žumer 2017). And it is here at this point that our account of the state of affairs for describing *bibliographic aggregates* actually turns to them specifically.

2.4. Bibliographic Aggregates

2.4.1. Bibliographic Aggregates as Bibliographic Entities

Several bibliographic entities that libraries commonly collect that are intrinsically composite in nature. Some examples are anthologies, serial publications (like journals and newspapers), bound journal volumes (multiple journal issues bound together as a single monograph-like object), and microfilms (the practice of mounting photographs of the pages of multiple newspaper issues into a single reel of film). Even library (to say nothing of archival and museum) collections themselves can be viewed as bibliographic entities. Indeed, with the advent of digital collections, which can be copied en masse from one computer to another, an entire industry of creating both collections and metadata describing them is evolving (Hunter, Legg, and Oehlerts 2010; Lewis 2013, Barbakoff 2017).

These are all examples of composite bibliographic objects which group together what would, in other circumstances, be singular bibliographic objects. All of these more complex bibliographic objects are generally understood to be forms of *bibliographic aggregates*.

Bibliographic aggregates, especially in the forms of curated collections and serialized publications, have a long history as objects of policy, e.g., collection development policy, and cataloging practice, e.g., through rules and guidelines such as the American Library Association's (ALA) Anglo-American Cataloging Rules (AACR2 [ALA 2002]) and Resource Description and Access (RDA) from the Joint Steering Committee for Development of RDA (JSCD-RDA 2014).

Conceptual debates regarding the nature of library and archival collections in particular can be found in both library-centric (Lee 2000, 2005; Palmer 2004; Palmer et al. 2006; Palmer, Zavalina, and Fenlon 2010) and archive-centric (Currall, Moss, and Stuart 2004; Yeo 2012) LIS literature. How to best model collections as first-class bibliographic entities in IR and metadata management systems has also been discussed in the LIS literature (Lagoze and Fielding 1998, Gonçalves et al. 2004, Galton 2010, Wickett et al. 2013, Wickett et al. 2014, Jett 2015).

Similarly, despite the acknowledged need to describe and represent serialized publications as bibliographic objects within library IR systems, a satisfactory conceptual model eluded the AACR2, as evidenced by the LIS domain's need for supplantation in the form of the Library of Congress's CONSER Cataloging Manual. CONSER has always been an imperfect solution. Indeed, the historical approach to cataloging serials now presents some significant problems for using serials catalog records to support text analytics endeavors. This has resulted in calls to reform

serials metadata practices (Antelman 2004, Kemp 2008, Krier 2012). More recently, there has also been conceptual work on a number of related kinds of *bibliographic aggregates*, such as video games that are compounding resources (Lee, Jett, and Perti 2015), series as bibliographic entities (Jett et al. 2017), and overarching “superworks” that collocate all resources in particular popular cultural domains (Kiryakos and Sugimoto 2018, Sugimoto et al. 2018, Lee et al. 2018).

As we mentioned in the previous section, the WEMI model set forth in FRBR seems to break down when we move from metadata common, to all bibliographic entities, to those needed to sufficiently describe more specific kinds of bibliographic entities, like video games. Similarly, it has also been noted that IFLA’s FRBR framework has difficulty in accommodating *bibliographic aggregates* generally and serialized publications specifically (FRBR Working Group on Aggregates [FRBR-WGA] 2011). As such, they propose a model for aggregates that does not depend on parts and wholes.

However, Barbara Tillett and others suggest in their 2014 “Letter to the Editor” of *Cataloging & Classification Quarterly* that:

“FRBR allows aggregates and components of any entity. The choice of which to identify should be paired with application design decisions and cataloging instructions and policies.” – Tillett et al. 2014, p 360.

In theory, the problem the FRBR-WGA is addressing is not a problem that should have occurred, given that *bibliographic aggregates* in general are relatively generic, and the basic FRBR standard

already possesses relationships (i.e., *part-of*) through which metadata that link bibliographic entities comprising more bibliographic entities to one another (e.g., that links collections and items together) can be recorded. In theory, we should only experience problems when we try to articulate something about more specific *bibliographic aggregates*, such as digital library collections. Nonetheless, a WEMI-centric solution (the *content-artifact aggregate model*) to the purported problem was suggested. As an aggregate model, the *content-artifact aggregate model* is not without its own forbears.

As early as 1987, in the pages of Tillett's own dissertation thesis, an aggregate model that is suspiciously similar to the *content-artifact aggregate model* is described. Or as Tillett says then:

“The whole-part (or part-whole) relationship holds between a component part of a bibliographic item or work and its whole, such as between a short story and the anthology in which it is contained. The components may be parts of some particular physical manifestation of a work, that is, parts of a bibliographic item, or they may be parts of some abstract work. For instance, The Wife of Bath's Tale is a component part of The Canterbury Tales. When a library has a separately published edition of The Wife of Bath's Tale and wants to show its relation to The Canterbury Tales, the relationship may be understood to hold between a physical item (the edition the library has) and the work as an abstract whole.” – Tillett 1987 p 59.

As a means to delineate when parts and wholes are both “items” and when the parts are “items” and the wholes are “works,” she sets forth a taxonomy of three whole-part relationships:

“containing relationships,” “extractive relationships,” and “abstract relationships.” Of the first relationship, Tillett tells us:

“The category ‘containing relationship’ specifically refers to those relationships involving the component parts of a physical unit, other than extracted parts. A containing relationship characterizes monographs and their individual chapters, published sets and their individual volumes, as well as series and their subseries. The series-subseries relationship typically is more complex than the other two examples of containing relationships, because a series may include collections or sets of monographs, or a series may be part of a larger series in a series hierarchy. In any case, the use of the term ‘containing relationship’ to identify this category connotes actual parts of some physical unit.” – Tillett, pp 59-60.

It would seem that common-place, specialized bibliographic entities like digital collections, anthologies, series, etc. are all kinds of containers, or at least they are if we take the term “container” to possess the meaning and connotations that we are typically familiar with.

With regard to “extractive relationships” Tillett tells us:

“When the parts of an item have been extracted and issued separately as individual selections, the relationship between the extracted items and the whole is categorized as an ‘extractive relationship.’ This category obviously excludes exact reprintings of a whole edition. Such reprintings are considered equivalent works, whereas

extracts must be considered precisely equivalent only to passages, lines, or other small portions of a work. As for detached copies that are parts of a larger work, their relationship to the part they copy is also an equivalence relationship, while their relationship to the whole work from which they are detached is whole-part.”

– Tillett, p 60.

This is an unusual category and seems to have some overlap with a non-part/whole category of relationships that Tillett calls “derivative relationships.” Indeed, it is probably best articulated as a kind of derivation, as it is difficult to claim that a short story reprinted separately as an individual selection is really part of the same contextual unit that it was originally published with. However, this is similar space from which the *content-artifact aggregate model* appears to evolve as the “copies” in the quoted text seem to indicate “items” that have been detached from a “whole work.”

One of the problems here is to what extent context matters. If Palmer (2004) and Palmer, Zavalina, and Fenlon (2010) are correct, then the very act of aggregating creates a context distinct from that in which the individual items exist. To some extent, this seems to bear out when Tillett says that the individual “items” are part of the “whole work,” even though they have been detached. The context provided by the “whole works” seems indispensable for the meaning of the separated “items.” As we will see later, this is actually the opposite position of the *content-artifact aggregate model* set forth by the FRBR-WGA (2011) and canonicalized in LRM (Riva, Le Boef, and Žumer 2017). But for now, we have this odd part/whole relationship between what must be abstract entities, “works,” and concrete (or at least textual) entities, “items.”

Finally, with regard to “abstract relationships,” Tillett tells us:

“The ‘abstract relationship’ holds between parts of a work and the work. Work here is to be understood as an abstraction. The term ‘abstract relationship’ is used therefore to convey the connotation of a relationship to some abstract whole rather than some physical item. This relationship is further described in the discussion of the linking device, uniform titles (see III. C. 4. c. 1)).” – p 61.

This is once again, a part/whole relationship, like the “container relationships” that operates with a specific aspect of a bibliographic entity, that of “work.” However, like the previous relationship, there seems to be some intellectual overlap with another kind of relationship in Tillett’s overall taxonomy—linking relationships. Unfortunately, it is not particularly informative beyond telling us that the part/whole relationship between parts of works and whole works is abstract. On the whole, it seems as though one of the problems we face when developing and discussing our conceptual models for things like *bibliographic aggregates* is that the definitions and examples often employ part-whole terms that are not actually synonymous with one another (e.g., containers, components, members, etc.).

2.4.2. The Problem of Paratext

Paratextual features present an additional problem for cataloging efforts in general and for the description of *bibliographic aggregates* in particular. As Coyle (2016b) points out in her chapter, “Some issues that arise in FRBR,” paratext (Genette 1997) can pose a significant problem for the cataloging of aggregates. Coyle says:

“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.” – Coyle 2016b, p 131.

Coyle presents paratext as a problem for the FRBR (or really the WEMI) model in particular; however, this does seem to be a problem for cataloging in general. The question is whether paratextual content is significant enough to change the “primary” content of a *work*. When it is deemed so (as in the case of a scholarly or critical edition) then it seems clear that the *work* can be treated as a derivative work. When it is not, then it seems altogether safe to not record any metadata regarding it at all. This position may have particular implications for cataloging in general; however, for our efforts in this dissertation, we are going to treat paratext as something that does not significantly impact our understanding of what *bibliographic aggregates* are. Thereby, we will ignore those examples in the four standards examined in this dissertation that seem to be paratextual features (e.g., a table of contents in the role of being a bibliographic part). We will revisit this issue in our concluding chapter.

2.4.3. An Initial Analysis of Bibliographic Aggregates

2.4.3.1. *Confounding Factors*

One problem with using different part-whole terms indiscriminately is that it causes confusion with respect to what the essential conceptual nature of the whole (i.e., the aggregate) is in relation to the part. We can showcase an example of this confusion by examining the FRBR standard's treatment of "Aggregate and Component Entities." What FRBR states is:

"The structure of the model [...] permits us to represent aggregate and component entities in the same way as we would represent entities that are viewed as integral units. That is to say that[,] from a logical perspective the entity *work* for example, may represent an aggregate of individual *works* brought together by an editor or compiler in the form of an anthology, a set of individual monographs brought together by a publisher [or author] to form a series, or a collection of private papers organized by an archive as a single *fond*." – IFLA 1998, p 29.

We are referred to subsequent sections (which we examine in more detail in Chapter 4) for additional examples corresponding to each of the aspects of the bibliographic entities articulated in the WEMI model. From the example in the text above, the authors mean that the "components" of *works* are other *works*. However, this does not seem to correspond to some of our commonsense notions of the components of a *work*. Surely it is the case that we would normally consider the main character (i.e., a person, fictional or otherwise) of a *work* to be a substantial *component* of that *work*. But, on no account would we conceptualize a person to be a *work* in their own right. We do not typically hold it to be true that people (fictional or otherwise) are merely some kind of intellectual content.

What is likely meant by the word “component” in this context is some other part/whole relationship. But how can we determine which one?

One method is to reuse a tried and true solution. While the LIS domain may be ill-equipped for the description of foundational concepts, such descriptions are a matter of course in the philosophical domain.

2.4.3.2. Collection-Item Metadata Relationships – A Brief Case Study

During the late 2000s, a group of researchers led by Allen Renear at the University of Illinois at Urbana-Champaign analyzed the relationship between metadata describing digital collections and metadata describing the digital objects gathered into them (Renear et al. 2008a; Renear et al. 2008b; Wickett, Renear, and Urban 2010). For this conceptual analysis work, they employed an ontological analysis approach (OntoClean) set forth by Guarino and Welty (2004). The approach focuses on using formalisms in first-order logic to interrogate the metaproperties of relationships between entities, as defined in formal ontologies such as the Dublin Core Metadata vocabulary.

Guarino and Welty were particularly interested in determining when relationships in ontologies represented contingent roles that entities were playing. However, as a formal analysis technique, the OntoClean method also has clear applications for ontology alignment (through metaproperty alignment) and broader analysis of metaproperties possessed by relationships. Or in simpler terms, Renear et al. were focused on the roles of entities and relationships in the conceptual models and the participation constraints that governed them. An approach different from, but complementary to, reductionist accounts of conceptual spaces like Svenonius’s.

The work ultimately resulted in a framework describing how particular attributes of the members of a collection are reflected by attributes of the collection through propagation relationships (Wickett 2018). For our purposes, we want to narrowly focus on one particular paper produced by this group—“Are Collections Sets?” (Wickett, Renear, and Furner 2011). In this paper, Wickett, Renear, and Furner are engaging with an ongoing issue in ontologies and conceptual models—the use of sets as a basis for entities. Specifically, they are contemplating whether or not the notion of collections set forth in several schemas (Powell, Heaney, and Dempsey 2000; Shreeves and Cole 2003; DCMIDTG 2007)¹⁸ are best represented as sets, a conclusion reached by several of the approaches to characterizing digital collections at the time (Lagoze and Fielding 1998; Gonçalves et al. 2004; Meghini, Spyrtos, and Yang 2010).¹⁹

2.4.3.3. *Part/Whole Conceptual Spaces*

As we already mentioned, the LIS domain (and other domains too) frequently conflates conceptually distinct terms when it defines *bibliographic aggregates*. To overcome this issue in “Are Collections Sets?” Wickett, Renear, and Furner import an already well-developed formal conceptual space that describes parts and wholes as primary objects of research. This conceptual space is called Classical Extensional Mereology (**CEM** [Simons 1987; Varzi 1996; Varzi 2016]).²⁰

Importantly, **CEM** provides a basic framework for analyzing part/whole relationships using axioms, definitions, and theorems that are already well-established. We explicate a truncated

¹⁸ All of which are based on an ER-diagram for collections set forth by Heaney (2000).

¹⁹ And as we shall see in Chapter 5, this approach is still alive and well.

²⁰ They don't tell us that they're using this well-developed vocabulary until deep into the analysis.

version of **CEM**—[Core] Mereology (**M**)—here, to gain a better sense of the conceptual space that Wickett, Renear, and Furner employed.

As a great deal of work has already been established considering the formal status of parts and wholes, this dissertation employs existing work as framework for interpreting the various part/whole relationships set forth in the four conceptual models selected for analysis. The part/whole framework—**M**—will provide a thorough background for developing a more precise account of aggregate entities, like those set forth in the *mereological aggregate models*. The following text reproduces well-established axioms, definitions, and theorems describing *parthood* and *proper parthood* from sources like Simons (1987) and Varzi (2016).

As Simons, Varzi, and other philosophers tell us, the *part-of* relationship typically possesses the following metaproperties:

- *Reflexive* – something is always part of itself
- *Antisymmetric* – the *part-of* relationship is directed, and its domain and range must be different entities
- *Transitive* – parts of parts are also part of the whole

The following three axioms express that the *part-of* relationship possesses the *reflexive*, *antisymmetric*, and *transitive* metaproperties.

A.1 (Reflexive Parthood): $\forall x(\text{partOf}(x, x))$

Axiom A.1 states that an entity, x , is always part of itself.

A.2 (Antisymmetric Parthood): $\forall x \forall y ((partOf(x, y) \wedge partOf(y, x)) \rightarrow (x = y))$

Axiom A.2 states that if an entity, x , is a part of an entity, y , and that entity, y , is also a part of entity x , then entity x is identical to entity y .

A.3 (Transitive Parthood): $\forall x \forall y \forall z ((partOf(x, y) \wedge partOf(y, z)) \rightarrow part(x, z))$

Axiom A.3 states that if an entity, x , is a part of an entity, y , and that entity, y , is, in turn, a part of another entity, z , then entity x is also part of entity z .

Together, these three axioms form **M** (Varzi 2016). These are not actually all the axioms, definitions, and theorems that Wickett, Renear, and Furner employ in their analysis (recall they are using **CEM**), but they are sufficient for us to get a sense of what they are doing here. Rather interestingly, they find that the *is-gathered-into* relationship they are investigating is more comparable to Frænkel and Bar-Hillel's ZFC *set-member-of* relationship. Like the *set-member-of* relationship (which is always an intransitive relationship), they find that the *is-gathered-into* relationship is sometimes an intransitive relationship (in effect noting that axiom A.3 above cannot hold for all instances of the *is-gathered-into* relationship). The implication is that the *is-gathered-into* relationship is not related to the *part-of* relationship described by **M**.

As Wickett, Renear, and Furner observe, “Allowing collections to be members of collections with transitivity would distinguish collections and [the] *isGatheredInto* [relationship]²¹ from sets and set relationships” – p 4. However, they ultimately reject this position, saying:

²¹ Note that we use a slightly different label here: *is-gathered-into* vs *isGatheredInto*, nevertheless these labels are intended to invoke the same relationship.

“If collections can be gathered into collections, as opposed to the items of some distinguished collections being gathered into a collection, then this creates a hierarchical structure within the collection. Allowing *isGatheredInto* to be transitive then collapses this structure. In order to preserve the intentions of curators who choose to gather whole collections instead of individual items from collections, we can consider transitivity not to hold for *isGatheredInto*.” – p 4.

Ultimately, they reject a variation of axiom A.3 specific to the *is-gathered-into* relationship. Their reason for doing so hinges on two arguments:

1. That transitivity collapses the hierarchical structure of collections, and
2. That rejecting transitivity helps preserve the intentions of curators who choose to gather whole collections instead of individual items from collections.

However, there is a problem here. They no longer seem to be talking about instances of the *is-gathered-into* relationship (i.e., instances of collections), but instead, they seem to be talking about situations in which collections might find themselves. The apparent hierarchical structure of a collection is only useful when agents, like us, want to discuss apparent structures within the collection (regardless of whether such structures are there).

We can contrast this with another example where it does seem to be the case that some instances possess a specific metaproperty and other do not: overlap. We can formally define overlap using the following axiom which says that two wholes *overlap* when they have a part in common.

D.1 (Definition of Overlapping Parts):

$$\forall x \forall y \left(\text{overlaps}(x, y) =_{df} \exists z (\text{partOf}(z, x) \wedge \text{partOf}(z, y)) \right)$$

Having defined overlap in general for part/whole relationships, let us propose a variation of the definition specific to the *is-gathered-into* relationship which simply tells us that when something appears in two different collections then those collections overlap.

D.1-C (Definition of Overlapping Collection Members):

$$\forall x \forall y \left(\text{overlaps}(x, y) =_{df} \exists z (\text{isGatheredInto}(z, x) \wedge \text{isGatheredInto}(z, y)) \right)$$

The question becomes, is it possible for two different collections to share the exact same item? If we consider the works in a digital library collection like the HathiTrust's, are there any instances where they are shared with other collections? The answer to this question is a resounding yes. Yes, because the HathiTrust has a tool that lets its users gather together digital objects from its corpus and curate them into their own personalized digital collections.

As an aside, we might point out that the HathiTrust has such a tool because it is trying to support additional primitive user activities—specifically Unsworth's *referring*, *sampling*, and *representing* activities (2000). These activities may be a step beyond the kinds of functionality (*discovery*) that is traditionally thought to be important for libraries to support (Verona 1959; Wilson 1968; Lubetzky 1969; IFLA 1998); however, as Unsworth points out, these are vital primitive activities that all researchers perform and that libraries should be supporting. Coyle (2016b) also remarks on Unsworth's primitives (p 43), but the timbre of her prose suggests that she is not convinced that libraries should support primitives beyond *discovery*.

Nonetheless, the important thing to note here is that we have clear evidence that there exist some collections (i.e., some instances of the *is-gathered-into* relationship) that evince the metaproperty of overlap (because they have overlapping members). This is not the case at all with transitivity. In the transitivity case, it is simply that sometimes, we want to examine the internal hierarchy of a collection, and as such, all instances of the *is-gathered-into* relationship must necessarily lack the property of transitivity, else the hierarchy necessarily collapses, and the examination becomes impossible. At other times, we want to make claims about what items are in the collection regardless of whether or not it is actually in a sub-collection within the collection; in these cases, all instances of the *is-gathered-into* relationship must possess the metaproperty of transitivity, otherwise no claims that the items in the sub-collection are also part of the parent collection can be made.

The real problem is that axiom A.3 (and all possible versions of it) is that it is too coarse of a tool for the analysis Wickett, Renear, and Furner are carrying out, and so they are unable to reconcile those situations in which any given instance of the *is-gathered-into* relationship seems to be transitive with situations where it seems to be intransitive.

If we cannot reconcile these two situations, then it would seem we have to reject some of our commonsense intuitions regarding collections and the things in them. For instance, if we accepted that the *is-gathered-into* relationship is an intransitive one, then we would no longer be able to say things like the library's collection includes the book *Moby Dick* if, in fact, that book is actually included in the library's special collection of 19th Century literature.

This problem is solvable, although probably not in the space that Wickett, Renear, and Furner have been allotted, as “Are Collections Sets?” is a conference paper. Specifically, there are additional tools suggested in the existing mereological literature that are expressly designed to cope with the first argument (hierarchy collapse) that Wickett, Renear, and Furner set forth.

The collapse of hierarchy is a problem that Bittner and Donnelly (2005) faced when they analyzed the differences between the *containment* and *componenthood* relationships. As Bittner and Donnelly realized, sometimes we want to talk about properties of part/whole relationships that are clearly transitive and sometimes we want to talk about properties of part/whole relationships, like hierarchy, that are clearly intransitive. Their solution for *containment* and *componenthood* relationships is also applicable to the *is-gathered-into* relationship. We discuss that solution further in the next chapter.

2.5. Chapter Summary

As we have seen in this chapter, adequate description of bibliographic entities through metadata is the central feature of the library catalog. Beginning as early as the eighteenth century, the task library catalogs are designed to accomplish was severely complicated by the advent of the mass-publication industry. These complications were compounded throughout the nineteenth and twentieth centuries by the continuous invention of new media forms and communication methods. In turn, these complications have become the driving force in the further development of library-centric conceptual models for bibliographic entities.

Bibliographic aggregates present a specialized sub-class of bibliographic entities wherein the wholes and their parts require equal descriptive accounts in order for users to make full use of

them as information resources. Just as the overall account of bibliographic entities has evolved over time, so too has the general account of *bibliographic aggregates*. However, the language employed by conceptual standards is frequently unclear, since it uses a variety of part/whole terms (e.g., using terms like “component” and “member” interchangeably) that actually play distinct roles in part/whole conceptual frameworks. Conceptual analysis work like Wickett, Renear, and Furner’s showcases how the application of informal formal methods can help to clarify our understanding of complex concepts like *bibliographic aggregates*.

In the next chapter, we review the existing mereological literature in order to define a fine-grained set of tools that will help us showcase the differences between the *componenthood* relationship, the kinds of part/whole relationships that our bibliographic standards likely mean, and an array of other part/whole relationships. To some extent, this means temporarily stepping away from LIS literature to better focus on the mereological literature that is going to provide us with a rich enough conceptual space to describe *bibliographic aggregates* with enough detail to distinguish them from other kinds of aggregates.

3. Part/Whole Conceptual Frameworks

3.1. Chapter Overview

In this chapter, we examine existing conceptual accounts of parts and wholes, in an attempt to reconcile the use of part/whole terms employed in our bibliographic standards documents (and their conceptual models) with their traditional semantics in mereological contexts. Thereby, the chapter begins with a deeper discussion of the terms being used in the context of the standards documents they are employed in. This is followed by a longer treatise on the use of existing mereological theory, specifically Extensional Mereology (**EM**). Following this explication, we examine Winston, Chaffin, and Herrmann's (1987) taxonomy of mereonymic (i.e., part/whole) relationships. As we note, modern approaches in computer science contexts expand upon this basic hierarchy.

Having established the basic framework for the analysis, we then examine the properties (a.k.a. metaproperties) of each relationship, to both differentiate each named relationship from one another and to identify those metaproperties that are of particular interest in our analysis. **EM** is then extended to incorporate those metaproperties and an initial, deeply informal, analysis is made to narrow the number of candidate part/whole relationships to just those that possess most, or all, of the metaproperties thought to play important roles for *bibliographic aggregates*. These are found to be the *component-of*, *contained-in*, and *member-of* relationships. The chapter concludes with a deeper analysis of all three of the candidate part/whole relationships.

3.2. Parts, Wholes, and Bibliographic Aggregates

As we noted in the previous chapter, many of our standards use part/whole terms without regard for the fact that in many instances, such terms indicate very different kinds of part/whole

relationships. The state of being a component is very different than the state of being in a container. In this chapter, we are going to closely examine some of the existing mereological literature to suggest a number of metaproperties (i.e., dimensions or facets) through which we can illustrate the differences among various part/whole relationships.

We are doing this so that we can account for the differences among the part/whole terms used in the bibliographic standards we are examining. For instance, the DC-CAP uses the label *is-gathered-into* to make *bibliographic aggregates* distinct from other kinds of parts and wholes, for which they have a more general *part-of* relationship to describe (DCCDTG 2007). Perhaps contrarily, IFLA's FRBR standard uses the label *part-of* but places narrow constraints on the range and domain of the *part-of* relationship (IFLA 1998) and, as we saw in the previous chapter, frequently uses the term "component" in language describing the parts (e.g., the section heading "3.3[:] Aggregate and Component Entities"). Similarly, FRBRoo uses three different labels—*component-of*, *member-of*, and *incorporates*—to signal apparently different senses of *part-of* with regard to three specific, apparently different, kinds of *bibliographic aggregates* (Bekiari et al. 2016).

In the subsequent sections of this chapter, we will develop a conceptual framework through which we will be analyze the different part/whole relationships set forth in our existing library metadata standards. We will do this by first considering other examples where different labels have been used to differentiate among different kinds of part/whole relations. Specifically, we will do two things:

1. We will consider a series of formal axioms and definitions in modal first-order logic with equivalence from the existing mereological literature with which to differentiate the properties (or more properly the meta-properties) possessed by various part/whole relations.
2. We will also examine the rich taxonomy of part/whole relationships set forth by Winston, Chaffin, and Herrmann (1987) who postulate the following part/whole relationships (component-object relations, member-collection relations, portion-mass relations, stuff-object relations, feature-activity relations, and place-area relations) that are related to but still different from the general *part-of* relation.²²

Afterwards, in Chapter 4, we will employ modified versions of the existing axioms and definitions from this chapter to suggest initial formalizations for potential metaproperties that the relationships used by the *mereological aggregate models* seem to possess. We will use these initial formalizations as the basis for our analyses in Chapter 4.

3.3. Winston, Chaffin, and Herrmann's Meronymous Relationships

3.3.1. Establishing an Initial Formal Framework

As we saw in the previous chapter, Wickett, Renear, and Furner used a formal framework called Classical Extensional Mereology (**CEM**) to provide a conceptual space in which to carry out their analysis. Our first task here is to provide the additional axioms, definitions, and theorems to build a similar conceptual space; however, unlike Wickett, Renear, and Furner, we are going to build

²² We should note that Winston, Chaffin, and Herrmann's taxonomy is a feature in several ontological approaches to describing parts and wholes (see for example, Bittner and Donnelly 2005; Guizzardi 2005; Keet 2006a, 2006b, among others).

our conceptual space using a simplified formal framework for mereology, called Extensional Mereology (**EM**).²³ We will begin by adding additional definitions to the three axioms (**M**) that we provided in the previous chapter, precisely as Varzi proceeds in his explanation (2016).

In his 2016 explication of Mereology for the Stanford Encyclopedia of Philosophy, Varzi also uses **M** (i.e. A.1~A.3) to define several additional mereological predicates. We reproduce them here as definitions D.2~D.6, as they contribute towards arriving at **EM**, which is the mereological framework used in the analysis carried out in the following chapter.

D.2 (Definition of Equality):

$$\forall x \forall y (equalTo(x, y) =_{df} (partOf(x, y) \wedge part(y, x)))$$

Definition D.2 states that an entity, x , is equivalent to another entity, y , when entity x is a part of entity y , and entity y is also a part of entity x .

The addition of definition D.2 to **M** provides the *casus belli* for a new relationship—*proper-part-of* that is irreflexive in nature instead of reflexive. In other words, *proper parts* are not parts of themselves.

²³ We will note here that our reason for doing this is that **CEM** also employs two axioms from which the entire mathematics of algebra can be derived. Similarly, a related formal framework to **CEM**, General Extensional Mereology (**GEM**), has an axiom from which functions can be fashioned and which can be further extended to derive the mathematics of set theoretics. **GEM** and closely related frameworks (e.g. **GEM+**, **AGEM**, and **AGEM+**) in particular is often employed in computer-science-oriented explications of part/whole relationships (see for example Guizzardi [2005] and Keet [2006a, 2006b]).

D.3 (Definition of Proper Parthood):²⁴

$$\forall x \forall y \left(\text{properPartOf}(x, y) =_{df} \left(\text{partOf}(x, y) \wedge \sim \text{partOf}(y, x) \right) \right)$$

Definition D.3 states that some entity, x , is a proper part of some other entity, y , if it is the case that entity x is part of entity y and it is not the case that entity y is part of entity x .

The next definition tells us that for every part, there is some whole which it is part of.

D.4 (Definition of Proper Extension):

$$\forall x \forall y \left(\text{properExtension}(x, y) =_{df} \left(\text{partOf}(y, x) \wedge \sim (x = y) \right) \right)$$

Definition D.4 states that an entity, x , is the proper extension of an entity, y , when that entity, y , is a part of entity x and it is not the case that entity x is identical to entity y .

From here Varzi (and others) usually introduce the notion of overlapping wholes or wholes that share parts.

D.5 (Definition of Overlap):

$$\forall x \forall y \left(\text{overlaps}(x, y) =_{df} \exists z \left(\text{partOf}(z, x) \wedge \text{partOf}(z, y) \right) \right)$$

Definition D.5 is a reprise from the previous chapter and, for the record, states that an entity, x , overlaps with another entity, y , when there exists some entity, z , such that entity z is part of entity x and entity z is also part of entity y .

²⁴ Note that this is subtly different from a variant of D.3 which can be written: $PPxy =_{df} Pxy \wedge \sim (x = y)$ –see Varzi (1996) for a full explication of the variant’s equivalence to the formalization in D.3. These two accounts are sometimes distinguished by adding the word “strict” to the formalism appearing in D.3.

Varzi (and others) usually also introduce a scoping definition designed to tell us when two parts are part of the same whole. They call this concept “underlap.”

D.6 (Definition of Underlap):

$$\forall x \forall y (\text{underlaps}(x, y) =_{af} \exists z (\text{partOf}(x, z) \wedge \text{partOf}(y, z)))$$

Definition D.6 states that an entity, x , underlaps another entity, y , when there exists some entity, z , such that entity x is part of entity z and entity y is also part of entity z .

These five definitions (D.2~D.6) form a quintet of basic mereological relations. A more intuitive sense of what each relationship is intended to communicate with regard to mereological status of entities can be gained from the Venn diagram in Figure 3.1 (below).




Figure 3.1: Basic Patterns of Mereological Relationships²⁵

If the entity y in the definitions we have already described (D.2~D.6) is the same as B in Figure 3.1 above, then we can see that C is a proper part of B. Similarly, B demonstrates equivalence as it is defined in D.3. As defined, only A is a proper extension of B. With regard to overlap, the

²⁵ Adopted from Varzi (2016, p 15).

diagram visually demonstrates that only E does not overlap with B.²⁶ Finally, as it has been defined, all of the circles underlap with B in the context of A.²⁷

These basic definitions play an important role in characterizing what happens when one says an item, x , *is-gathered-into* a collection, y . However, these four additional definitions are not sufficient in and of themselves to produce **EM**. To arrive at the framework for our analysis, we must also supplement **M** with two additional axioms describing various aspects of overlapping.

A.4 (Weak Supplementation):

$$\forall x \forall y (properPart(x, y) \rightarrow \exists z (partOf(z, y) \wedge \sim overlaps(z, x)))$$

Axiom A.4 states that if an entity, x , is a proper part of another entity, y , then there exists some other entity, z , such that z is part of y and does not overlap with x . We should also note that by adding axiom A.4 to our conceptual space we have arrived at a new formal framework for part/whole relationships typically called Minimal Mereology (**MM**). We can add one an additional axiom to finally arrive at **EM**.

A.5 (Strong Supplementation):

$$\forall x \forall y (\sim part(y, x) \rightarrow \exists z (partOf(z, y) \wedge \sim overlaps(z, x)))$$

²⁶ By some accounts, entity z can be thought of as a “null item” which is a part of everything, in which case all of the circles in the diagram overlap. As Varzi (2016) points out, this is a controversial position and any discussion of it is outside the scope of the dissertation.

²⁷ Similar to the overlap case, by some accounts entity z can be thought of as a “universal entity” of which everything is a part. In this case though, the existence of such an entity does not affect underlap relationships illustrated in the figure. The existence of a “universal entity” is similarly controversial (Varzi 2016).

Axiom A.5 states that if an entity x is not a part of another entity y , then there exists some other entity, z , such that z is part of y and does not overlap with x . When axioms A.4 and A.5 are added to **M** the result is **EM**.

From here we want to focus on the notion *proper parts* that definition D.3 introduces and produce new versions of the axioms of the axioms used in **M**. We know from definition D.3 that a new version of axiom A.1 will not work, since by definition, the *proper-part-of* relationship is not a reflexive relationship. So, we will need to propose an entirely new axiom that says nothing can be a proper part of itself.

A.6 (Irreflexive Proper Parthood): $\forall x(\sim \text{properPart}(x, x))$

Axiom A.6 states that an entity, x , cannot be a proper part of itself.

Conversely, nothing about definition D.3 implies that the *proper-part-of* relationship is not antisymmetric or transitive like our existing *part-of* relationship. Our next two axioms encapsulate these similarities.

A.7 (Antisymmetric Proper Parthood):
 $\forall x \forall y \left((\text{properPart}(x, y) \wedge \text{properPart}(y, x)) \rightarrow (x = y) \right)$

Axiom A.7 states that if an entity, x , is the proper part of an entity, y , and entity y is also a proper part of entity x , then entity x is identical to entity y .

A.8 (Transitive Proper Parthood):
 $\forall x \forall y \forall z \left((\text{properPart}(x, y) \wedge \text{properPart}(y, z)) \rightarrow \text{properPart}(x, z) \right)$

Axiom A.8 states that if an entity, x , is a proper part of an entity, y , and that entity, y , is, itself, a proper part of an entity, z , then entity x is also a proper part of entity z .

From **EM** and our three additional axioms, we can produce a table (Table 3.1 below) of metaproperties that various part/whole relationships may or may not possess. Here a “+” indicates that all instances of that relationship necessarily possesses a particular metaproperty, a “-” indicates that while **specific instances** of that relationship may or may not possess the metaproperty, they do not all necessarily possess that metaproperty.²⁸

Table 3.1: Part/Whole Relationship Metaproperties

Metaproperty	part-of	proper-part-of
<i>Reflexive</i>	+	-
<i>Irreflexive</i>	-	+
<i>Symmetrical</i>	-	-
<i>Asymmetrical</i>	-	+
<i>Antisymmetrical</i>	+	+
<i>Transitive</i>	+	+
<i>Weakly Supplementing</i>	-	+
<i>Strongly Supplementing</i>	+	+

As Table 3.1 illustrates, the *proper-part-of* relationship is a much more constrained relationship than the more general *part-of* relation. This is so because the *proper-part-of* relationship possesses more metaproperties necessarily than the general *part-of* relation.

3.3.2. Initial Distinctions

For our purposes, we want to focus more narrowly on the specific kinds of examples employed in our existing bibliographic conceptual models. Specifically, we want to focus on those examples

²⁸ Note that some of the metaproperties in Table 3.1 are mutually exclusive, e.g., reflexive and irreflexive.

where it seems the presence of a full-fledged, highly descriptive metadata record is going to be useful to library users. We already know from Brack, Palmer, and Robinson (2000), Sweet and Thomas (2000), Foulonneau et al. (2005) and similar sources that such metadata records offer library users several benefits when assessing large multi-part resources like scholarly research collections for new or additional uses. The separate creation of metadata describing journal series as grand, overarching information resources, and metadata specific to journal articles, has been common practice for a long period of time. Hence, we also want to identify which of the examples employed by our conceptual standards can meaningfully be identified as *bibliographic aggregates*, and which ones seem to actually be indicating things for which we do not need full-fledged bibliographic control in the form of metadata records.

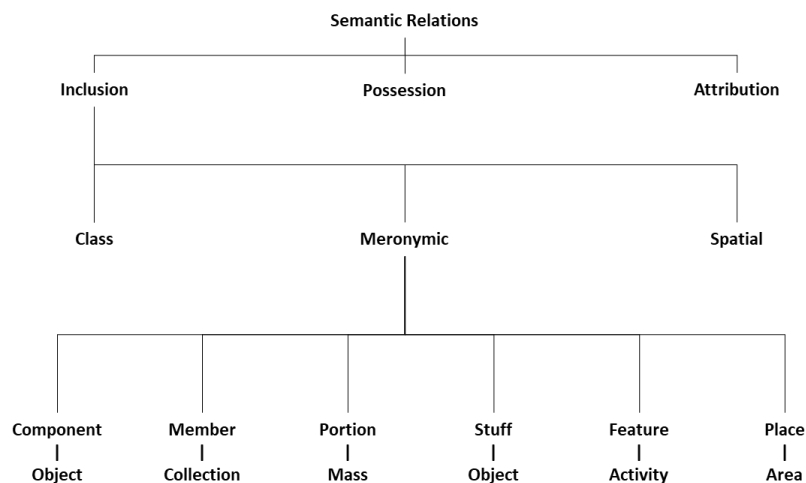


Figure 3.2: Winston, Chaffin & Herrmann’s Partial Classification of Semantic Relations

Work by Winston, Chaffin and Herrmann (1987) demonstrates that a wide variety of what they call *meronymous* relationships abound. Several of the relationships we noted that our *mereological aggregate models* used (e.g., FRBRoo’s *component-of*, *member-of*, and *incorporates*) seem to be

good matches for meronymous relationships named by Winston, Chaffin, and Herrmann in Figure 3.2 (above).

However, before we go on, we need to note that modern interpretations of Winston, Chaffin, and Herrmann’s mereonymic relationships are frequently rearranged into the hierarchy illustrated in Figure 3.3. This is because the sense of the spatial is not so easily subtracted from relationships like *component-object* or *place-area*.

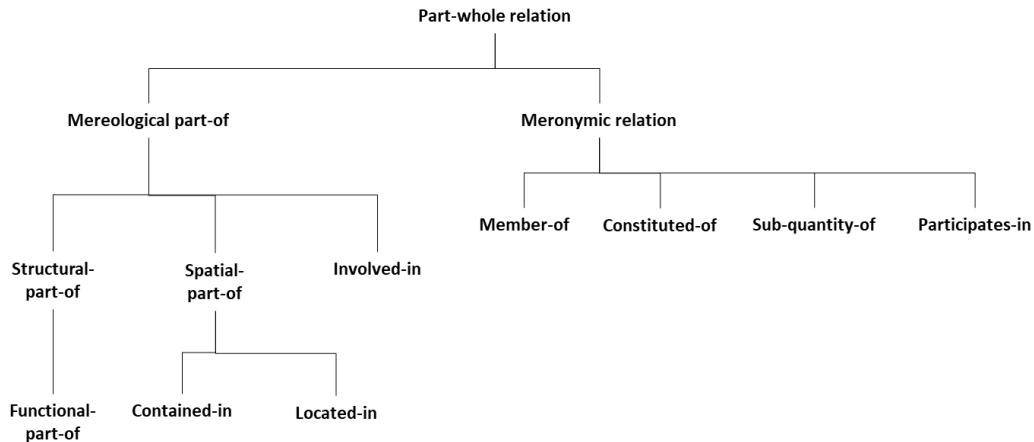


Figure 3.3: Alternative View of Winston, Chaffin, and Herrmann’s Hierarchy²⁹

Functionally, the differences between Figures 3.2 and 3.3 are not going to impact the analytic framework we are developing here. This is because all the relationships in the hierarchy are really meronymic relationships, in the sense that Winston, Chaffin, and Herrmann intend—we will not be accounting for the spatial aspects (or lack thereof) that may be present in some of the relationships. It is possible to distinguish them from one another using a common set of axioms, definitions, and theorems. Importantly though, we are going to add one additional meronymic relationship not in Figure 3.2 (but present in Figure 3.3) to our framework development process,

²⁹ Reproduced from Keet 2006a, p 15.

the *contained-in* relationship. We are not necessarily adding the *involved-in* relationship to the discussion in the following sections, although it might be possible to interpret it as the *functional-feature-of* relationship noted in Table 3.2 below. Note also that the *feature-of* relationship should be interpreted as a being the same relationship as the *participates-in* relationship that appears in Figure 3.3.

Table 3.2 (below) provides a complete list of both metaproperties and specific part/whole relationships. However, not all of the listed metaproperties or specific part/whole relationships are going to be useful towards our discussion of *bibliographic aggregates*. In the subsequent section, we will provide a general description of the additional metaproperties beyond those supplied by **EM** (e.g., for Individual Functional Dependence [IFD], Dense, Discrete, and the like). The section immediately following that will discuss which of the part/whole relationships listed across the top row in Table 3.2 are the most likely candidate relationships for *bibliographic aggregates*, so that we can focus our efforts both here and in Chapter 4. In the final section of this chapter, we will develop axioms and definitions specific to the subset of metaproperties that seem to be most important for drawing distinctions among our candidate part/whole relationships. These will narrow down which of the three is most likely meant to be invoked by our bibliographic standards.

However, before we move on we will address one potential problem with the framework of metaproperties illustrated in Table 3.2. Specifically, the keen-eyed reader will observe that two of the metaproperties (weak supplementation and strong supplementation) that were previously listed in Table 3.1 are no longer among the metaproperties in our conceptual framework. The primary reason for their removal is that they do not strongly figure into the analysis being carried out in

either this chapter or the one that follows. However, we do want to point out that continued controversy surrounding these axioms (Donnelly 2011, Cotnoir & Bacon 2012, Beaney 2016, Cotnoir 2018) was also a consideration. As we explain below, another metaproperty, No Partial Overlap (NPO), is actually much more applicable to the cases we are analyzing throughout this dissertation.

Table 3.2: Metaproperties Needed to Differentiate Among Part/Whole Relationships

Metaproperty	part-of	proper-part-of	member-of	sub-quantity-of	constitutes	feature-of	component-of	contained-in	functional-component-of	functional-feature-of	located-in
<i>Reflexive</i>	+	-	-	-	-	-	-	-	-	-	-
<i>Irreflexive</i>	-	+	+	+	+	+	+	+	+	+	+
<i>Symmetrical</i>	-	-	-	-	-	-	-	-	-	-	-
<i>Asymmetrical</i>	-	+	+	+	+	+	+	+	+	+	+
<i>Antisymmetrical</i>	+	+	+	+	+	+	+	+	+	+	+
<i>Transitive</i>	+	+	+	+	+	+	+	+	+	+	+
<i>IFD</i>	-	-	-	-	-	-	-	-	+	+	-
<i>Dense</i>	-	+	+	+	+	+	+	-	+	+	+
<i>Discrete</i>	-	-	+	-	-	+	+	+	+	+	+
<i>NPO</i>	-	-	-	+	+	+	+	-	+	+	-
<i>NSIP</i>	-	-	+	+	+	+	+	-	+	+	-
<i>SIS</i>	-	-	-	+	+	+	+	-	+	+	-
<i>Homeomorous</i>	-	-	-	+	-	-	-	-	-	-	+
<i>Separable</i>	-	-	+	-	-	-	+	+	+	-	-

3.3.3. General Descriptions of Metaproperties for Part/Whole Relationships

From Winston, Chaffin, and Herrmann's point of view, the distinctions laid out in Figures 3.2 and 3.3 rely on three additional metaproperties—*functionality*, *homeomerosity*, and *separability*. Formalization of these three metaproperties requires extending **EM** with additional definitions and axioms, which we do in the following sections.

Of these three, *functionality* is the most difficult to characterize and we discuss it at greater length in Appendix A. The metaproperty “IFD” stands for Individual Functional Dependence and was suggested by Vieu and Aurnague (2007). IFD is a metaproperty designed to strongly indicate when a part is functionally necessary to its whole, so much so that neither the part nor the whole can survive separation from one another. Examples of this kind of metaproperty include things like mammals and their hearts (precluding any interventions by modern medical techniques). Notice that separability is a closely linked issue here. It is so closely linked that definitions and axioms for separability and inseparability will be drawn from the literature that discusses functionality axioms like the one for IFD.

It is unclear when functionality plays a distinct role in a *bibliographic aggregate*. Some examples of things said to be parts of *bibliographic aggregates*, e.g., a table of contents in a text or a character in a story, seem to be quite functional with respect to the whole that they are part of. Conversely, things that are more typical *bibliographic aggregate* examples, e.g., the short stories in an anthology, the items in a collection, the articles in a journal article, etc., do not seem to have clear functions with regard to the whole they are part of.

Homeomerosity is a metaproperty designed to signal when a part shares virtually all or most of its most important properties with its whole. Examples include things like slices of pie, spoonfuls of soup or pudding, and tracts of land (e.g., the lot a house sits on). The metaproperty of homeomerosity does not seem like one we might be overly concerned with, unless we intend to make arguments like, the short story's text is a portion of the anthology's text or the short story's content is a portion of the anthology's content. We do not typically talk about short stories, articles, or similar bibliographic entities in this fashion. Instead, we more commonly talk about them as distinct entities in their own right.

Separability is a metaproperty designed to signal when a part can be removed from the context of the whole it is a part of, without injuring either its own identity or the identity of the whole. Some examples of separable parts include a person in a jury, a student in a class, an egg in a carton, etc. Importantly, despite the existence of examples like the aforementioned table of contents, many of the things we typically say are parts of *bibliographic aggregates*, e.g., short stories, articles, etc. seem to possess this kind of metaproperty, as we can readily see through publications like selections or data stores like article databases.

Beyond these three metaproperties are five metaproperties (denseness, discreteness, not partially overlapping, possessing no single immediate predecessor, and possessing only a single immediate successor) that specifically characterize the hierarchical structure that factors into our descriptions of parts and wholes. Bittner and Donnelly (2005) develop these metaproperties as part of their discussion of the *contained-in* and *component-of* relationships.

- **Denseness** is a metaproperty that characterizes when we (human beings) can conceptually apply some arbitrary part/whole hierarchy to a whole. Bittner and Donnelly use the example of their car, and their ability to divide it again and again into smaller and larger arbitrary portions like the front half, the back one-fifth, etc.
- **Discreteness** is a metaproperty that characterizes when an existing hierarchy of parts can easily be identified. In Bittner and Donnelly's case, they might employ the example of their car door being an easily identified, discrete part of their car. They might go on to say that the car door's window, the button for rolling the window up and down, the door's handle, etc. are all easily identified, discrete parts of their car's door.
- **Not Partially Overlapping (NPO)** is a metaproperty that indicates that none of the whole's parts overlap with other wholes. An example of in Bittner and Donnelly's case might be the fact that no two cars use the exact same door mirrors even if those mirrors are the same model of mirror.
- **No Single Immediate Predecessor (NSIP)** is a metaproperty that indicates the minimum number of parts that a whole may have, i.e., two or more. Bittner and Donnelly introduce an axiom for this metaproperty because they are not interested in cases where a whole has only a single component.
- **Single Immediate Successor (SIS)** is a metaproperty that indicates the maximum number of wholes (one) that a part may be part of. Like the preceding metaproperty, Bittner and Donnelly introduce this meta property because they are not interested in entities which share the exact same parts. Indeed, containers sharing things in them and entities with shared components seem contrary to both containers and things with components.

For our context, several of these metaproperties seem rather important. Denseness and discreteness seem particularly pertinent since we have several library catalogue-centric use cases that rely on our ability to readily identify individuals within a hierarchy, e.g., the articles in a journal issue, the items in a collection, etc., and to apply arbitrary hierarchies to those things, e.g., only articles on certain topics, only authors with certain letters in their names, etc.

The NSIP metaproperty also seems as though it will be important to us, as we probably don't want to consider *bibliographic aggregates* that only have one part. For instance, can we really say that an anthology with only one short story in it is different from a short story, novelette, or novella in and of itself? In this type of case, it seems better to err on the side of precision and name an entity as what it appears to be rather than what it might claim to name itself.

The NPO and SIS metaproperties seem less useful for our consideration here though. This is primarily because we have strong evidence of content both abstract (e.g., artistic content, propositional content, etc.) and symbolic (e.g., text, images, etc.) being shared across multiple entities. A good example of this is the HathiTrust which provides its users with a tool for fashioning their own digital collections from the objects already in the HathiTrust corpus. Once created, the objects in these user-generated digital collections are shared by both the individual user-generated digital collections and the HathiTrust's overarching corpus. It seems unlikely, then, that we can claim that *bibliographic aggregates* never overlap as metaproperties like NPO and SIS would require us to.

In the next section, we use these general insights about the metaproperties of part/whole relationships to help us determine which part/whole relationships we should focus our efforts on. Through this reduction of the part/whole relationships listed in Table 3.2, we hope to decrease the overall number of supporting axioms and definitions that need to be developed for the analysis in Chapter 4.

3.3.4. Candidate Part/Whole Relationships for Bibliographic Aggregates

As we noted in the previous sections, not all of the part/whole relationships suggested by Winston, Chaffin, and Herrmann are appropriate for our *bibliographic aggregate* modeling use case. In particular, it does not seem as though any of the part/whole relationships that possess the homeomeric metaproperty (the metaproperty of being made of the “same stuff”) is going to be helpful. While bibliographic entities like articles, short stories, and the like might be fashioned from the same kinds of “stuff,” e.g., propositions, text, etc., it usually isn’t the case that they are fashioned from the exact same propositions, text, etc. So, we will discount the *sub-quantity-of* and *located-in* relationships from our analysis.

We should note that we are not removing them from consideration because bibliographic entities like anthologies and short stories do not participate in such relationships, but because they do not speak to the *bibliographic aggregate* role we are interested in. They do seem more relevant to analyses of specific bibliographic entities, e.g., the novel *Moby Dick*, wherein each chapter might be modeled as a *sub-quantity-of* the novel’s content as a whole, a portion of the text denoting the novel’s content, a range of paper leaves on which text denoting the novel’s content is inscribed, etc. Similarly, a part/whole relationship like *located-in* can be very valuable for models of a library

collection's physical plant (i.e., the ranges of shelves upon which bibliographic entities are organized).

As we also noted in the preceding section, we are interested in those part/whole relationships which can give us a good sense of a *bibliographic aggregate*'s internal hierarchies through the fact that the bibliographic entities that *bibliographic aggregates* group together are easily identified (i.e., it is obvious what the items of a collection, the short stories in an anthology, the articles in a journal issue, etc., are. So, those part/whole relationships that necessarily possess the **discrete** metaproperty should be included among our candidate relationships. This allows us to eliminate the *constitutes* relationship, along with the more general *proper-part-of* and *part-of* relationships from consideration.

Here we need to make an important distinction. Just because a general relationship like *proper-part-of* does not necessarily possess a metaproperty like discreteness does not mean that some of its instances (or even a great many of them) do not exhibit this metaproperty as a characteristic. It is well within the realm of possibility for the *proper-part-of* relationship to possess sub-properties whose instances necessarily possess the metaproperty of discreteness. It is exactly those kinds of sub-properties (i.e., relationships) that we are interested in.

Similarly, we believe that the **separability** metaproperty is going to be a necessary one, as we have many examples where things can be removed from *bibliographic aggregates* (such as when a book is deaccessioned from a library collection) without destroying either *bibliographic aggregate* or

itself as bibliographic entities. This will allow us to remove the *feature-of* and *functional-feature-of* relationships from our group of candidate part/whole relationships.

Finally, as we previously mentioned, we do not believe that the “parts” of *bibliographic aggregates* necessarily play any specific functional role within the context of the *bibliographic aggregate* itself. Hence, we are not going to be interested in any part/whole relationships that necessarily possess the **IFD** metaproperty. Thereby, we can also eliminate the *functional-component-of* relationship from our consideration.

Table 3.3: Candidate Part/Whole Relationships for Bibliographic Aggregates

Metaproperty	member-of	component-of	contained-in
<i>Reflexive</i>	-	-	-
<i>Irreflexive</i>	+	+	+
<i>Symmetrical</i>	-	-	-
<i>Asymmetrical</i>	+	+	+
<i>Antisymmetrical</i>	+	+	+
<i>Transitive</i>	+	+	+
<i>Dense</i>	+	+	-
<i>Discrete</i>	+	+	+
<i>NPO</i>	-	+	-
<i>NSIP</i>	+	+	-
<i>SIS</i>	-	+	-
<i>Separable</i>	+	+	+

Through these eliminations, we are left with a narrower list of both part/whole relationships and metaproperties as illustrated in Table 3.3 above. In fact, we are left with just three possible part/whole relationships: *contained-in*, *component-of*, and *member-of*. This is not so surprising, since as we already noted, our existing bibliographic standards (IFLA 1998; Bekiari et al. 2016) and past analyses of *bibliographic aggregates* (Tillett 1987) all employ labels like “container relationships,” “components,” “members,” etc. in their accounts of *bibliographic aggregates*. In

the next section we will extend **EM** with additional axioms and definitions to help us analyze the extent to which these three part/whole relationships are interchangeable and where they are not, and to better identify which of them, if any, our bibliographic standards really mean to invoke.

3.3.5. Further Extensions to EM

3.3.5.1. *Immediate Parts*

As we noted, the internal hierarchy of a *bibliographic aggregate* is vitally important to us. The ability to distinguish the bibliographic entities comprised by a *bibliographic aggregate* along organic boundaries is a necessary feature of any *bibliographic aggregate* model. Unfortunately, all three of our candidate part/whole relationships also possess the metaproperty of transitivity. As Wickett, Renear, and Furner (2011) point out, this has the negative effect of collapsing any internal hierarchy. Fortunately, a solution to this problem has been proposed by Bittner and Donnelly (2005).

Bittner and Donnelly are carrying out a close analysis of the *contained-in* and *component-of* relationships in order to better distinguish them from one another. As part of their analysis, they develop five additional relationship elements as an aid in determining when something is a *structural-component-of* a whole and when it is merely *contained-in* a whole. These additional metaproperties focus on two areas: *immediate parthood* (i.e., when a part is one or at most two hierarchical steps from the whole) and parts with arbitrary boundaries (i.e., things that lack bona fide boundaries, like tracts of land).

Bittner and Donnelly are pursuing this avenue of conceptualization because they have already accepted that the part/whole relations they are examining (*component-of* and *contained-in*) are

sub-properties of the *proper-part-of* relationship. As the *proper-part-of* relation always possesses the metaproperty of transitivity, so too do the *component-of* and *contained-in* relationships (which they define to sub-properties of the *proper-part-of* relationship).

There are several important implications here for our analysis of bibliographic aggregates:

- If our bibliographic standards really mean relationships like *component-of* when they use the term “component,” then bibliographic aggregates like digital collections cannot possibly be sets like Lagoze and Fielding (1998); Gonçalves et al. (2004); or Meghini, Spyrtos, and Yang (2010) suggest.
- If this is the case, then Wickett, Renear, and Furner’s (2011) soft conclusion that collections are not sets is greatly strengthened, because the *set-member-of* relationship does not possess the metaproperty of transitivity.
- However, it might be the case that collections are actually containers as Hadro (2015) suggests.³⁰
- Bittner and Donnelly have proposed a solution to the exact problem that Wickett, Renear, and Furner faced in their analysis of collections as sets. In this case, Bittner and Donnelly are already ontologically committed to transitive part/whole relationships. They still need a tool that allows them to remark on the internal hierarchy of the parts and their whole. The *immediately-part-of* relationship accomplishes this because it is by definition an *intransitive* relationship.

³⁰ We should note though that Hadro might actually be using the terms for “collection” and “container” that are used in object-oriented programming contexts and which are, in fact, sets (Hughes 1997).

Returning to our review of Bittner and Donnelly, we see they begin by discussing a necessary metaproperty of the *contained-in* relation—*discreteness*. Regarding the *contained-in* relation, Bittner and Donnelly tell us that, “Containment structures are *discrete*” -- p 383. They give us the following example:

“[I]f x is contained in y then either x is an[mm] immediately contained in y or (a) there exists a z such that x is an[mm] immediately contained in z and z is contained in y , and (b) there exists a z such that x is contained in z and z is immediately contained in y .” -- Bittner & Donnelly (2005), p 383.

What they are aiming to show is that we know precisely what the contents of the container is. When they say that the relationship is “discrete,” they mean that we can easily articulate the hierarchy between the container and containees.

However, for this to work, Bittner and Donnelly need to define an additional relationship—*immediately-contained-in*—which is an intransitive version of the *contained-in* relation. Since we are discussing parthood relationships, the following definitions and axioms are generalized adaptations of more specific ones provided by Bittner and Donnelly.

D.7 (Immediate Parthood):

$$\forall x \forall y \left(\text{immediatelyPartOf}(x, y) =_{df} \left(\text{partOf}(x, y) \wedge \sim \exists z \left(\text{partOf}(x, z) \wedge \text{partOf}(z, y) \right) \right) \right)$$

Definition D.7 states that for all entities x and y , entity x is *immediately-part-of* entity y if and only if it is the case that entity x is *part-of* entity y , and there exists no entities z such that entity x is part

of entity z and entity z is in turn part of entity y . In other words, there are no transitive parthood relationships between a part (x) and a whole (y).

Unlike the *part-of* relationship, the *immediately-part-of* relationship is intentionally an intransitive relationship, allowing internal part/whole hierarchies to be illustrated. Bittner and Donnelly supply the following formalism to represent this metaproperty of the *immediately-part-of* relationship.³¹

$$\mathbf{A.9 \text{ (Intransitive Immediate Parthood): } \forall x \forall y \forall z \left(\left(\text{immediatelyPartOf}(x, y) \wedge \text{immediatelyPartOf}(y, z) \right) \rightarrow \sim \text{immediatelyPartOf}(x, z) \right)}$$

Axiom A.9 states that for all entities x , y , and z , if it is the case that entity x is an *immediately-part-of* entity y and entity y , in turn, is an *immediately-part-of* entity z then it cannot be the case that entity x is an *immediately-part-of* entity z .

3.3.5.2. Discrete Parts

Discreteness can now be formalized through a series of axioms that Bittner and Donnelly develop.

$$\mathbf{A.10 \text{ (Up-Discreteness): } \forall x \forall y \left(\text{partOf}(x, y) \rightarrow \left(\text{immediatelyPartOf}(x, y) \vee \exists z \left(\text{partOf}(x, z) \wedge \text{immediatelyPartOf}(z, y) \right) \right) \right)}$$

Axiom A.10 states that for all entities x and y , if entity x is *part-of* entity y , then it is the case that entity x is *immediately-part-of* entity y or there exists some entity z such that entity x is *part-of* entity z and entity z is *immediately-part-of* entity y .

³¹ We should note that they go through the steps to prove the formalism in axiom A.9 as a theorem that is partially dependent on their discreteness axioms. As we do not labor to demonstrate this here, we represent it as an axiom instead.

$$\mathbf{A.11 \text{ (Down-Discreteness):}} \forall x \forall y \left(partOf(x, y) \rightarrow \left(immediatelyPartOf(x, y) \vee \exists z (immediatelyPartOf(x, z) \wedge partOf(z, y)) \right) \right)$$

Axiom A.11 states that for all entities x and y , if entity x is *part-of* entity y then either entity x is *immediately-part-of* entity y or there exists an entity z such that entity x is *immediately-part-of* entity z and entity z is *part-of* entity y .

$$\mathbf{A.12 \text{ (Discrete Parthood):}} \forall x \forall y \left(discretelyPartOf(x, y) \rightarrow \left(UD(x, y) \wedge DD(x, y) \right) \right)$$

Axiom A.12 states that for all entities x and y , if entity x is *discretely-part-of* entity y then entity x is both **Up-Discrete (UD)** and **Down-Discrete (DD)** with respect to entity y .

What is really happening here is that the analysis is making an ontological commitment to only remark upon one level of internal hierarchy at a time. So, if a container contains several objects and several additional containers, we are committed to saying nothing about the contents of those sub-containers, until such a time as we examine them individually, as wholes in their own right. In an IR-system context, this metaproperty can be useful because it suggests the scope of the information about a whole's parts that should be communicated to an end user at any given time.

Bittner and Donnelly go on to point out that the *structural-component-of* relationship is also discrete, so their account of the differences between *structural-component-of* and *contained-in* has not yet gained much traction, as the two relationships both have the metaproperty of being *discrete*.

3.3.5.3. Fiat Parts, also known as Denseness

Bittner and Donnelly next point out that both *proper parthood* and *structural componenthood* are dense relationships, due, as they say, “to the existence of *fiat* parts (parts which lack a complete bona fide boundary) [Smith 2001]” – pp 343-4. Referencing Smith, they make an argument that containers and the things contained within them do not have *fiat* parts.

Regarding the nature of *fiat* parts, they expand their explanation with the following example:

“Consider my car and its proper parts. My car does not have an immediate proper part—What-ever proper part x we chose, there exists another slightly bigger proper part of my car that has x as a proper part.” – Bittner & Donnelly 2005, p 344.

An adaptation of the axiom Bittner and Donnelly use to formalize the *dense* metaproperty appears below.

A.13 (Dense Parthood): $\forall x \forall y (partOf(x, y) \rightarrow \exists z (partOf(x, z) \wedge partOf(z, y)))$

Axiom A.13 states that for all entities x and y , if entity x is *part-of* entity y , then there exists some entity z such that entity x is *part-of* entity z and entity z is *part-of* entity y .

What Bittner and Donnelly are trying to point out here is that some part/whole relationships are such that we can apply hierarchies of our own devising, by drawing arbitrary distinctions through the parts of the whole. This metaproperty actually reflects a common functionality that is seen in most digital libraries and many OPACs, that of refining a corpus (either resulting from browsing

or searching) through facets. At least with regard to digital collections, this seems to be a very likely metaproperty that their part/whole relationships possess.

However, there is a potential problem here. If axiom A.13 holds true then it would seem to be the case that we can slice and dice our car an infinite amount of times, almost in the same manner that one uses to estimate the area under a curve using an integral. This seems to violate some intuitions we have about parts and wholes. We consider a narrow slice of the front of their car, say the front $1/127^{\text{th}}$ of the car or something equally arbitrary that cuts through all the things we would normally consider to be a component, e.g., the headlights, the grill, the bumper, etc. so that none of these components are wholes in themselves—now we seem to have an arbitrary component which has no components. Or using the language of axiom A.13 we have invented a z which seems to have no valid x as a part. And similarly, there must be some maximum z such that it is in fact identical to y . So, it must be the case there are some, yet unspecified by Bittner and Donnelly, additional constraints on what values of z are going to be valid z 's. There must be some constraints on axiom A.13 which they do not discuss but nonetheless prevent the enterprise represented by the axiom from slipping into absurdity. Recognizing this problem, we discuss additional axioms that can be used to express these constraints during our discussion of the *denseness* metaproperty in the context of *bibliographic aggregates* in the next chapter.

3.3.5.4. *Partial Overlap*

Bittner and Donnelly (2005) next introduce a metaproperty that they call *no-partial-overlap* (*NPO*). This metaproperty is a definitive one for their *structural-component-of* property. The example they use is, “Two distinct car components share a component only if one is a

subcomponent of the other” – p 343. The following axiom is adapted from their formalization for *NPO*.

A.14 (No Partial Overlap [NPO]): $\forall x \forall y \left(\text{overlaps}(x, y) \rightarrow (x = y \vee \text{partOf}(x, y) \vee \text{partOf}(y, x)) \right)$

Axiom A.14 states that for all entities x and y , if entity x overlaps with entity y then it must be the case that entities x and y are the same entity, entity x is *part-of* entity y , or entity y is *part-of* entity x .

We showed a counter-example to this in the previous chapter, in the transitivity case that Wickett, Renear, and Furner were examining, which did not seem to correspond to other metaproperty cases (i.e., they switched from instances of relationships possessing metaproperties to situations instances of those relationships were found in). There are countless other examples such as journal articles re-published in additional venues (e.g., in the digital collection of an institutional repository) or short stories republished in new anthologies and so, it seems quite doubtful that any of the part/whole properties in our bibliographic standards are going to possess a metaproperty like the one set forth in axiom A.14.

However, Bittner and Donnelly point out that the *component-of* relationship does possess such a metaproperty. This immediately implies that when bibliographic standards documents (IFLA 1998, Bekiari et al. 2015) use a term like “component,” we should not interpret them as invoking Winston, Chaffin, and Herrmann’s *component-of* relationship. In actuality, they likely mean some other part/whole relationship than the *component-of* relationship.

3.3.5.5. Minimum Numbers of Parts & Maximum Numbers of Wholes

To better distinguish componenthood from containment, Bittner and Donnelly (2005) define two additional metaproperties, *no-single-immediate-predecessor (NSIP)* and *single-immediate-successor (SIS)*. Notably they say that the *contained-in* relationship lacks both of these metaproperties. The following axioms are adaptations of the axioms they present in their formalization of NSIP and SIS.

A mark of componenthood then is that every integral whole y has at least two components (x and z).

A.15 (No Single Immediate Predecessor [NSIP]):

$$\forall x \forall y (immediatelyPartOf(x, y) \rightarrow \exists z (immediatelyPartOf(z, y) \wedge \sim x = z))$$

Axiom A.15 states that for all entities x and y , if entity x is *immediately-part-of* entity y then there exists an entity z such that entity z is *immediately-part-of* entity y and it is the case entities x and z are not the same entity.

Again, here is a useful metaproperty that has direct implications for metadata that describes bibliographic aggregates. If our analysis of the part/whole relationships of bibliographic aggregates (in the next chapter) finds that those relationships possess this metaproperty, then we can strongly make a case that bibliographic aggregates possessing no items or only one item are not actually bibliographic aggregates at all, and so no metadata should be recorded. This can be used as a threshold for determining when something requires bibliographic description.

Bittner and Donnelly's (2005) *SIS* metaproperty is designed to constrain relationships like proper

parthood and componenthood, so that their immediate parts do not overlap with other aggregates. The following axiom is an adaptation of their formalization.

A.16 (Single Immediate Successor [SIS]): $\forall x\forall y\forall z \left((immediatelyPartOf(x, y) \wedge immediatelyPartOf(x, z)) \rightarrow y = z \right)$

Axiom A.16 states that for all entities x , y , and z , if it is the case that entity x is *immediately-part-of* both entity y and entity z , then it is the case that entities y and z are the same entities.

This is an important distinction for Bittner and Donnelly as they hold that all instances of the *structural-component-of* relationship will possess this metaproperty but only some cases of the *contained-in* relationship will possess it. They then give an example of an instance of the *contained-in* relationship that lacks the *SIS* metaproperty, saying the following:

“Consider the tool box in the trunk of my car. It is also contained in my car. My car and the trunk of my car are distinct immediate containers for my tool box.” – Bittner & Donnelly, p 344.

However, it is clear that in some instances, the *contained-in* relationship will possess the *SIS* metaproperty. Similarly, it is also clear that in many instances, the *contained-in* relationship is going to possess the *NSIP* metaproperty as well, and so neither of these metaproperties is as strongly indicative of the *componenthood* relationship as Bittner and Donnelly first hoped. In any event, Bittner and Donnelly provide enough formalization to provide further formal accounts of *structural componenthood* and *containment*.

3.3.5.6. Separable Parts

Unfortunately, axioms regarding separability rarely tell us much about separability as a phenomenon. They speak more to the inseparability of a part, which in turn, is used in discussions of that part's functionality within the context of its whole. We provide a fuller discussion of this in Appendix A. For our purposes here, we will reproduce several definitions and an axiom from Guizzardi (2005) which provide a formal basis for separable parts.

Guizzardi's first relevant definition here tells us that something is *existentially-dependent* on something else if it is necessarily the case that when one exists then so does the other. He introduces a primitive *exists* predicate to aid with the definition.

D.8 (Existential Dependence):

$$\forall x \forall y (\text{existentiallyDependentOn}(x, y) =_{df} (\text{exists}(x) \rightarrow \text{exists}(y)))$$

Definition D.8 states that for all entities x and y , entity x is *existentially-dependent* on entity y , if and only if, it is necessarily the case that the existence of entity x implies the existence of entity y .

This next definition tells us that if something is an *essential-part-of* something else, then it must be the case that something else is *existentially-dependent* on that that thing and the thing is necessarily *part-of* it.

D.9 (Essential Parthood):

$$\forall x \forall y (\text{essentialPartOf}(x, y) =_{df} (\text{existentiallyDependentOn}(y, x) \wedge \square \text{partOf}(x, y)))$$

Definition D.9 states that for all entities x and y , entity x is an *essential-part-of* entity y , if and only if, entity y is *existentially-dependent-on* entity x and it is necessarily the case that entity x is *part-of* entity y .

In some cases, parts are existentially co-dependent with their wholes. A common example in the literature is that of a brain and the person it is *part-of*. In these cases, we say that the thing is an *inseparable-part-of* the whole.

D.10 (Inseparable Parthood):

$$\forall x \forall y \left(\text{inseparablePartOf}(x, y) =_{df} \left(\text{existentiallyDependentOn}(x, y) \wedge \Box \text{partOf}(x, y) \right) \right)$$

Definition D.10 states that for all entities x and y , entity x is an *inseparable-part-of* entity y , if and only if, entity x is *existentially-dependent-on* entity y and it is necessarily the case that entity x is *part-of* entity y .

At this point, we can say that when something is *separably-part-of* a whole (e.g., like a book in a library collection or a tool in a toolbox), then it cannot be *inseparably-part-of* the whole.

A.17 (Separable Parthood): $\forall x \forall y (\text{separablyPartOf}(x, y) \rightarrow \sim \text{inseparablyPartOf}(x, y))$

Axiom A.17 states that for entities x and y , if entity x is *separably-part-of* entity y , then it is not the case that entity x is *inseparably-part-of* entity y . Rather importantly, this axiom relies on definition D.10 which tells us that *inseparable* parts are both *existentially-dependent* on the whole

and necessarily *part-of* it. Moreover, from definition D.9, we see that no separable part can also be an essential part.

3.3.6. Formal Analysis of Containment

We will consider the *contained-in* relationship as the first candidate part/whole relationship for establishing what our bibliographic standards actually mean with their disparate labels. To aid in this, we follow the adaptation of definitions and axioms from those in our extended **EM** framework. We do this in the exact same manner in which they are set forth by Bittner and Donnelly (2005, p 345) to provide a formal account of the metaproperties possessed by the *contained-in* relation. Unlike Bittner and Donnelly, we will signal the changes from the base framework’s axioms and definitions by retaining their numbering but adding the letters “CI” (for *contained-in*) to their labels.

To help us probe the internal hierarchies of containers, we will first adapt definition D.7 so that it is specific to the *contained-in* relationship.

D.7-CI (Immediate Containment):

$$\forall x \forall y \left(\text{immediatelyContainedIn}(x, y) =_{df} \sim \exists z (\text{containedIn}(x, z) \wedge \text{containedIn}(z, y)) \right)$$

Definition D.7-CI states that for all entities x and y , entity x is *immediately-contained-in* entity y if and only if there exist no entities z such that entity x is *contained-in* entity z and entity z is, in turn, *contained-in* entity y .

We should recall that the *immediately-contained-in* relationship is an intransitive relationship.

A.9-CI (Intransitive Immediate Containment):

$$\forall x \forall y \forall z \left((immediatelyContainedIn(x, y) \wedge immediatelyContainedIn(y, z)) \rightarrow \sim immediatelyContainedIn(x, z) \right)$$

Axiom A.9-CI states that for all entities x , y , and z , if it is the case that entity x is an *immediately-contained-in* entity y and entity y , in turn, is an *immediately-contained-in* entity z then it cannot be the case that entity x is an *immediately-contained-in* entity z .

Conversely, the *contained-in* relationship is a transitive relationship. So, if something is contained in a container and that container is contained in another bigger container then the first thing is also contained in the larger container.

$$\text{A.3-CI (Transitive Containment): } \forall x \forall y \forall z \left((containedIn(x, y) \wedge containedIn(y, z)) \rightarrow containedIn(x, z) \right)$$

Axiom A.3-CI states that for all entities x , y , and z , if it is the case that entity x is *contained-in* entity y and entity y is, in turn, *contained-in* entity z , then it is also the case that entity x is *contained-in* entity z .

$$\text{A.13-CI (Discrete Containment): } \forall x \forall y \left(containedIn(x, y) \rightarrow \left(immediatelyContainedIn(x, y) \vee \left(\exists z (immediatelyContainedIn(x, z) \wedge containedIn(z, y)) \wedge \exists z (containedIn(x, z) \wedge immediatelyContainedIn(z, y)) \right) \right) \right)$$

Axiom A.13-CI states that for all entities x and y , if entity x is *contained-in* entity y , then it must be the case that entity x is *immediately-contained-in* entity y , or it is the case there exists some entity z such that entity x is *immediately-contained-in* entity z . In turn, entity z is *contained-in*

entity y , and there exists some entity z such that entity x is *contained-in* entity z and, in turn, entity z is *immediately-contained-in* entity y . Note that we have expanded the Up-Discrete and Down-Discrete portions of the axiom.

Notice here that we have followed Bittner and Donnelly’s advice and have not claimed that the *contained-in* relationship is a “dense” relationship. This is the primary distinction Bittner and Donnelly draw between components of things and things in containers. That the *contained-in* relationship lacks this metaproperty implies that the part/whole relationships used in our bibliographic standards are not the *contained-in* relationship.

This is significant because it implies that when Hadro (2015) says something like, “this collection is divided into containers by country,” he does not mean to invoke the *contained-in* relationship. Rather he is using the label “container” to refer to some other part/whole relationship. This is precisely because the items in the “sub-containers” he speaks of can be further sub-divided according arbitrary criteria. In other words, we can apply an additional fiat hierarchy to them (or as Bittner and Donnelly would say, the sub-containers are dense [and not containers at all]).

A.18 (Parts contained within Wholes): $\forall x \forall y \forall z \left((properPartOf(x, y) \wedge containedIn(y, z)) \rightarrow containedIn(x, z) \right)$

Axiom A.18 states that for all entities x , y , and z , if it is the case that entity x is a *proper-part-of* entity y and entity y is *contained-in* entity z , then entity x is also *contained-in* entity z . Axiom A.17 is caused by A.3-CI which says that the *contained-in* relationship is a transitive relationship.

A.19 (Parts of Containers are Containers): $\forall x \forall y \forall z \left((\text{contained}(x, y) \wedge \text{properPartOf}(y, z)) \rightarrow \text{containedIn}(x, z) \right)$

Axiom A.19 states that for all entities x , y , and z , if it is the case that entity x is *contained-in* entity y and entity y is a *proper-part-of* entity z , then it is also the case that entity x is *contained-in* entity z . Axiom A.18 also follows from axiom A.3-CI.

Finally, we must take Winston, Chaffin, and Herrmann's stipulation regarding separability into account. We can enforce separability by adapting Axiom A.16 (separable parts) to our purposes.

A.17-CI (Containees are Separable Parts): $\forall x \forall y (\text{containedIn}(x, y) \rightarrow \text{separablyPartOf}(x, y))$

Axiom A.17-CI states that for all entities x and y , if entity x is *contained-in* entity y then it is the case that entity x is *separably-part-of* entity y .

Overall, the *contained-in* relationship seems like a good potential fit for *bibliographic aggregates*. However, its one notable lack, that not all of its instances necessarily possess the metaproperty of denseness is going to be a stumbling block. Through OPACs and digital library IR systems, it is easy to see that collection part/whole relationships, at the very least, will all exhibit the denseness metaproperty, as a great deal of IR functionality requires the ability to draw arbitrary distinctions among groups of things.

If we carefully consider other examples like journal issues or anthologies, it seems as though all of their part/whole relationships that are in the context of their being *bibliographic aggregates* also exhibit this propensity. It is easily possible to arbitrarily group and regroup the articles comprised

by a journal issue or the short stories in an anthology by arbitrary distinctions like the letters in their titles, the letters in their first author’s surname, their topical aboutness, and so on and so forth. It does not seem that any instances of part/whole relationships with respect to *bibliographic aggregates* lack the metaproperty of denseness. And so, we may remove the *contained-in* relationship from our list of candidate part/whole relationships.

3.3.7. Formal Analysis of Componenthood

We can use the same process to analyze how well the *component-of* relationship fits our bibliographic standards’ examples of *bibliographic aggregates*. The following axioms, definitions, and theorems adapt Bittner and Donnelly’s axioms, definitions, and theorems for structural componenthood (see Bittner & Donnelly [2005], pp 344-5. Like our previous case, we add additional letters (“CO” in this case) to the definition and axiom labels to differentiate them from the base *part-of* case.

Because overlapping components is not a desirable situation for Bittner and Donnelly, they begin their analysis by first adapting definition D.5 from **EM**.

D.5-CO (Definition of Overlapping Components):

$$\forall x \forall y \left(\text{hasOverlappingComponentWith}(x, y) =_{df} \exists z \left(\text{componentOf}(z, x) \wedge \text{componentOf}(z, y) \right) \right)$$

Definition D.5-CO states that for all entities *x* and *y*, entity *x* has an *overlapping-component-with* entity *y* if and only if there exists some entity *z*, such that entity *z* is a *component-of* both entities *x* and *y*.

A.3-CO (Transitive Componenthood): $\forall x \forall y \forall z \left((componentOf(x, y) \wedge componentOf(y, z)) \rightarrow componentOf(x, z) \right)$

Axiom A.3-CO is a specialization of theorem A.3 (transitive parts). Axiom A.3-CO states that for all entities x , y , and z , if it is the case that entity x is a *component-of* entity y and entity y is, in turn, a *component-of* entity z , then it is also the case that entity x is a *component-of* entity z .

Next, Bittner and Donnelly adapt definition D.7, because internal hierarchy is once again a core concern.

D.7-CO (Immediate Componenthood):

$\forall x \forall y \left(\begin{array}{c} \text{immediateComponentOf}(x, y) =_{df} \\ \left(componentOf(x, y) \wedge \sim \exists z (componentOf(x, z) \wedge componentOf(z, y)) \right) \end{array} \right)$

Definition D.7-CO states that for all entities x and y , entity x is an *immediate-component-of* entity y if and only if it is the case that entity x is a *component-of* entity y and there exists no entities z such that entity x is a *component-of* entity z and entity z , in turn, is a *component-of* entity y .

A.12-CO (Discrete Componenthood): $\forall x \forall y \left(componentOf(x, y) \rightarrow \left(\begin{array}{c} \exists z (immediateComponentOf(x, z) \wedge componentOf(z, y)) \wedge \\ \exists z (componentOf(x, z) \wedge immediateComponentOf(z, y)) \end{array} \right) \right)$

Axiom A.12-CO states that for all entities x and y , if entity x is a *component-of* entity y , then it is the case that there exists an entity z such that entity x is an *immediate-component-of* entity z and entity z , in turn, is a *component-of* entity y , and it is also the case that there exists an entity z such that entity x is a *component-of* entity z and entity z , in turn, is an *immediate-component-of* entity y .

A.9-CO (Intransitive Immediate Componenthood):

$$\forall x \forall y \forall z \left((\text{immediateComponentOf}(x, y) \wedge \text{immediateComponentOf}(y, z)) \rightarrow \sim \text{immediateComponentOf}(x, z) \right)$$

Axiom A.9-CO states that for all entities x , y , and z , if it is the case that entity x is an *immediate-component-of* entity y and entity y , in turn, is an *immediate-component-of* entity z then it cannot be the case that entity x is an *immediate-component-of* entity z . This is a reminder that while the *component-of* relationship is a transitive relationship, the *immediate-component-of* relationship is an intransitive relationship.

Here, we need to note that Bittner and Donnelly believe a distinguishing characteristic of the *component-of* relationship is that it is a dense relationship. We can showcase this by adapting axiom A.13.

A.13-CO (Dense Componenthood): $\forall x \forall y \left(\text{componentOf}(x, y) \rightarrow \exists z (\text{componentOf}(x, z) \wedge \text{componentOf}(z, y)) \right)$

Axiom A.13-CO states that for all entities x and y , if entity x is *component-of* entity y , then there exists some entity z such that entity x is *component-of* entity z and entity z is *component-of* entity y .

The next axiom has to do with overlapping components, or more to the point, the fact that Bittner and Donnelly do not see overlapping components as a possibility. It does make some sense, since even where a component might be shared between two or more wholes, it is likely playing different roles with respect to each whole.

A.14-CO (NPO Componenthood):

$$\forall x \forall y \left(\text{hasOverlappingComponentWith}(x, y) \rightarrow ((x = y) \vee \text{componentOf}(x, y) \vee \text{componentOf}(y, x)) \right)$$

Axiom A.14-CO states that for all entities x and y , if it is the case that entity x is an *overlapping-component-with* entity y , then it is the case that entities x and y are the same entity, entity x is a *component-of* entity y , or entity y is a *component-of* entity x .

A whole with only one component does not seem to be something that needs to be discussed in whole/component terms.

A.15-CO (NSIP Componenthood): $\forall x \forall y \left(\text{immediateComponentOf}(x, y) \rightarrow \exists z (\text{immediateComponentOf}(z, y) \wedge \sim(z = x)) \right)$

Axiom A.15-CO states that for all entities x and y , if entity x is an *immediate-component-of* entity y then there exists some entity z such that entity z is also an *immediate-component-of* entity y and it is the case that entities x and z are not the same entity.

The next axiom is also a consequence of there being no overlap among components and wholes.

A.16-CO (SIS Componenthood): $\forall x \forall z_n \left((\text{immediateComponentOf}(x, z_1) \wedge \text{immediateComponentOf}(x, z_2)) \rightarrow (z_1 = z_2) \right)$

Axiom A.16-CO states that for all entities x and z_n , if it is the case that entity x is an *immediate-component-of* both entities z_1 and z_2 then it is the case that entities z_1 and z_2 are the same entity.

Essentially, any particular component contributes to but a single whole.

At this point, we should recall a stipulation from Winston, Chaffin, and Herrmann’s taxonomy that components are separable from their integral wholes. We can provide an additional axiom to enforce this presence of the *separability* metaproperty.

A.17-CO (Separable Componenthood): $\forall x\forall y(\text{componentOf}(x,y) \rightarrow \sim\text{inseparableComponentOf}(x,y))$

Axiom A.17-CO states that for all entities x and y , if entity x is a *component-of* entity y , then it is also the case that entity x is not an *inseparable-component-of* entity y .

At first, the *component-of* relationship also seems to be a good candidate for explaining *bibliographic aggregates*. Recall from our previous discussions that bibliographic entities or aspects of them (e.g., *works*) overlap among *bibliographic aggregate* contexts. However, the regime of axioms and definitions that Bittner and Donnelly have developed for the *component-of* relationship specifically exclude overlapping situations. This would seem to preclude the *component-of* relationship from our list of viable candidate part/whole relationships. It would not seem to be the case that when our bibliographic standards use a label like “component,” they actually mean the *component-of* relationship. This leaves only the *member-of* relationship as a candidate explanation for how the parts of *bibliographic aggregates* relate to the *bibliographic aggregate* as a whole.

3.3.8. Formal Analysis of the Membership Relationship

We will now examine Winston, Chaffin, and Herrmann’s member-collection pairing, which we are calling the *member-of* relation. For Winston, Chaffin, and Herrmann, the key metaproperty that this relationship possesses is *separability*. But we can also apply the metaproperties developed

by Bittner and Donnelly. To further distinguish the term “collection” from the way Winston, Chaffin, and Herrmann use it and the manner in which LIS professionals are accustomed to employing it, we will use the label “collective” to invoke the former’s use cases and retain the label “collection” for the latter’s.

Let us consider two examples where the *member-of* relationship is likely to be employed:

1. The novel *Moby Dick* is a member of the University Library’s collection, and
2. The economics professor is a member of the University’s faculty.

If we consider the density metaproperty, we find that there will always be some intervening, arbitrary sub-grouping (e.g., alphabetically, by year of membership start, by age, etc.) which in turn will have additional sub-groupings (e.g., before a particular letter, after a particular year, etc.) that *Moby Dick* and the economics professor will, respectively, be members of. Therefore, the *member-of* relationship indicates that some members will lack bona fide boundaries (Smith 2001). Hence, the membership relationship possesses the *density* metaproperty.

It is also clear that either a member of a collective is an immediate member of that collective or it is a member of another collective, which itself is an immediate member of the original collective and so on. Hence, the *membership* relationship has the *discrete* metaproperty.

However, when we consider whether the *membership* relationship has the *NPO* metaproperty, we can see that it is possible for a collective’s member to be simultaneously the member of another, different collective. For example, the book *Moby Dick* can be both a member of a library’s special

collection of 19th Century novels and a member of a library’s collection of books about whaling. Similarly, the economics professor might simultaneously be a member of the faculties of the business and the mathematics departments. So, the *membership* relationship lacks the *NPO* metaproperty.

Considering whether collectives can possess only a single member, it does not seem to be the case that they do. Library collections always seem to have more than a single bibliographic resource in them, and groups like faculties always seem to have more than a single member. Hence, the *membership* relationship possesses the *NSIP* metaproperty.

But, because the *membership* relationship lacks the *NPO* metaproperty, it also lacks the *SIS* property. This makes sense since any particular member can be in more than one collective simultaneously.

We can represent the membership relationship through the following axioms and definitions. We will again add letters (“MO” in this case) to differentiate these axioms and definitions from those for the general *part-of* case.

D.5-MO (Definition of Overlapping Members):

$$\forall x \forall y \left(\text{hasOverlappingMemberWith}(x, y) =_{df} \exists z (\text{memberOf}(z, x) \wedge \text{memberOf}(z, y)) \right)$$

Definition D.5-MO states that for all entities *x* and *y*, entity *x* has an *overlapping-member-with* entity *y* if and only if there exists some entity *z* such that entity *z* is a *member-of* both entities *x* and *y*.

A.3-MO (Transitive Membership): $\forall x \forall y \forall z \left((memberOf(x, y) \wedge memberOf(y, z)) \rightarrow memberOf(x, z) \right)$

Axiom A.3-MO states that for all entities x , y , and z , if it is the case that entity x is a *member-of* entity y and entity y is, in turn, a *member-of* entity z then it is also the case that entity x is a *member-of* entity z .

D.7-MO (Immediate Membership):

$\forall x \forall y \left(\begin{array}{l} immediateMemberOf(x, y) =_{df} \\ \left(memberOf(x, y) \wedge \sim \exists z (memberOf(x, z) \wedge memberOf(z, y)) \right) \end{array} \right)$

Definition D.7-MO states that for all entities x and y , entity x is an *immediate-member-of* entity y if and only if it is the case that entity x is a *member-of* entity y and there exists no entities z such that entity x is a *member-of* entity z and entity z , in turn, is a *member-of* entity y .

A.12-MO (Discrete Membership): $\forall x \forall y \left(memberOf(x, y) \rightarrow \left(\begin{array}{l} \exists z (immediateMemberOf(x, z) \wedge memberOf(z, y)) \wedge \\ \exists z (memberOf(x, z) \wedge immediateMemberOf(z, y)) \end{array} \right) \right)$

Axiom A.12-MO states that for all entities x and y , if entity x is a *member-of* entity y , then it is the case that there exists an entity z such that entity x is an *immediate-member-of* entity z and entity z , in turn, is a *member-of* entity y , and it is also the case that there exists an entity z such that entity x is a *member-of* entity z and entity z , in turn, is an *immediate-member-of* entity y .

A.9-MO (Intransitive Immediate Membership):

$\forall x \forall y \forall z \left((immediateMemberOf(x, y) \wedge immediateMemberOf(y, z)) \rightarrow \sim immediateMemberOf(x, z) \right)$

Axiom A.9-MO states that for all entities x , y , and z , if it is the case that entity x is an *immediate-member-of* entity y and entity y , in turn, is an *immediate-member-of* entity z , then it cannot be the case that entity x is an *immediate-member-of* entity z .

Here, we need to recall that we established through two examples above that the *member-of* relationship is a dense relationship. Thereby, we need an adaptation of axiom A.13.

A.13-MO (Dense Membership): $\forall x \forall y \left(\text{memberOf}(x, y) \rightarrow \exists z (\text{memberOf}(x, z) \wedge \text{memberOf}(z, y)) \right)$

Axiom A.13-CO states that for all entities x and y , if entity x is *member-of* entity y , then there exists some entity z such that entity x is *member-of* entity z and entity z is *member-of* entity y .

We also want to establish that collectives with only one member (or no members) do not seem to be collectives at all. So, we will need an adaptation of axiom A.15.

A.15-MO (NSIP Membership): $\forall x \forall y \left(\text{immediateMemberOf}(x, y) \rightarrow \exists z (\text{immediateMemberOf}(z, y) \wedge \sim(z = x)) \right)$

Axiom A.15-MO states that for all entities x and y , if entity x is an *immediate-member-of* entity y then there exists some entity z such that entity z is also an *immediate-member-of* entity y and it is the case that entities x and z are not the same entity.

Finally, we need to represent that the members of collectives seem perfectly separable from the collective. Removing a tree from a forest neither causes the tree to stop being a tree nor the forest to stop being a forest.

A.17-MO (Separable Members): $\forall x\forall y(memberOf(x,y) \rightarrow \sim inseparableMemberOf(x,y))$

Axiom A.17-MO states that for all entities x and y , if entity x is a *component-of* entity y then it is also the case that entity x is not an *inseparable-component-of* entity y .

Here, we seem to have finally arrived at one possible explanation for the part/whole relationships in our bibliographic standards. When IFLA (1998) and Bekiari et al. (2015) use a label like “component,” the part/whole relationship they might be invoking is the Winston, Chaffin, and Herrmann’s *member-of* relationship. More broadly, it seems that we have arrived at an initial answer for research question 1a, which asks if our standards mean distinct part/whole relationships when they use terms like “container” or “component” (they do not), and research question 1b, which asks which part/whole relationship is being referred to by these labels. As we have seen in this chapter, what they really seem to mean is Winston, Chaffin, and Herrmann’s *member-of* relationship.

3.4. Chapter Summary

As we have shown, part/whole relationships have received a great deal of close study by both philosophers and computer scientists. Each part/whole term corresponds to a particular part/whole relationship which possesses particular characteristics in the form of metaproperties that can be used to distinguish it from other part/whole relationships. By examining the examples used in our

bibliographic conceptual standards, we were able to identify several metaproperties that seem to play key roles for distinguishing *bibliographic aggregates* from other kinds of aggregates. These were:

- The possibility for *bibliographic aggregates* to possess overlapping parts (such as a digital item being in more than one digital collection or a comic book issue being in more than one comic book series [such as in the case of cross-over issues]).³²
- The ability to identify existing hierarchical structures in a *bibliographic aggregate*, i.e., possession of the *discreteness* metaproperty (which also requires the concept of immediate parts).
- The ability to apply arbitrary hierarchical structures to a *bibliographic aggregate*, i.e., possession of the *denseness* metaproperty.
- The necessity that a *bibliographic aggregate* possess more than one part, i.e., the *NSIP* metaproperty.
- And, the need for parts to be separable from the *bibliographic aggregate*.

Through these metaproperties, we were able to discount the *contained-in* relationship as being a successful candidate for the part/whole relationship possessed by *bibliographic aggregates*, on account of the *contained-in* relationship precluding the possession of the *denseness* metaproperty. Therefore, *bibliographic aggregates* are not the same as containers. Similarly, we were able to discount the *component-of* relationship as being a successful candidate for the part/whole relationship possessed by *bibliographic aggregates*, because it requires the *NPO* metaproperty,

³² Marvel Comics' *Age of Apocalypse* is a good example of this multiple series membership phenomenon.

and would thereby exclude those examples of *bibliographic aggregates* that possess overlapping parts. Finally, we confirmed that the *member-of* relationship is a good candidate for the part/whole relationship possessed by bibliographic aggregates, since all of its instances seem to possess those metaproperties thought necessary to account for the various examples of *bibliographic aggregates* used in our bibliographic standards.

In the following chapter, we will be developing additional adaptations of the axioms and definitions from our established mereological framework as part of the process of analyzing the relationships used in the *mereological aggregate models*. We will specifically focus on the following relationships: DC-CAP's *is-gathered-into*, FRBR's *part-of*, and a trio of relationships defined in FRBR₀₀. Unlike the rote process we have gone through here, accepting each proposed axiom, or definition more-or-less whole cloth for the sake of developing a framework of distinguishing characteristics for part/whole relationships, we will instead interrogate (i.e., analyze) each axiom and definition with real-world examples to determine if the part/whole relationship being analyzed actually possesses a particular metaproperty or stands in a particular relationship to other part/whole relations. Where we find that strong counter-examples exist precluding one of the part/whole relationships we are examining from possessing a particular metaproperty in all of its instances, we will conclude that it does not necessarily possess that metaproperty.

4. Analysis of Mereological Aggregate Models

4.1. Chapter Overview

In this chapter, we closely examine the part/whole relationships employed in the aggregate models set forth in three high-level conceptual standards: DC-CAP, FRBR, and FRBR_{oo}. We begin with a brief discussion of the type of aggregate model all three of the mentioned conceptual standards employ—the *mereological aggregate model*. This is a model that depends on conceptualizations of parts and wholes. We argue that all the *mereological aggregate models* employed by these three conceptual standards are in fact the same *mereological aggregate model*. Further we argue that the part/whole relationships employed by these conceptual standards are either synonymous with Winston, Chaffin, and Herrmann’s *member-of* relationship or are sub-properties of it.

4.2. Mereological Aggregate Models

Recall from Chapters 1 and 2 that we want to discuss the way our bibliographic standards conceptualize *bibliographic aggregates*. The primary reason for doing this is to see what opportunities there are among the various bibliographic standards to describe *bibliographic aggregates* as first-class bibliographic entities, directly alongside more familiar bibliographic entities like books. Among other things, we want to as Jett, Fenlon, and Downie (2018) point out, capture the topicality or theme of composite bibliographic entities like collections.

While the members of the LIS and allied domains might disagree on the precise nature of a collection as a bibliographic entity (Lagoze and Fielding 1998; Lee 2000, 2005; Currall, Moss, and Stuart 2004; Gonçalves et al. 2004; Palmer 2004; Palmer et al. 2006; Galton 2010; Yeo 2012), there is general acknowledgement that they are bibliographic entities and, as such, require metadata to aid users in accomplishing the essential user tasks of finding, identifying, selecting,

and obtaining (Brack, Palmer, and Robinson 2000; Sweet and Thomas 2000; Foulonneau et al. 2005; Palmer et al. 2006). Similarly, other kinds of *bibliographic aggregates*, such as serials (Antelman 2004, Krier 2008), series (Jett et al. 2017), video games (Lee, Jett, and Perti 2015), and overarching “superworks” (Kiryakos and Sugimoto 2018, Sugimoto et al. 2018, Lee et al. 2018) also require sufficient metadata for users to find, identify, select, and obtain them.

In this chapter, we will set aside specifics of what each of the bibliographic entities named above are and accept that in general, each of them is a kind of *bibliographic aggregate*. Thereby, we will analyze the *mereological aggregate models* set forth in three specific high-level bibliographic standards (i.e., the Dublin Core Collections Application Profile [DC-CAP], FRBR, and Object-Oriented FRBR [FRBR_{OO}]) as case studies of how our standards model *bibliographic aggregates* in general. Our analysis here reexamines the initial findings we set forth in Chapter 3 in order to confirm our intuition that all three of the standards examined in this chapter really mean Winston, Chaffin, and Herrmann’s *member-of* relationship. By doing this, we will be able to demonstrate if the aggregate models used in DC-CAP, FRBR, and FRBR_{OO} are comparable with one another, or even exactly the same as one another.

One important factor for us to acknowledge is that our three conceptual model cases evolve from different intellectual traditions and are intended to provide similar metadata solutions for different computational environments. In particular, DC-CAP is a digital library-focused model that extends the Dublin Core metadata standard. As Dublin Core is by intent, designed to provide a simple, highly general “lingua franca” for communicating metadata about a wide variety of objects both digital and physical it should come as little surprise later in the chapter when we show that the DC-

CAP aggregate model generalizes beyond the digital library use cases it was initially designed to accommodate. Similarly, it is important to recall from our literature review in Chapter 2, that the FRBR conceptual model arises from the normalization of MARC records from library OPACs in the late 1980s and early 1990s. Despite this, it adopts a similar approach to modeling aggregates by employing the use of part/whole terms to define aspects of the aggregate model. Finally, we should also note that FRBRoo is the result of a harmonization effort (a kind of specialized mapping/interpretation activity) designed to produce an ontological vision of FRBR that is compliant with the worldview of the CIDOC Conceptual Reference Model (Le Boef et al. 2018).

Despite the different origins and intended use contexts of these three high-level conceptual models, all three of them employ an approach to modeling bibliographic aggregates that relies on part/whole concepts. The part/whole relationships they employ as aggregating relationships are central to our analysis in this chapter. More specifically, we will focus on Winston, Chaffin, and Herrmann's *member-of* relationship as a point of comparison for each of the aggregating relationships used in the mereological aggregate models, since we found it was a good candidate explanation for what various part/whole terms and labels used in the standards actually mean. Also recall that, in the previous chapter, we found that the *member-of* relationship necessarily possesses the following metaproperties:

- Irreflexivity – Collectives are not members of themselves.
- Asymmetry (and thereby Antisymmetry) – Collectives are not members of their individual members.
- Transitivity – If a (sub)collective is a member of a collective, then that (sub)collective's members are also members of the parent collective.

- Discreteness – A collective has an existing hierarchy of identifiable members (which may be collectives in their own right).
- Denseness – The nature of a collective is such that arbitrary internal hierarchies may be applied to it without affecting its status as a collective.
- No Single Immediate Predecessors (NSIP) – Collectives have at least two immediate members (no empty or singleton member collectives).
- Separability – The identifiable members of a collective may be separated from the collective without damaging the existence of either the collective or the former member.

The *member-of* relationship does not necessarily possess any of the following meta-properties:³³

- Reflexivity – The metaproperty of being both the subject and object of a relation. The *member-of* relationship necessarily lacks this metaproperty since it is exclusive with the *irreflexivity* metaproperty.
- Symmetry – The metaproperty of a relationship that indicates that both entities are objects of the relationship (e.g., marriage). The *member-of* relationship necessarily lacks this metaproperty since it is exclusive with the *asymmetry* metaproperty.
- Individual Functional Dependency (IFD) – It is not typically the case that collective's members serve particular functions that are dependent upon one another for their functionality (such as the heart and lungs in a living human body would be).
- Not Partially Overlapping (NPO) – Distinct collectives do not partially overlap with one another through their members. As we note above, we can show that sometimes collectives

³³ Recall also that the lack of all instances necessarily possessing a metaproperty does not imply that no instances possess that metaproperty.

do partially overlap with one another because they share members and so, it is not always the case that the members of collectives never partially overlap.

- Single Immediate Successors (SIS) – The immediate members of a collective are not also, simultaneously, immediate members of another collective. Requires that collectives do not share members.

Our goal here is to determine if these aggregating relationships have the same or similar meta-properties to the *member-of* relationship, or if their meta-properties are more similar to another of Winston, Chaffin, and Herrmann’s meronymic relationships, or, if the relationships are more similar to the more familiar *set-membership* relationship as set out in Fränkel and Bar-Hillel’s (1958) Zermelo-Fränkel (ZFC) axiomization of the *[set-]member-of* relationship. An initial comparison of the last is made to the Dublin Core Collections Application Profile’s *is-gathered-into* relationship by Wickett, Renear, and Furner (2011).³⁴

We also want to compare the aggregating relationships set forth by FRBR and FRBR_{OO}. In FRBR’s case, the relationship used is a rather broad and general sounding “part-of” relationship. However, it is narrowly scoped such that its domain and range are always the same kinds of things. Specifically, the kinds of things in FRBR aggregates are always *works*, *expressions*, *manifestations*, or *items*. And, it is always the case that the wholes and parts spoken of in FRBR are the same kind of thing (e.g., *works* are *part-of* a (larger) *work*, and so on and so forth). FRBR’s “part-of” relationship does not cross the large boundaries set by the individual *work*, *expression*, *manifestation*, and *item*. Similarly, FRBR_{OO} uses a trio of relationships (*has-component*, *has-*

³⁴ This exploration is both discussed further and extended below.

member, and *incorporates*) that narrowly link together FRBR *works* with (larger) FRBR *works* and FRBR *expressions* with (larger) FRBR *expressions*. At a first glance, from the above description of scope alone, it appears that FRBRoo aggregating relationships might be specialized cases of FRBR’s aggregating relationship.

The implications from the analysis will inform our metadata practices by clarifying which “parts” are *members-of* an aggregate and thereby require their own distinct metadata records and which “parts” fall under the purview of a different mereological relationship (e.g., *component-of*) and thereby might better be modeled as an attribute (or property) of the aggregate or one of its members.

4.3. Aggregates in the Dublin Core Collection Application Profile

4.3.1. DC-CAP Aggregate Model

The Dublin Core Collection Application Profile (DC-CAP [DCCDTG 2007]) provides a conceptual model for collections that uses a unique relation—*is-gathered-into*. The intended goal of this model is to allow metadata describing collections in general as opposed to a specific conceptualization as set forth by more technical authors like Lagoze and Fielding (1998), Gonçalves et al. (2004), or Galton (2010), or domain-specific authors such as Lee (2000, 2005), Currall, Moss, and Stuart (2004), Palmer (2004), Palmer et al. (2006), or Yeo (2012).

As the DC-CAP’s documentation remarks in its footnotes, the *is-gathered-into* relationship is directly based upon one employed in the United Kingdom’s Office for Library and Information Networking (UKOLN) collections model (Heaney 2000). Interestingly, while the DC-CAP reuses Heaney’s *is-gathered-into* relation, it never goes so far as defining it. Presumably, this implies that

those using the DC-CAP model are expected use Heaney’s definition for the *is-gathered-into* relationship (below).

“[Is-Gathered-Into:] A relationship between Items or Item-Components and a Collection specifying the manner in which Items or Components are or have been gathered into the Collection” – Heaney 2000, p 17.

Examining the definition further, we find that Heaney has defined or noted several additional things that could prove problematic should we adopt his definition. One issue is that he names several attributes for this relationship:

- accrual method
- accrual periodicity
- density
- identifier

The presence of attributes for a relationship in a conceptual model developed using Chen’s ER model (1976) is typical. But, one question it raises for us is whether or not we should understand *is-gathered-into* as a property of an entity (e.g., A has the property of being *gathered-into-B*), or if we would better understand the *is-gathered-into* relationship as a specialized kind of entity (e.g., an Event) that depends on other entities (A and B) for its existence.³⁵ In so far as Chen’s ER model

³⁵ Entire standards, such as the W3C’s Web Annotation standard have been written to expand relationships, like *annotates*, into entities of their own, such as *annotation*. See for example: <https://www.w3.org/TR/annotation-model/>

is concerned, both entities and relationships are sets of tuples (typically stored as tables),³⁶ and so, the ER diagram provided in the DC-CAP documentation (Figure 4.1 above) is not actually as informative as we might hope, as we have lingering questions regarding the conceptual connections between the tables.

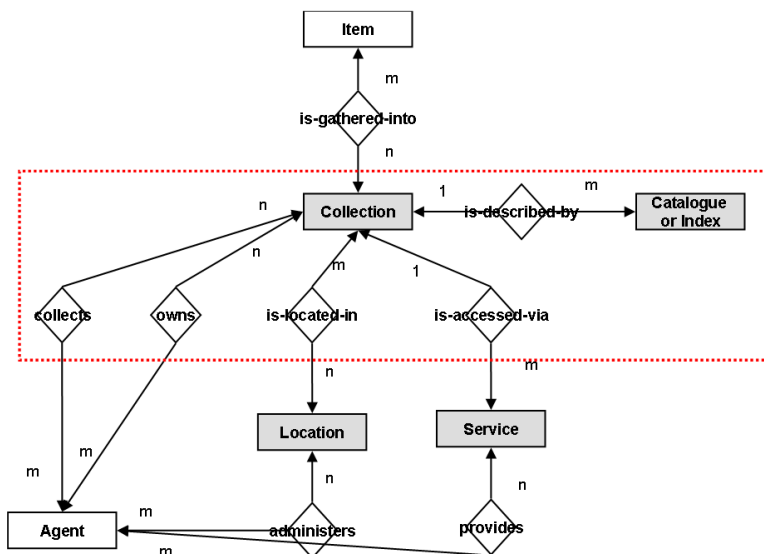


Figure 4.1: DC-CAP Data Model³⁷

The DC-CAP conceptual model attempts to sidestep this potential problem with Heaney’s definition by leaving it to implementers to specify the definition of the *is-gathered-into* relationship (assuming they don’t reference the footnote in the documentation and adopt Heaney’s definition directly). This leaves the door open to researchers like Wickett, Renear, and Furner (2011) to model *is-gathered-into* as a (more generic) relationship (sans attributes) rather than the

³⁶ A confounding factor here is that the Chen model forms a continuum of tabular relationships which run the gamut from entity to relationship with the entity-like/relationship-like “associative entity” falling in the middle. Associative entities are even symbolized by the diamond shape (typically used for relationships in diagrams that use Chen notation) bound within an entity’s box-shape.

³⁷ Image retrieved from <http://dublincore.org/groups/collections/collection-application-profile/>

relationship (set of tuples) set forth by Heaney, and thereby interrogate its nature using first-order logic. This same process also allowed Wickett et al. to develop an understanding of the *is-gathered-into* property as a sub-property of Dublin Core’s *is-part-of* relationship (2013, 2014) for the Europeana Data Model (EDM [Europeana 2016]).

4.3.2. Initial Analysis of “*is-gathered-into*” Relationship

An initial formalization of this model appears in Wickett, Renear, and Furner (2011) and is further refined in Wickett et al. (2013, 2014) and Jett et al. (2016a). We only provide a brief gloss of its formalization here. Taking the entity *collection* to be a primitive one, Wickett, Renear, and Furner provide the following formal definition that says that when something is gathered into something else, then that something else must be a collection.

$$\mathbf{D.11 \text{ (Collections Defined): } \forall x \left(\text{Collection}(x) =_{df} \exists y(\text{isGatheredInto}(y, x)) \right)}$$

Definition D.11 states that for all entities x , x is a *collection* if and only if there exists some entity y , such that entity y is *gathered-into* entity x .

From the definition that Wickett, Renear, and Furner provide us, we can see that nothing would logically prevent entity y from being a collection in and of itself. Similarly, with regard to cardinality, it would seem that nothing would prevent us from having collections consisting of only a single item. Importantly, we can see that it is also the case that *collections* are defined by the relationship *is-gathered-into*. And so, a formal understanding the scope and nature of the *is-gathered-into* relationship is also needed.

In their analysis, Wickett, Renear, and Furner set forth an initial formalization of the *is-gathered-into* relationship by matching it the Zermelo-Fraenkel (ZFC) axiomization of the *[set-]member-of* relationship by Fraenkel & Bar-Hillel (1958). ZFC narrowly scopes the membership relation's range to sets. In other words, whatever thing is a member of a collection, is also a member of a set. This is a narrower account of membership than Winston, Chaffin, and Herrmann's (1987) member-collection pairing, unless we accept that the latter's *collection* is conceptually synonymous with ZFC's notion of *set*. However, we already accepted in the previous chapter (noted above) a strongly transitive sense of Winston, Chaffin, and Herrmann's membership relation. So, we will use a slightly different label (*set-member-of*) to differentiate the ZFC-sense of membership from the one we have already established.

The task set before us is to develop a similar set of axioms, definitions, and theorems for the *is-gathered-into* relation. Wickett, Renear, and Furner (2011) tell us that “the relationship properties of *isGatheredInto* align with the relationship properties of the set theoretic relationship *memberOf*” – p 4.

Regarding their “*memberOf*” relation, they tell us:

“The set membership relationship *memberOf* is irreflexive (no set can be a member of itself), asymmetric (if x is a member of set y then y cannot be a member of x), antisymmetric (again, a trivial consequence of asymmetry) and not transitive.” – Wickett, Renear, and Furner (2011), p 4.

This allows us to take the first step in applying the framework of metaproperties set forth in the previous chapter to the *is-gathered-into* relation. We can reproduce the axioms and theorems that Wickett, Renear, and Furner use for their analysis.

Set membership is irreflexive. Or for our purposes no aggregates can be gathered into themselves.

A.6-G (Irreflexive Gathering): $\forall x(\sim isGatheredInto(x, x))$

Axiom A.6-G states that no entity x can be *gathered-into* itself.

Set membership is asymmetric. Or for our purposes nothing gathered into an aggregate, gathers the aggregate into itself.

A.20 (Asymmetric Gathering): $\forall x\forall y(isGatheredInto(x, y) \rightarrow \sim isGatheredInto(y, x))$

Axiom A.20 states that for all entities x and y , if it is the case that entity x is *gathered-into* entity y , then it is not the case that entity y is *gathered-into* entity x .

Trivially, set membership is also antisymmetric (and our aggregates too).

T.1 (Antisymmetric Gathering): $\forall x\forall y((isGatheredInto(x, y) \wedge isGatheredInto(y, x)) \rightarrow (x = y))$

Theorem T.1 follows trivially from axiom A.20. Theorem T.1 states that for all entities x and y , if it is the case that entity x is *gathered-into* entity y and it is also the case that entity y is *gathered-into* entity x , then it is the case that entities x and y are the same entity.

Finally, set membership is intransitive. Here, we have a distinction between our approach which focuses on Winston, Chaffin, and Herrmann’s *membership* relationship and the *set membership* used by Wickett, Renear, and Furner for their initial analysis of the *is-gathered-into* relation. What we are saying here is that if we treat aggregations as though they are sets, then the members of sub-collections are not members of the parent collection that gathers the sub-collection.

$$\mathbf{A.21 \text{ (Intransitive Gathering): } \forall x \forall y \left(isGatheredInto(x, y) \rightarrow \sim \exists z (isGatheredInto(x, z) \wedge isGatheredInto(z, y)) \right)}$$

Axiom A.21 states that for all entities x and y , if entity x is *gathered-into* entity y , then there exists no entity z such that entity x is *gathered-into* entity z and entity z , in turn, is *gathered-into* entity y .

Ultimately, Wickett, Renear, and Furner reject the notion that aggregates (or more specifically the collections they are examining) are the same as sets and also reject axiom A.21, instead adopting a weak notion of transitivity for collections. One of the reasons for doing this is to match our anecdotal accounts of collections wherein various arbitrary sub-collections are frequently discussed. However, in this case, transitivity is not a metaproperty that the *is-gathered-into* relationship necessarily possesses.

It is important to note here that we have already discussed axiom A.21 at some length in Chapter 2 during our initial study of Wickett, Renear, and Furner’s work and concluded that it is safe to reject axiom A.21 as we can employ additional tools (developed in Chapter 3) to help clarify the situations in which it seems as though instances of the *is-gathered-into* relationship possesses the

quality of intransitivity. However, for the sake of our analysis here, we will proceed for the time being by accepting that the *is-gathered-into* relationship is not necessarily a transitive relationship.

Table 4.1 (below) reproduces Wickett, Renear, and Furner's table comparing relationship properties. We have added a partial account of our previous understanding of Winston, Chaffin, and Herrmann's member-collective pairing (i.e., the *member-of* relation) to the table. Note that the *member-of* relation, unlike the *set-member-of* and *is-gathered-into* relations, necessarily possesses the metaproperty of transitivity. For now, we are going to set aside this distinction and focus on any other evidence that there are distinctions between the *member-of* and *set-member-of* relationships beyond transitivity.

Table 4.1: Comparison of Relationship Properties³⁸

Metaproperty	part-of	proper-part-of	member-of	set-member-of	subset-of	is-gathered-into
<i>Reflexive</i>	+	-	-	-	+	-
<i>Irreflexive</i>	-	+	+	+	-	+
<i>Symmetrical</i>	-	-	-	-	-	-
<i>Asymmetrical</i>	-	+	+	+	+	+
<i>Antisymmetrical</i>	+	+	+	+	+	+
<i>Transitive</i>	+	+	+	-	+	-

³⁸ Adapted from Wickett, Renear, and Furner (2011), p 4.

4.3.3. Homogeneity Among Collection Members

At this point, we have to admit that Winston, Chaffin, and Herrmann's member-collection pairing possesses a number of metaproperties common to both the ZFC sense of *set membership* (ZFC-*member-of*) and ZFC's *subset-of* relation. However, looking at Winston, Chaffin, and Herrmann's examples (Figure 4.2 below) we can see that they are narrowly focusing on groupings whose memberships can be represented by sets at particular points in time, but which otherwise fluctuate over time.

They enjoin us not to confuse this kind of membership with membership in a class, saying:

“Collections³⁹ must be distinguished from classes. The class-member relationship is not a meronymic relationship because it is not expressed by ‘part’ but by ‘is,’ as in,

(2d) The Nile is a river.

(2e) Fido is a dog.

Classes differ from collections in that membership in a class is determined on the basis of similarity to other members, while membership in a collection is determined on the basis of spatial proximity or by social connection.” – Winston, Chaffin, and Herrmann 1987, p 423.

³⁹ Recall from the previous chapter that we are using the label “collective” to capture the sense of “collection” that Winston, Chaffin, and Herrmann mean.

- (2a) A tree is part of a forest.
- (2b) A juror is part of a jury.
- (2c) This ship is part of a fleet.

Figure 4.2: Winston et al.'s Member-Collective Examples (1987, p 423)

Accepting that collective-membership is completely disjoint with class-membership, we must consider the two bases that Winston, Chaffin, and Herrmann set forth. Social connection seems an obvious enough connective force. As does spatial proximity, but producing axioms, definitions, and theorems to enforce them seems difficult.

We might be better served by observing that the item-members of Winston, Chaffin, and Herrmann's collective-member relationship (i.e., *member-of* in the table above) are all homogenous in some aspect of their nature. For example, if we define a forest as a group of trees then we can see that all of the forest's parts are also members of the same class of things, i.e., they are all trees. The same will be true of the ships in a fleet. So, while the forest is not the class of trees, all of its members, trees, do in fact share a kind of class membership.

Homogeneity across some particular aspect or aspects shared by the members of the collective seems to be the hallmark that differentiates the *member-of* relationship from the *set-member-of* and *subset-of* relationships, as it appears to be present in every example of collective-membership that we typically encounter.

The things in a forest are trees. The things in a jury are jurors. The things in a fleet are ships.

This homogeneity clearly extends to the kinds of aggregates this dissertation focuses on—bibliographic aggregates. For instance, the things in an anthology are all short stories. The things in a journal issue are articles. The things in a journal series are journal issues. The things in a library collection are media objects. The things in an archival collection are also media objects but are simply organized according to different principles (i.e., they are fonds). The things in a museum collection depend on the museum type, but they also form a homogeneous mass across one or more dimensions. For example, the items in a geology collection are all geological objects.

We have additional evidence that supports the idea that the items gathered into collections are all homogenous with regard to some criteria in the form of accrual and collection development policies. These policies frequently spell out, at least, the general nature of the things being gathered into a collection, if not the precise nature of a collection’s members. These same policies also provide further anecdotal evidence that collections are not sets, because they frequently also include information regarding when things cease being members of a collection (such as through deaccession policies).

We can use an axiom to help analyze this concept. What we want to say is that all the things in a collection are the same kind of thing. We can do this with the following axiom which takes the *classified-as* relationship to be primitive.

$$\mathbf{A.22 \text{ (Homogeneous Collection Membership): } \forall x \forall y \forall z \left((isGatheredInto(x, z) \wedge isGatheredInto(y, z)) \rightarrow \exists X \exists Y \left(\left((classifiedAs(x, X) \wedge classifiedAs(y, X) \wedge classifiedAs(z, Y)) \wedge \sim(x = y = z) \right) \right) \right)}$$

Axiom A.22 states that for all entities x , y , and z , if it is the case that entity x is *gathered-into* entity z and it is also the case that entity y is *gathered-into* entity z then there exists some class X such that it is the case that entities x and y are both *classified-as* class X and entity z is a *classified-as* class Y , and it is also the case that none of the entities are the same entity.

Here, we likely need to be clearer about what class Y in axiom A.22 is. Taken together, definition D.11 and axiom A.22 imply that class Y must be *Collection*.⁴⁰ We can represent this with the following axiom (note that we have started our variable assignments over from x and X , as is common practice).

$$\mathbf{A.23 \text{ (Collection is a Class): } \forall x \forall y \forall X \left(\left(isGatheredInto(x, y) \wedge \right. \right. \\ \left. \left. (classifiedAs(y, X)) \rightarrow (classifiedAs(y, X) = Collection(y)) \right) \right)}$$

Axiom A.23 states that for all entities x and y and classes X , if it is the case entity x is *gathered-into* entity y and it is also the case that entity y is *classified-as* [an instance of] class X , then it is the case that that classification is the same as claiming that entity y *is-a* *Collection*. Or in simpler words, our class X , is the class, *Collection*.

It is important to note that axiom A.22 and axiom A.23 seem to satisfy both the examples of *Collections* that Winston, Chaffin, and Herrmann are examining and also those typical of the *mereological aggregate models* being examined here. Unfortunately, the label, “Collection” is doing double duty. Winston, Chaffin, and Herrmann are using it to refer to collections in general,

⁴⁰ Here we will deviate from our use of the label “collective” and use “Collection” in the general sense that Winston, Chaffin, and Herrmann mean, as it figures predominately in the example here at the bottom of the page.

and Wickett, Renear, and Furner are using it to refer more specifically to the kinds of bibliographic collections that are the true focus here.

However, both uses suffer from the same practicalities injected by natural language usage. For instance, we sometimes speak of sub-groupings within larger groupings such that the following membership examples will also be true.

1. A grove of trees is part of a forest.
2. A squadron of ships is part of a fleet.
3. A special collection is part of a collection.

In these cases, “a grove” can easily be interpreted to mean “a tiny forest,” “a squadron” can be interpreted to mean “a small fleet of ships,” and a special collection can be interpreted to mean “a small collection.”

To some extent, this should not be a problem for axiom A.22 because there are no additional axioms that would prevent the classes X and Y from being the exact same class. Or as is implied by axiom A.23, our xs and ys might also be collections, exactly as Wickett, Renear, and Furner (2011) observe with regards to definition D.11. So, the three examples above should be accounted for by axiom A.22, exactly how it is.

However, it is frequently the case, especially with regard to collections in archives, libraries, and museums, that we describe them as gathering together both resources and smaller collections

simultaneously. For instance, the library’s collections might consist of [“adult”]-fiction, [“adult”]-non-fiction, children’s, young adult, and reference collections. In turn, the library’s reference collection may consist of the media objects in it, and additional special collections such as a microfilm collection and a genealogy collection.

We will adjust the scope of our axiom A.22 to accommodate this mixed aggregate model where both individual whole media objects are collected and aggregates like collections are collected. So, a collection may consist of both collection-members (i.e., “items”) and other collections or all of one or the other of them.⁴¹

A.22’ (Collection-Inclusive Homogeneous Collection Membership):

$$\forall x \forall y \forall z \left((isGatheredInto(x, z) \wedge isGatheredInto(y, z)) \rightarrow \right. \\ \left. \exists X \exists Y \left(\left((classifiedAs(x, X) \wedge classifiedAs(y, X) \wedge classified(z, Y)) \vee \right. \right. \right. \\ \left. \left. \left(classifiedAs(x, X) \wedge classifiedAs(y, Y) \wedge classified(z, Y) \right) \vee \right. \right. \\ \left. \left. \left(classifiedAs(x, Y) \wedge classifiedAs(y, X) \wedge classified(z, Y) \right) \wedge \sim(x = y = z) \right) \right) \right)$$

Axiom A.22’ states that for all entities x , y , and z , if it is the case that entity x is *gathered-into* entity z and it is also the case that entity y is *gathered-into* entity z , then there exists some classes X and Y such that it is the case that entities x and y are both *classified-as* class X and entity z is *classified-as* class Y , or entity x is *classified-as* class X and both entities y and z are *classified-as*

⁴¹ We should note that in some sense this revised axiom A.22’ could also be applied to sets where class X is understood to be the class of *Set-Members* and class Y is understood to be the class of *Sets*. We should also note that the first part of the preceding observation is a trivial implication of *Parthood* and mereological relationships in general, and thereby is not informative in the same way that all of the members of the collection being trees or forests (where the overarching collection type is forest) would be. This is because of the things one finds in forests that are excluded from participating in the collection ‘forest’, e.g., animals, fungi, bacteria, plants, and other things that one finds exclusively in forest settings.

class Y , or entity y is *classified-as* class X and both entities x and z are *classified-as* class Y , and it is also the case that none of the entities are the same entity.

Axiom A.22' seems to match our intuitions about collection membership—i.e., that collections can be members of (larger) collections in the exact same manner that individual items are members of collections. And so, we have an axiom that satisfies the counter-examples we raised against the initial version (axiom A.22).

4.3.4. Transitivity with Regard to Collection Membership.

Now that we have addressed both the homogenous nature of the things gathered into collections and the capacity of collections to simultaneously gather other collections into their whole, we need to more closely inspect the notion of transitivity with regard to the *is-gathered-into* relationship.

Recall from Chapter 2 that transitivity and intransitivity were a stumbling block for the analysis Wickett, Renear, and Furner carried out. Recall also that when we initially examined it, we employed an example of a common phenomenon in the library domain—the book *Moby Dick* is both a member of a library's special collection of 19th Century literature and a member of the library's collection. In turn, the library's special collection is also a member of the library's collection.

An obvious question is, if the book *Moby Dick* is a member of the library's special collection of 19th Century literature, then how is it also a member of the library's collection if it is also the case that the *is-gathered-into* relationship is not a transitive relation? We know from a stewardship perspective that the library maintains the objects that have been gathered into its collection,

regardless of ownership status and regardless of how many of those things are gathered into special collections.

A.3-G (Transitive Collection Membership): $\forall x \forall y \forall z \left((isGatheredInto(x, y) \wedge isGatheredInto(y, z)) \rightarrow isGatheredInto(x, z) \right)$

Axiom A.3-G states that for all entities x , y , and z , if it is the case that entity x is *gathered-into* entity y and it is also the case that entity y is *gathered-into* entity z , then it is the case that entity x is *gathered-into* entity z .

Similarly, we can equip ourselves to analyze the hierarchy within collections and thereby respond to one of Wickett, Renear, and Furner's arguments against the *is-gathered-into* relationship being a transitive relationship with the following definition and axiom.

D.7-G (Immediate Collection Membership):
 $\forall x \forall y \left(isImmediatelyGatheredInto(x, y) =_{df} \left(isGatheredInto(x, y) \wedge \sim \exists z (isGatheredInto(x, z) \wedge isGatheredInto(z, y)) \right) \right)$

Definition D.7-G states that for all entities x and y , entity x is *immediately-gathered-into* entity y if and only if it is the case that there exists no entity z such that entity x is *gathered-into* entity z and entity z , in turn, is *gathered-into* entity y .

A.9-G (Intransitive Immediate Collection Membership):
 $\forall x \forall y \left(isImmediatelyGatheredInto(x, y) \rightarrow \sim \exists z (isImmediatelyGatheredInto(x, z) \wedge isImmediatelyGatheredInto(z, y)) \right)$

Axiom A.9-G states that for all entities x and y , if entity x is *immediately-gathered-into* entity y , then there exists no entity z such that entity x is *immediately-gathered-into* entity z and in turn, entity z is *immediately-gathered-into* entity y .

Regarding Wickett, Renear, and Furner's second argument against transitivity for the *is-gathered-into* relationship (that it might injure the intentions of curators when they gather a whole collection rather than cherry-pick certain objects for inclusion in their collections), we can observe that in some instances curators do seem to reject gathering a whole collection in preference for gathering particular items. Indeed, the entire practice of artifact lending among museums is a strong example that this kind of preferential selection of items from collections frequently occurs. It is likely the case that where a curator has chosen to gather an entire collection into an existing collection, the inclusion does, in fact, meet a particular curatorial goal.

This is not to say that Wickett, Renear, and Furner's analysis is incorrect. Rather, these observations serve as further evidence for their conclusion that collections are not sets because they change over time. It is also the case that collections are not sets because the *is-gathered-into* relationship necessarily possesses the metaproperty of transitivity rather than intransitivity and for the same reasons as they give for assuming collections change over time—because we anecdotally speak of them as though it were the case. Thereby, the *is-gathered-into* relationship being transitive in nature matches our commonsense understanding of collections.

We should further note that nothing about the perspective on collections laid out here damages a view that collections are sets in a role and that at different times t , different sets fulfill the role of

being a particular collection. However, at this point, we must admit that the *is-gathered-into* relationship is beginning to resemble Winston, Chaffin, and Herrmann's *member-of* relation. Indeed, for the rest of this analysis, we will consider Wickett, Renear, and Furner's *is-gathered-into* relationship to be synonymous with Winston, Chaffin, and Herrmann's *member-of* relation.

4.3.5. Separable Collection Membership

Recall from the development of our analytical framework in the previous chapter that Bittner and Donnelly were concerned with the separability of both things in containers and components of things. Similarly, in library and similar cultural heritage settings we often anecdotally speak of particular members of collections (or aggregates) in terms of them being separable members from the collection as a whole.⁴² In this sub-section we will probe whether or not our intuitions about the *is-gathered-into* relationship allow this kind of affordance.

We know from actual collections in actual archives, libraries, and museums that collections change over time. Not only are members added but members are also removed. This last practice is a particular feature of circulating fiction collections in public libraries. Because physical-shelf space is limited,⁴³ older works whose copies circulate with great infrequency are often deaccessioned to provide room for new works that might circulate more frequently. In these cases, should the deaccessioned material still be in good physical condition, it is a common practice to resell to consumers. The important part to note here is that neither the former member of the collection nor the collection itself cease to exist.

⁴² Indeed, this is the whole point of Tillett's (1987) rather confusing "extractive [part/whole] relationship" in her analysis of bibliographic relationships.

⁴³ This is likely going to eventually be true for digital collections of digital resources too, because storage solutions are not infinite (or cost free).

We can probe this by defining what it might mean to be an inseparable member of a collection by adapting definition D.10 (inseparable parts) from the previous chapter.

D.10-G (Inseparable Collection Membership):

$$\forall x \forall y \left(\text{inseparablyGatheredInto}(x, y) =_{df} \left(\text{existentiallyDependentOn}(x, y) \wedge \Box \text{isGatheredInto}(x, y) \right) \right)$$

Definition D.10-G states that for all entities x and y , entity x is *inseparably-gathered-into* entity y if and only if it is the case that both entity x is *existentially-dependent-on* entity y and entity x is necessarily *gathered-into* entity y .

We can safely reject definition D.10-G on both counts. From the commonplace examples of collections and their former members, we can see that said former members are not *existentially-dependent-on* the collection. But we also note that the fact that any given member of a collection being a member of that collection is a contingent fact. It is not going to be the case that that particular collection member is a member of that particular collection in all possible worlds. Instead, we can firmly say that the members of collections are separable from those collections and represent this with the following axiom.

A.17-G (Separable Collection Membership): $\forall x \forall y (\text{isGatheredInto}(x, y) \rightarrow \sim \text{inseparablyGatheredInto}(x, y))$

Axiom A.17-G states that for all entities x and y , if entity x is *gathered-into* entity y , then it is the case that entity x is not *inseparably-gathered-into* entity y .

We might be tempted to object to axiom A.17-G. Say, for example, we have created something expressly for inclusion in our collection. We might be tempted to say that since it is intended for the collection it is necessarily gathered-into the collection. This might be true; however, it is still not the case that this new collection member is *existentially-dependent-on* the collection. The collection might cease to exist, but the former collection member remains and vice versa.

We can take this proposed counter-example a step further. Suppose that it is something that is functionally dependent on the collection for its existence, for instance, say it is the table of contents for an anthology. Even so, the short stories that comprise the anthology are still not functionally dependent upon one another.⁴⁴ Indeed, the short stories that are members of the anthology and the table of contents for that anthology actually seem to be participating in different relationships with respect to the anthology. In the former case, that relationship is clearly *is-gathered-into*; however, in the case of the table of contents, since it is functionally dependent upon the anthology and its members, another of Winston, Chaffin, and Herrmann's relationships is suggested, one that Bittner and Donnelly spend a great deal of effort refining, namely the *component-of* relation.

Moreover, it is going to be the case that we can remove particular members from an aggregate without destroying either the aggregate or the former member. Wickett, Renear, and Furner (2011) speak of this anecdotally in their conclusion, that collections are not sets. And we have previously noted that there exist deaccession policies that govern when an item is no longer a member of a

⁴⁴ We should note that we could spend some time and space here reproducing specialized versions of the definitions, axioms, and theorems that Vieu and Aurnague (2007) set forth in their mediation on **Individual Functional Dependency (IFD)**, but we would not normally ascribe particular functions to the various members of collections, especially with respect to and dependent upon, one another. That seems contrary to our everyday intuitions with respect to what membership in a collection entails.

collection. Similarly, we can observe that the particular articles in a journal issue, short stories in an anthology, or songs on an album would not necessarily cease to exist if that journal issue, anthology, or album ceased to exist, nor would those journal issues, anthologies, or albums cease to exist if particular articles, short stories, or songs did not exist. Indeed, we know from common publication practices that articles, short stories, and particularly songs, are often aggregated anew as members of other media objects or made available as individual selections (Tillett 1987).⁴⁵ And so, it necessarily seems to be the case that the members *gathered-into* collections are separable from those collections. As such, we have no reason to reject axiom A.62 and accept that collections and collection members are separable from one another.

Now that we have answered whether the members of collections can be separated from the collection, we want to return to the matters of internal hierarchy that Wickett, Renear, and Furner want to preserve. Bittner and Donnelly approach modeling internal hierarchy using two very different means: *discreteness* and *density*.

4.3.6. Discrete and Dense Collection Membership

4.3.6.1. Discrete Collection Membership

Discreteness as Bittner and Donnelly set it forth is an attempt to scope our examinations of the apparent hierarchy of elements within wholes. They examine containers and components specifically, but their approach will also work for collections. To some extent, this is because we often speak of collections using terms that imply that they are interchangeable with containers.

⁴⁵ And as evinced by the existence of the entire article database industry.

The axiom that Bittner and Donnelly provide allows for a stepwise consideration of container and component hierarchies. Put simply, it is a check to verify that the things in a container are immediate containees or immediate (sub-)containers that contain more immediate containees. We should note here that nothing prevents the containees of immediate containers from being (sub-sub)-containers themselves.⁴⁶ The axiom provided by Bittner and Donnelly only examines a single two-step of hierarchy at a time. Below is an axiom that adapts Bittner and Donnelly’s discreteness axiom for our collections case. In simple terms, it says that the things *gathered-into* a collection are either items or (sub-)collections with their own items (which very well might be sub-sub-collections).

$$\mathbf{A.12-G (Discrete Collection Membership):} \forall x \forall y \left(isGatheredInto(x, y) \rightarrow \left(immediatelyGatheredInto(x, y) \vee \exists z \left(\left(immediatelyGatheredInto(x, z) \wedge isGatheredInto(z, y) \right) \wedge \left(isGatheredInto(x, z) \wedge immediatelyGatheredInto(z, y) \right) \right) \right) \right)$$

Axiom A.12-G states that for all entities x and y , if entity x is *gathered-into* entity y , then it is the case that either entity x is *immediately-gathered-into* entity y , or there exists an entity z such that entity x is *immediately-gathered-into* entity z and entity z , in turn, is *gathered-into* entity y , and it is also the case that there exists an entity z such that entity x is *gathered-into* entity z and entity z , in turn, is *immediately-gathered-into* entity y .

⁴⁶ We should keep in mind that they also have a version of this example that works with components.

As we would only expect collection-members or (sub-)collections at any given place in a collection hierarchy, axiom A.12-G seems plausible. However, we should also note that this may imply that collections are synonymous with containers⁴⁷ since so far, other than homogeneity, collections and containers have all of the same metaproperties. This is significant because it might be the case that *collection membership* is simply a specialized version of *containment*. At this point, it might also be the case that *collection membership* is a specialized version of *componenthood*.

4.3.6.2. Dense Collection Membership

A much more difficult question to answer is whether or not Wickett, Renear, and Furner's *collection membership* relationship (*is-gathered-into*) is dense (i.e., possesses fiat parts) in the same manner that Bittner and Donnelly put forth. The goal that Bittner and Donnelly have in mind is whether or not it seems possible to apply an external and arbitrary hierarchy to a whole that has components. They use the example of a car where it is possible to divide the car arbitrarily into halves, quarters, and so on and each component of the whole (i.e., each half, quarter, so on) has the components that it has (and still functions).

Importantly, this application of arbitrary hierarchy does not work for containers. It is not the case that we can simply call the left half of a container a container in its own right. If we simply tilt the container to the right, then, with a little help from gravity, we will find that the left half has failed to contain anything at all. This is an extremely important distinction to Bittner and Donnelly who hold it up as one of the key differences between the *containment* and *componenthood* cases.

⁴⁷ Recall also that "container" is a term that Tillett (1987) invokes through her "containing [part/whole] relationship."

Beginning our argument from the same place that Bittner and Donnelly begin theirs: if we consider the members of a collection to be proper parts, then it does seem that whatever proper part x we choose from the collection, there will exist another slightly bigger proper part of the collection that has our x as a proper part too.

For example, consider a university faculty. We can apply a number of different fiat sub-groupings to the faculty. We might, for instance group faculty members alphabetically according to their surnames. In this case there might be a sub-grouping for faculty members whose names begin with the letter “A.” This sub-grouping also belongs to the sub-grouping of all faculty members whose surnames begin with the letters “A” or “B.” And this second sub-grouping belongs to the sub-grouping of all faculty members whose surnames begin with a letter from the first quarter of the alphabet, and so on, and so forth.

Similarly, say we have a digital collection mounted in a content management system. It is possible using the tuples describing the various items in the collection to formulate arbitrary queries such that arbitrary sub-collections can be fashioned out of the digital collection. In some cases, these arbitrary sub-collections might actually be gathered into a new collection for some purpose the curator has in mind. Many digital libraries, especially those that support digital humanists and similar scholars, such as the HathiTrust Research Center’s (HTRC) workset-centric⁴⁸ (Jett 2015, Jett et al. 2016a) scholar support systems, already have infrastructure that support taking advantage of our capacity to apply arbitrary hierarchies to collections.

⁴⁸ As Jett (2015) and Jett et al. (2016a) note, worksets are a specific kind of scholar-built thematic research collection.

It seems as though the *is-gathered-into* relationship is a dense relationship in the manner that Bittner and Donnelly think of dense relations. We can represent this through the following axiom.

$$\mathbf{A.13-G \text{ (Dense Collection Membership): } \forall x \forall y (isGatheredInto(x, y) \rightarrow \exists z (isGatheredInto(x, z) \wedge isGatheredInto(z, y)))}$$

Axiom A.13-G states that for all entities x and y , if entity x is *gathered-into* entity y , then it is the case that there exists some entity z such that entity x is *gathered-into* entity z and entity z , in turn, is *gathered-into* entity y .

We have already discussed many of the relevant examples that make us believe that it is possible to apply arbitrary hierarchies to collections, but digital collections provide some additional confirmation that axiom A.13-G is difficult to refute. As Lagoze and Fielding (1998) note, collection processes in digital libraries are no less reliant on the important role of *selecting* the members of the collection according to some criterion. However, an advantage that digital library collections enjoy is that it is easier to apply arbitrary hierarchies to them. This is because the metadata describing the members of digital collections (and often the members themselves) are stored in some form of (often relational) database system. Database systems, in turn, possess manifold tools for applying *additional criteria* that can be used to arbitrarily divide and sub-divide a digital collection into arbitrary sub-collections, thus making the application of arbitrary hierarchies much easier to illustrate.

However, this is not to say that axiom A.13-G is not without some important implications. First and foremost, we should interpret axiom A.13-G as remarking on something that we necessarily

require metadata to describe, unlike axiom A.12-G. Axiom A.12-G can specifically be used to give us a sense of the extent of a *bibliographic aggregate* (i.e., how many things are in the aggregate). Instead, what axiom A.13-G is intended to articulate is that an agent, like an end user can apply arbitrary hierarchies that facilitate two important activities:

1. Exploration of the aggregate's contextual mass by creating arbitrary groupings within it (e.g., grouping items by date, topic, etc.) as is demonstrably done by library OPAC faceted search features and,
2. Creating new *bibliographic aggregates* from portions of old *bibliographic aggregates* as is demonstrably done by collection building tools like those provided by the HathiTrust digital library (see <https://babel.hathitrust.org/cgi/mb?a=listcs&colltype=featured>).

Equally important is that axiom A.13-G does not exist in a vacuum from the other axioms used to describe *bibliographic aggregates*. It is true that it provides an alternate view of *bibliographic aggregates* than axiom A.12-G but, the provided view must still obey the limitations established for Winston, Chaffin, and Herrmann's *member-of* relationship (i.e., the *is-gathered-into* relationship). Of these limitations, we have already discussed one in section 4.3.3. through axiom A.22' and we discuss the other fully in section 4.3.8. below (where we will re-examine axiom A.13-G in its context).

With regards to axiom A.13-G, we can see that axiom A.22' (reproduced in the footnotes for convenience⁴⁹) implies that there is a minimal unit which is aggregated by a *bibliographic aggregate*.⁵⁰ Thereby, there are limitations to what kinds of entities can actually be included in the hierarchy.⁵¹ For example, if our class Y is the class of “anthologies” (i.e., “anthology”) and our class X is the class of “short stories” (i.e., “short story”) then it is not the case that we can arbitrarily divide the instances of the class of “short story” so that they were in effect instances of the class of “half of a short story” and so on and so forth and in effect create an infinite number of sub-anthologies out of our initial anthology.

However, we should also note that nothing would prevent us from adding additional classes to axiom A.22'. We could thusly add the class of “half of a short story” and facilitate the creation of additional combinatorials producing sub-anthologies. The limitation is that we in effect must name (and in essence, number) the classes whose instances are valid members of any putative hierarchy that we can build by employing axiom A.13-G. In practice, we might expect such an expansion of classes to be used to name similar but distinct entity types that are allowed to participate as members of a *bibliographic aggregate* rather than finer-grained portions of an already named

⁴⁹ **A.22' (Collection-Inclusive Homogeneous Collection Membership):** $\forall x \forall y \forall z \left((isGatheredInto(x, z) \wedge isGatheredInto(y, z)) \rightarrow \exists X \exists Y \left(\left((classifiedAs(x, X) \wedge classifiedAs(y, X) \wedge classified(z, Y)) \vee (classifiedAs(x, X) \wedge classifiedAs(y, Y) \wedge classified(z, Y)) \vee (classifiedAs(x, Y) \wedge classifiedAs(y, X) \wedge classified(z, Y)) \wedge \sim(x = y = z) \right) \right) \right)$

⁵⁰ There are actually a number of more general implications from this axiom for the entire mereological approach on the whole but, as we are focused on bibliographic aggregates here we do not discuss them beyond the mention in this footnote as they require a more philosophy-oriented context to discuss more fully (i.e., further discussion is simply out-of-scope here and one imagines would require an entire dissertation or similar effort to give adequate treatment to).

⁵¹ We should note that this is also an issue for Bittner and Donnelly's analysis that provides something of a counter-example since they in part are missing an axiom like axiom A.22' and in part because they do not discuss the interaction between density and the NSIP concept discussed in section 4.3.8 below.

entity type. For example, we might expand a version of axiom A.22' designed expressly to model anthologies with classes for novelettes and novellas such that an anthology comprises a combination of short stories, novelettes, novellas, and (sub)-combinations (i.e., sub-anthologies) of instances of these three classes.

Since it is always going to be the case that we can apply additional grouping criteria to a collection (within the limitations set forth by axiom A.22' and axiom A.15-G [below]), there seem to be no valid instances where we cannot apply an arbitrary hierarchy to it. Hence, we are not in a position to reject axiom A.13-G. This is important because *density* is not a metaproperty the Winston, Chaffin, and Herrmann's *containment* relationship possesses. Recall that its lack of *density* was one of the hallmarks that Bittner and Donnelly (2005) used to showcase the *containment* relationship's difference from the *componenthood* relationship. In the same way, since the *is-gathered-into* relationship does seem to necessarily evince this metaproperty then we can no longer propose that *collection membership* is a specialized version of the *containment* relation. As this is the case, and since the *is-gathered-into* relationship also necessarily possesses the metaproperty of separability which the *componenthood* relationship lacks, it is clear that the *is-gathered-into* relationship (i.e., Wickett, Renear, and Furner's *collection membership* relation) is a distinct relationship from both the *containment* and *componenthood* relationships.

4.3.7. Overlapping Collection Members

The next question to consider is whether or not the members of collections partially overlap with one another. At first, it does not seem very likely that collections have overlapping items. However, let us consider the following example provided by Valentine Charles at Europeana.

One example [of item overlap] is of an Art Nouveau vase, *Vaas met pauwen en bladeren*, delivered by the Rijksmuseum to Europeana. [*Vaas met pauwen en bladeren* is aggregated into a virtual exhibition on Art Nouveau.] As part of the data aggregation process the vase is represented as a new Europeana object, which is then featured in a Europeana Pinterest Board dedicated to Art Nouveau cultural heritage objects focusing on the theme of Nature. Later, this Pinterest Board aggregated in the [same] virtual exhibition on Art Nouveau.—Valentine Charles, 2017.⁵²

From this example, it seems clear that, especially for digital objects, it can be the case that one can gather the same item into a single collection on separate occasions. In the example above, the vase, or more precisely an image of it, is first gathered into both a virtual exhibition and a Pinterest Board. Later the Pinterest Board is gathered into the virtual exhibition. Two collections, the virtual exhibition and the Pinterest Board, partially overlap through the vase's image. It does not seem to be the case that the *is-gathered-into* relationship possesses the no-partial-overlap (NPO) metaproperty. At this point, we can also definitively say that the *is-gathered-into* relationship is a distinct relationship from the *component-of* relationship as components do not overlap with one another.

4.3.8. Immediate Predecessors and Immediate Successors

For Bittner and Connelly, the consideration of how many components a whole has and how many wholes a single component can be a part of is vital to any analysis of components and wholes. It

⁵² This text was provided as part of a still unpublished article on bibliographic aggregates that compares the FRBR and LRM models.

is important to them that anything that has components, always has at least two of them. It is equally important to them is that components are not shared by wholes. They develop two meta-properties to help them enforce these rules: **no-single-immediate-predecessor (NSIP)** and **single-immediate-successor (SIS)**. An axiom for **NSIP** appears below. In layman's terms it simply states that all collections have at least two members.

A.15-G (NSIP Collection Membership): $\forall x \forall y (immediatelyGatheredInto(x, y) \rightarrow \exists z (immediatelyGatheredInto(z, y) \wedge \sim x = z))$

Axiom A.15-G states that for all entities x and y , if entity x is *immediately-gathered-into* entity y then there exists some entity z such that entity z is *immediately-gathered-into* entity y and entities x and z are not the same entity.

If we accept axiom A.15-G, then the NSIP metaproperty would entail that all collections have at least two members. While content management systems (which treat collections as though they were sets) and some library scholars speak of collections with single (or even no) members, we wonder if such concepts have any value beyond aligning collection models with set models.

Since Wickett, Renear, and Furner (2011) note in their findings that collections are not sets, and since we have shown that the *is-gathered-into* relationship has the metaproperty of being transitive, there seems to be no advantage in accepting a view of collections that promotes single-item or empty collections.

Indeed, we know from sources like Palmer (2004) and Palmer, Zavalina, and Fenlon (2010) that researchers (and likely all users) benefit from the unified context that a collection of multiple items

brings into existence by being gathered together. Palmer and her collaborators call this unified context, *contextual mass*. We suspect that this notion of contextual mass both matches our commonsense understanding of collections and generalizes beyond the specific thematic research collections that Palmer et al. are examining.

We suspect this because of the existence of aphorisms like, “The whole is more than the sum of its parts.” It is frequently difficult to qualify, let alone quantify what the *moreness* resulting from gathering things together *is*. However, we can best preserve our commonsense intuitions about this *moreness* by accepting that the *is-gathered-into* relationship possesses the NSIP metaproperty and not rejecting axiom A.15-G.

Accepting that axiom A.15-G holds true for collectives has additional implications for axiom A.13-G. Specifically, when axiom A.15-G is true then axiom A.13-G necessarily fails when we name the collective as the sub-collective and sole member of that collective. This would also violate axiom A.6-G (collective membership is an irreflexive relationship) as the collective and its sub-collective are identical in this case. We can interpret this as a pragmatic limiting factor on the number of possible arbitrary hierarchies that can be produced through axiom A.13-G. This number is directly proportional to the number of members in a collection. In the case where a collective has only the minimum number of members required by axiom A.15-G then we can see that there are no combinatorials which are valid for axiom A.13-G. Every possible combinatorial either violates axiom A.6-G and also violates axiom A.15-G at some point in the putative hierarchy.

Thereby, the minimum number of members required for a collective before we can even begin to apply axiom A.13-G is three. If we name the members of our collective as A, B, and C then we can see it is only possible to produce four hierarchies: the given hierarchy where the collective's members are named A, B, and C (per axiom A.12-G; i.e., the discrete hierarchy) and the three putative hierarchies we can create using axiom A.13-G.

- Putative Collective 1 with members: A and sub-collective B-C (with members B and C)
- Putative Collective 2 with members: B and sub-collective A-C (with members A and C)
- Putative Collective 3 with members: C and sub-collective A-B (with members A and B)

It is possible for the number of combinatorials to grow very quickly. A collective (i.e., a *bibliographic aggregate*) like the HathiTrust digital library's corpus has (at the time of writing) a discrete hierarchy (per axiom A.12-G) of some 17 million members and thereby possesses a number putative hierarchies somewhere around 3.5×10^{49} according to axiom A.13-G (as constrained by axioms A.22' and A15-G). Obviously, where an axiom like axiom A.12-G directly informs the kind of metadata we want to record, axioms like A.13-G and A.15-G speak more towards specific functionalities we would want *bibliographic aggregate*-oriented user systems (like OPACs) to support. For instance, if a user wants to craft a digital collection of their own from a digital library's corpus they should be required to group at least two items of interest together in order to build their digital collection (so as not to violate axiom A.15-G). Similarly, we want our system to support the ability of users to build arbitrary hierarchies from the existing corpus as an aid in crafting their own digital collections (or to refine their search results, etc.) but, since it is going to be both computationally challenging and resource intensive to represent the myriad putative hierarchies that exist under the purview of axiom A.13-G we should not capture the *a*

priori metadata that describes these putative hierarchies. Instead, it is better to support recording the metadata for those instances that users determine are specifically of interest to them (and which are likely to number in much more manageable orders of magnitude than the true number of putative instances for any given corpus).

With regard to the SIS metaproperty, the example given by Valentine Charles (in Section 4.3.7 above) neatly showcases that, in fact, it can be the case, especially for digital objects, that overlap between collections that an item is gathered into can and does occur. And so, the SIS metaproperty is not one that the *is-gathered-into* relationship necessarily possesses. Since we lack clear cases showing a need for further analysis, we will not propose an axiom for SIS simply to reject.

4.3.9. Initial Findings

At this point we can produce a table (Table 4.2 below) that showcases the metaproperties possessed by the *is-gathered-into* relation. Notice that Wickett, Renear, and Furner's *is-gathered-into* relationship has all of the same metaproperties that Winston, Chaffin, and Herrmann's *member-of* relationship possesses.

From this analysis, we can conclude that not only is Wickett, Renear, and Furner's *is-gathered-into* (*collection membership*) not the same as the *set-member-of* (*ZFC set membership*) relation, but also that it is exactly the same, with regard to metaproperties, as Winston, Chaffin, and Herrmann's *member-of* (*membership*) relation.

The only clear distinction between the *is-gathered-into* and *member-of* relationships that we can potentially draw is one of scope. Winston, Chaffin, and Herrmann are speaking in general about

all types of collections (or collectives as we have been saying) while Wickett, Renear, and Furner are focused more narrowly on “curated” collections.⁵³

Table 4.2: Collection Membership (*isGatheredInto*) Relationship Metaproperties

Metaproperty	<i>member-of</i>	<i>is-gathered-into</i>
<i>Reflexive</i>	-	-
<i>Irreflexive</i>	+	+
<i>Symmetrical</i>	-	-
<i>Asymmetrical</i>	+	+
<i>Antisymmetrical</i>	+	+
<i>Transitive</i>	+	+
<i>Dense</i>	+	+
<i>Discrete</i>	+	+
<i>NPO</i>	-	-
<i>NSIP</i>	+	+
<i>SIS</i>	-	-
<i>Separable</i>	+	+
<i>Homeogeneous</i>	+	+

A key consideration is whether or not we would say that the forest example that Winston, Chaffin, and Herrmann give could be considered “curated.” At first, we might be tempted to reject the notion that a forest is actually a “curated” collection but consider the trees at the forest’s edge. There is certainly some amount of distance from some imaginary boundary after which some agent describing the forest decides which trees are within the forest and which trees are outside of the forest. Under no circumstances will it be the case that the forest will tell us such information of its own accord. Which trees are within the forest and which trees are outside of the forest is a result of a curatorial decision that some decision-making agent (i.e., a “curator”) has made.

⁵³ Their intent, and Dublin Core’s, is likely to only remark upon cultural heritage and scientific collections (including data collections), but nothing in their analysis or the Dublin Core definitions actually requires or implies that the *is-gathered-into* actually has so narrow a scope.

Hence, we conclude that all of Wickett, Renear, and Furner’s findings generalize to the cases that Winston, Chaffin, and Herrmann were examining, and the *is-gathered-into* relationship is precisely the same as the *member-of* relation.

4.4. Aggregates in the Functional Requirements for Bibliographic Entities

4.4.1. FRBR Aggregate Model

In this next section, we examine the *mereological aggregate model* set forth in IFLA’s Functional Requirements for Bibliographic Entities (FRBR) conceptual model (1998). Through our analysis we will consider whether or not FRBR’s aggregating relation, *part-of*, is actually related to the *is-gathered-into* relation.

FRBR is a high-level bibliographic conceptual model. Revised by IFLA in 2009, the FRBR standard sets forth an aggregate model that is both simple and easy to use. Unlike DC-CAP’s agnostic collection model (2007), the FRBR model clearly links its aggregate model to four primary entities—*work*, *expression*, *manifestation*, and *item*. Together, these four entities comprise FRBR’s Work-Expression-Manifestation-Item (WEMI) model. This model is a response to and an extension of the kinds of content-artifact conceptualizations that Verona (1959), Wilson (1968), and Lubetzky (1969) were all remarking upon.

An important factor for FRBR’s aggregate models is that aggregates never comprise entities that cross the boundaries between the WEMI entities. Or, as the IFLA author group of FRBR tells us:

“For the purposes of the model, entities at the aggregate or component level operate in the same way as entities at the integral unit level.” – IFLA 1998, p 29.

Thus, works compose aggregate works, expressions compose aggregate expressions, manifestations compose aggregate manifestations, and items compose aggregate items. Consequently, any analysis of FRBR's aggregates relies upon formalization of the WEMI model and its relationships if for no other reason than to demonstrate that FRBR's model for *bibliographic aggregates* actually operates at the levels of bibliographic extension proposed by WEMI. These are exactly like Tillett's "container" and "abstract" part/whole relationships (1987).

Fortunately, formalization of WEMI has already been described by Wickett and Renear (2009). They approach the formalization from what might be thought of as a constructivist point of view, rooted from the most abstract entity, FRBR's *work*, and building up to the most concrete entity, FRBR's *item*, using three relationships: *realizes*, *embodies*, and *exemplifies*.

4.4.2. Formal Account of FRBR WEMI Model

Wickett and Renear set forth the following three formal definitions (D.12-D.14) in their account of the constructed WEMI entities.

The first definition tells us that if something *realizes* a *work*, then that something is an *expression*.

D.12 (FRBR Expressions Defined): $\forall x (Expression(x) =_{df} \exists y (realizes(x, y) \wedge Work(y)))$

The second definition tells that if something *embodies* an *expression*, then that something is a *manifestation*.

D.13 (FRBR Manifestations Defined):

$$\forall x \left(\text{Manifestation}(x) =_{df} \exists y (\text{embodies}(x, y) \wedge \text{Expression}(y)) \right)$$

The third definition tells us that if something *exemplifies* a *manifestation*, then that something is an *item*.

D.14 (FRBR Items Defined): $\forall x \left(\text{Item}(x) =_{df} \exists y (\text{exemplifies}(x, y) \wedge \text{Manifestation}(y)) \right)$

Only the three relationships (*realizes*, *embodies*, and *exemplifies*) and the *work* entity are taken to be primitive by the formal analysis here.

As Wickett and Renear go on to point out, “there is an existence dependency between each defined entity type and the entity type that appears in its definition.” - [p 4]. The implication of this formalization is that while there might exist unrealized works, unembodied expressions, and unexemplified manifestations, every item has a corresponding chain of manifestation-expression-work that it is directly related to.

Wickett and Renear also formalize the one-to-many cardinality restraints set forth by the WEMI model with regard to the *realizes* and *exemplifies* relationships for their analysis (p 5). We reprise them here in axioms A.24 and A.25 below.

The first axiom tells us that an *expression realizes* exactly one *work*.

A.24 (Realization Relationship Cardinality): $\forall x \forall y \forall z \left((\text{realizes}(x, y) \wedge \text{realizes}(x, z)) \rightarrow (y = z) \right)$

The second axiom tells that an *item exemplifies* exactly one *manifestation*.

A.25 (Exemplification Relationship Cardinality): $\forall x \forall y \forall z \left((exemplifies(x, y) \wedge exemplifies(x, z)) \rightarrow (y = z) \right)$

As Wickett and Renear say of the *embodies* relation, “a many-to-many relationship may be view[ed] as unconstrained in terms of its cardinality, there is no corresponding cardinality axiom for the embodiment relationship.” -- [p 5]. The *embodies* relationship occupies a special place in the WEMI model. Its implications give rise to IFLA’s new Library Reference Model’s aggregate model (Riva, Le Bœf, and Žumer 2017). A much fuller analysis of the *embodies* relationship appears below.

Next, Wickett and Renear provide additional axioms (A.26-A.28) formalizing the domain and range of each of the three relationship primitives which I reprise here:

This first axiom states that only *expressions realize* something, and they only *realize works*.

A.26 (Domain & Range of Realization): $\forall x \forall y \left(realizes(x, y) \rightarrow (Expression(x) \wedge Work(y)) \right)$

This next axiom states that only *manifestations embody* something, and they only *embody expressions*.

A.27 (Domain & Range of Embodiment): $\forall x \forall y \left(embodies(x, y) \rightarrow (Manifestation(x) \wedge Expression(y)) \right)$

This third axiom states that only *items exemplify* something, and they only *exemplify manifestations*.

A.28 (Domain & Range of Exemplification): $\forall x \forall y (exemplifies(x, y) \rightarrow (Item(x) \wedge Manifestation(y)))$

As Wickett and Renear point out, given axioms A.26-A.28, the final clause of each of the definitions D.12-D.14 is redundant and can be removed, simplifying definitions D.12-D.14 into the definitions below (D.12', D.13', and D.14').

The first revised definition tells us that if something *realizes* anything, then that something is an *expression*.

D.12' (Expressions Redefined): $\forall x (Expression(x) =_{df} \exists y (realizes(x, y)))$

The second revised definition tells us that if something *embodies* anything, then that something is a *manifestation*.

D.14' (Manifestations Redefined):
 $\forall x (Manifestation(x) =_{df} \exists y (embodies(x, y)))$

The final revised definition tells us that if something *exemplifies* anything, then that something is an *item*.

D.15' (Items Redefined): $\forall x (Item(x) =_{df} \exists y (exemplifies(x, y)))$

Finally, Wickett and Renear provide three final axioms formalizing the assumption that the WEMI entities are disjoint from one another based upon discussions of the WEMI entities at that time (e.g., Tillet 2005). They are reprised here as axioms A.29-A.31:

The first axiom tells us that *works* are not the same as *expressions*, *manifestations*, or *items*.

$$\mathbf{A.29 \text{ (Work Disjointness): } \forall x \left(\text{Work}(x) \rightarrow \sim(\text{Expression}(x) \vee \text{Manifestation}(x) \vee \text{Item}(x)) \right)}$$

The second axiom tells us that *expressions* are not the same as *works*, *manifestations*, or *items*.

$$\mathbf{A.30 \text{ (Expression Disjointness): } \forall x \left(\text{Expression}(x) \rightarrow \sim(\text{Work}(x) \vee \text{Manifestation}(x) \vee \text{Item}(x)) \right)}$$

The third axiom tells us that *manifestations* are not the same as *works*, *expressions*, or *items*.

$$\mathbf{A.31 \text{ (Manifestation Disjointness): } \forall x \left(\text{Manifestation}(x) \rightarrow \sim(\text{Work}(x) \vee \text{Expression}(x) \vee \text{Item}(x)) \right)}$$

While this trio of axioms seems incomplete because there is no corresponding axiom for items, as Wickett and Renear point out (p. 6), if x is an *item* then as entailed by axioms A.29-A.31, x cannot be a *work*, *expression*, or *manifestation*.⁵⁴ So, an additional axiom formalizing the entailment would be redundant.

⁵⁴ It seems likely that axioms A.73 and A.74 could also be reduced for the same reasons but we do not labor to do that here.

4.4.3. Formal Analysis of FRBR Aggregates

Having accepted a particular formalization of the WEMI model, we will now produce an initial formalization of the FRBR aggregate model.

FRBR's Section 5.3.1.1 (pp 66-8) defines whole/part relationships at the WEMI *work* level. These parts come in two kinds, "dependent parts" (e.g., chapters, sections, issues, illustrations, etc.) and "independent parts" (e.g., monographs in a series, journal articles, etc.).

FRBR's Section 5.3.2.1 (p 71) defines whole/part relationships at the WEMI *expression* level. Like WEMI *works*, these parts also come in "dependent" and "independent" parts. The examples are like those used for works, but with regard to dependent parts, textual objects like tables of contents and amendments are said to also be examples of dependent parts in WEMI *expressions*.

FRBR's Section 5.3.4.1 (pp 75-6) defines whole/part relationships at the WEMI *manifestation* level. No distinction is made at this level with regard to examples but the text notes, "Physical content as represented by [a] manifestation can be divided in much the same way that intellectual content can be divided in the case of work and expression" -- p 75. Examples showcased in the text include singular volumes from a multi-volume manifestation, the soundtrack for a film on a separate medium, and the soundtrack for a film embedded in the film.

Finally, FRBR's Section 5.3.6.1 (p 78) defines whole/part relationships at the WEMI *item* level. Here again the examples do not detail any dependency or independency, like in the examples at the WEMI *work* and *expression* levels, but again the text notes that, "Parts of items can be discrete components or integral parts" -- p 78. The examples given include the physical components of a

copy and the binding of a book. Presumably, these examples encompass the after-market binding of journal issues together into bound volumes (or in some cases monographic items containing multiple volumes), which is common practice in many academic library settings to aid in their long-term preservation and shelving. Note that we don't believe that the intent was to include physical components such as the spine of a book, the end papers, etc., but the focus is on the physical section that bears inscriptions that correspond to particular units of media, such as the text in a singular volume, the soundtrack for a film, etc.

There is nothing in the text to suggest that there are whole/part relationships existing across multiple levels of the WEMI model. Any text supporting their existence is absent from the relevant sections describing Work-Expression and Manifestation-Item relationships.

It is also interesting to note that any text supporting the existence of any Expression-Manifestation relationships is absent from the FRBR document altogether. It may have been the case that the authors believed either that there are no such relationships beyond the embodiment relationship or that any such relationships were not important enough to merit inclusion in FRBR.

The representational information relating parts and wholes thereby seems to be unrelated to the representational information describing levels of intellectual abstraction (i.e., that something is a *work, expression, manifestation, or item*), beyond the fact that those relationships that we are given (on pp 66-8, 71, 75-6, and 78) are limited in scope to their specific relative levels of intellectual abstraction.

One might then begin a formal account of bibliographic aggregates by adding an additional primitive relationship to those we already have, the binary predicate *partOf*(*x*,*y*).

The account above suggests that a partial definition of aggregates might then be formalized along these lines.

D.15 (FRBR Aggregate Definition): $\forall x \left(\text{Aggregate}(x) =_{df} \exists y (\text{partOf}(y, x) \wedge \sim(x = y)) \right)$

Definition D.15 states that for all entities *x*, *x* is an *aggregate*, if and only if, there exists an entity *y* such that entity *y* is *part-of* entity *x* and it is not the case that entities *x* and *y* are the same entity.

A potential problem with definition D.15 is that *part-of* as FRBR sets forth might be confused with more generic mereological *parthood*. However, we know from the examples from the sections of FRBR cited above, that in actuality, we are dealing with a narrowly-scoped set of aggregates. Specifically, we are examining anthologies, collections, journal issues, and bibliographic series, or more precisely, we are examining *bibliographic aggregates*.

Since all of our examples are going to be the same, we have reason to suspect that FRBR's *part-of* relationship (hereafter referred to as *bibliographic-part-of* [*bPartOf*] to avoid confusion with general mereological *parthood*) possesses the same or similar metaproperties to the *is-gathered-into* relationship (a.k.a. Winston, Chaffin, and Herrmann's *member-of* relation).

We know from our examples that we would not normally say that an anthology is *bibliographically-part-of* itself, nor would we say that an anthology is *bibliographically-part-of* any of the short stories that form its parts. Hence, the *bibliographically-part-of* relationship seems to be both *irreflexive* and *asymmetric*.

We know from examples of bibliographic series that the articles that are *bibliographically-part-of* a journal issue are also regarded as *bibliographically-part-of* a journal volume and the journal series as a whole. The same is true from the short stories that are *bibliographically-part-of* an anthology which in turn is *bibliographically-part-of* a monograph series. Jim Butcher's anthology, *Side Jobs*, is an example of this, as it and the short stories in it are both part of Butcher's *Dresden Files* series. Hence, the *bibliographically-part-of* relationship seems to be *transitive*.

Similarly, we have already established that *bibliographic aggregates* both already have established hierarchies of identifiable members and can have arbitrary hierarchies applied to them. And so, the *bibliographically-part-of* relationship seems to be both *discrete* and *dense* in the same manner the *is-gathered-into* relationship was.

As it does not seem to be the case that we call individual short stories anthologies, journal articles journal issues, songs albums or items collections, there do not seem to be any instances where bibliographic aggregates are empty or possess only one member and so, the *bibliographically-part-of* relationship seems to possess the *NSIP* metaproperty.

All of the *bibliographic parts* of our *bibliographic aggregates* are such that they do not necessarily rely on the existence of the bibliographic aggregate for their own existence. Short stories, journal articles, songs, and bibliographic objects all exist on their own. Since this is the case, it seems that the *bibliographically-part-of* also possesses the metaproperty of *separable*.

Finally, all of the *bibliographic parts* of *bibliographic aggregates* are all part of their respective wholes for some curatorial (or editorial) reason. They are all carefully selected (and in some cases arranged) according to a curatorial (or editorial) criterion, which in many cases is set forth, e.g., through an accrual policy for instance. Thereby, it seems that the *bibliographically-part-of* relationship possesses the metaproperty of being *homogeneous* as we have defined it in the preceding section.

At this point, the *bibliographically-part-of* relationship resembles the *is-gathered-into* relationship to a great extent. However, from FRBR's text we see that its scope is much narrower. The homogeneous entities being aggregated are limited to WEMI's *works, expressions, manifestations, and items*. Let us propose then that the relationship we are speaking of is actually a narrower sense of Wickett, Renear, and Furner's *is-gathered-into* relation.

We will begin our formalization by noting that it is a sub-class of the *is-gathered-into* relation.

A.32 (Bibliographic Parthood as a kind of Collection Membership):

$$\forall x \forall y (bPartOf(x, y) \rightarrow isGatheredInto(x, y))$$

Axiom A.32 states that for all entities x and y , if entity x is *bibliographically-part-of* entity y then it is also the case that entity x is *gathered-into* entity y . One merit of this axiom is that we can reuse the axioms we have already developed for our analysis of the *is-gathered-into* relation.

For instance, we can adapt axioms A.26-A.28 to communicate the narrower domain and range of the *bibliographically-part-of* relation.

$$\mathbf{A.33 \text{ (Domain \& Range of Bibliographic Parthood): } \forall x \forall y \left(bPartOf(x, y) \rightarrow \right. \\ \left. \left((Work(x) \wedge Work(y)) \vee (Expression(x) \wedge Expression(y)) \vee \right. \right. \\ \left. \left. (Manifestation(x) \wedge Manifestation(y)) \vee (Item(x) \wedge Item(y)) \right) \right)$$

Axiom A.33 states that for all entities x and y , if entity x is *bibliographically-part-of* entity y then it is also the case that entities x and y are both *works*, both *expressions*, both *manifestations*, or both *items*.

This constrained primitive can be substituted back into definition D.15 to provide a refined version of it (and here we introduce a new label for aggregates in our context to differentiate them from general mereological cases—“*bAggregate*” [bibliographic aggregate]):

$$\mathbf{D.15' \text{ (Bibliographic Aggregates Defined):}} \\ \forall x \left(bAggregate(x) =_{df} \exists y (bPartOf(y, x) \wedge \sim(x = y)) \right)$$

D.15' states that for all entities x , entity x is a *bibliographic-aggregate*, if and only if there exists an entity y such that entity y is *bibliographically-part-of* entity x and it is not the case that entities x and y are the same entity.

This allows us to produce another table (Table 4.3 below). We will add an additional row to capture the narrow domain and range of the *bibliographically-part-of* relation, calling it “WEMI-Homogenous.” For clarity’s sake we will remove, with the exception of the *is-gathered-into* relation, the columns denoting Winston, Chaffin, and Herrmann’s (1987) meronymic relationships.

Our conclusion here hinges upon whether or not we suspect that there are any reasons to reject axiom A.32, which states that the *bibliographic-part-of* relationship is a sub-property of the *is-gathered-into* relation. One potential argument against accepting axiom A.32 is on the grounds that some works are dependent works (i.e., we will also reject axiom A.17-G regarding *separability*). Common examples are (so-called) *expression-level* features of works like indices and tables of contents that are commonly held to be dependent on their parent work for their existence.

The problem with this line of thought is that, as a text-free conceptual entity, *works* are not likely to be the kind of entity that one can compose an index or table of contents for and so these *expression-level* features are hardly dependent on the *work* for their existence. If anything, they would be dependent on a *work’s* corresponding *expression* for their existence.

Table 4.3: Bibliographic Parthood (*bPartOf*) Relationship Metaproperties

Metaproperty	<i>is-gathered-into</i>	<i>bibliographically-part-of</i>	<i>ZFC-member-of</i>	<i>subset-of</i>
<i>Reflexive</i>	-	-	-	-
<i>Irreflexive</i>	+	+	+	+
<i>Symmetrical</i>	-	-	-	-
<i>Asymmetrical</i>	+	+	+	+
<i>Antisymmetrical</i>	+	+	+	+
<i>Transitive</i>	+	+	-	+
<i>Dense</i>	+	+	+	+
<i>Discrete</i>	+	+	-	-
<i>NPO</i>	-	-	+	+
<i>NSIP</i>	+	+	+	+
<i>SIS</i>	-	-	+	+
<i>Separable</i>	+	+	-	-
<i>Homeogeneous</i>	+	+	-	-
<i>WEMI-Homogeneous</i>	-	+	-	-

As we know from digital documents and various publication practices (especially among academic journals), both indices and tables of contents can be published separately from the bibliographic entities they remark upon. And so, it is also not the case that we have any reason to reject axiom A.17-G (*separability*) from our account of the *bibliographically-part-of* relation.

Since bibliographic aggregates in the FRBR sense are primarily concerned with the gathering of the same kinds of things that Wickett, Renear, and Furner were looking at in their examples, and none of the counter-examples seem to negate our assumptions, we propose that there is no reason to reject axiom A.32 and accept that FRBR's *bibliographically-part-of* relationship is, in fact, just a narrower version of the *is-gathered-into* relation.

This greatly simplifies matters as the kinds of aggregates we are examining becomes quite focused and narrow. We are limited to things like library collections, archival collections, journal issues, monograph series, music compilations, and the like. In the specific case of FRBR, we are also somewhat concerned with their representation at the various WEMI models. But since the FRBR document has aggregate models at each corresponding level, there is no reason to suspect that any unusual problems will occur when selecting and recording metadata for their representation, at least not in the general case. When we begin to distinguish the different kinds of bibliographic aggregates from one another, we expect that FRBR's WEMI model will break down as the FRBR-WGA alludes to in their final report (2011) and as has been shown to be the case for other kinds of bibliographic entities (Jonsson 2005, Miller and Le Bœuf 2005, Nicolas 2005, Baca and Clarke 2007, McDonough et al. 2010, Lee et al. 2012).

4.5. Aggregates in Object-Oriented FRBR [FRBR_{OO}]

4.5.1. FRBR_{OO} Aggregate Model

In this final section, we examine the aggregate model set forth by the FRBR harmonization effort for CIDOC's Conceptual Reference Model (CIDOC-CRM)—the Object-Oriented FRBR (FRBR_{OO}) ontology (Bekiari et al. 2015). Oddly, Bekiari et al. do not employ CIDOC-CRM's existing E78 (Collection) entity for the purposes of modeling aggregates in FRBR_{OO}. We speculate that this might be because E78 (Collection) is defined in the following manner:

“This class comprises aggregations of instances of E18 Physical Thing that are assembled and maintained (“curated” and “preserved,” in museological terminology) by one or more instances of E39 Actor over time for a specific purpose and audience, and according to a particular collection development plan.”

– Le Bœuf et al. 2018, p 166.⁵⁵

The implication here is that Collections in the CIDOC-CRM sense are made up physical objects. Since after-market binding of journal issues into book format for better long-term shelving and storage is an extremely common practice in (especially academic) library settings, E78's absence from FRBR_{OO} is all the more remarkable. Given increasing evidence that metadata representing after-market bindings plays a valuable role in statistical analysis of digitized text, the absence of a ready method for capturing metadata at the FRBR *item*-level demonstrates at least one advantage that the FRBR aggregate model we analyzed in the previous section enjoys over the FRBR_{OO} aggregate model.

⁵⁵ We should note that as of the 6th version of CIDOC-CRM, E78's label has been changed from “Collection” to “Curated Holding.”

We might also remark that one easily extendable method to produce an aggregate model for FRBR-compliant bibliographic aggregates has been passed over.

Confusing the matter are the kinds of things which might be aggregates from FRBR₀₀'s point of view. The FRBR₀₀ standard has several candidate aggregate entities which include:

- F1 (Work) – can comprise multiple works
- F2 (Expression) – can comprise multiple expressions
- F14 (Individual Work) – can comprise multiple works
- F15 (Complex Work) – always comprises multiple works
- F16 (Container Work) – always comprises multiple works
- F17 (Aggregation Work) – representing the intellectual effort of aggregating⁵⁶
- F18 (Serial Work) – always comprises multiple works
- F19 (Publication Work) – can comprise multiple works
- F20 (Performance Work) – can comprise multiple works
- F22 (Self-Contained Expression) – can comprise multiple expressions
- F25 (Performance Plan) – can comprise multiple expressions

Examining this list of potential aggregate entity types, we should perhaps not be surprised that E78 (Collection) was excluded. It would seem that the editors of FRBR₀₀ believe that only bibliographic aggregates at the FRBR *work* and FRBR *expression* levels need to be modeled.

⁵⁶ Included in this list because it uses the word “aggregation.”

We have some reason to believe then, that the aggregation relationships that FRBR₀₀ sets forth narrower versions of the *is-gathered-into* relationship and perhaps even of the *bibliographic-part-of* relationship analyzed in the previous section. After all, we are once again limited to things like anthologies, collections, journal issues, monograph series, music compilations, and the like.

As FRBR₀₀ provides a much narrower account of bibliographic aggregates (hereafter referred to as the FRBR₀₀ account) several of the candidate entities (F1 and F14) can seemingly be dismissed out of hand since the focus in the FRBR₀₀ documentation is focused aggregating FRBR's *Works* into aggregate FRBR *Expressions*. Despite this, several candidate entities remain at the work level.

4.5.2. Formal Analysis

The text of FRBR₀₀ is silent with regard to the specifics of F2 (Expression) acting as an aggregate entity. It does, however, provide details on aggregation through the last candidate entity F22 (Self-Contained Expression). Here, an example from FRBR₀₀ is helpful. “The Italian text of Dante’s textual work entitled ‘Divina Commedia’ (F22) R5 has component the Italian text of Dante’s textual work entitled ‘Inferno’ (F22)” – p 74. The implication is that a specialized relationship *hasComponent* (R5) is used in cases where the aggregate Expression realizes comprises multiple individual “components.” We can formally analyze this through an additional axiom (A.34).

$$\mathbf{A.34 \text{ (Expression Componenthood): } \forall x \forall y \left(hasComponent(x, y) \rightarrow \left(SelfContainedExpression(x) \wedge SelfContainedExpression(y) \right) \right)}$$

Axiom A.34 states that for all entities x and y , if entity x *has-component* entity y then it is also the case that entities x and y are both *self-contained expressions*.

However, this is a very different sense of componenthood than we have hitherto postulated. Using our pre-existing notion of componenthood (see Chapter 3), we would normally understand it to be the case that the symbols that comprise words are (functional) components of an expression (vis-à-vis, expressions are textual objects). But it is not typically the case that any arbitrary symbol removed from its context can be said to realize a work, in whole or in part. So, the symbols that comprise the words that form an expression are not expressions in and of themselves. Hence, FRBR_{OO}'s R5 (hasComponent) relationship is not a *componenthood* relationship at all.

Going back to the example that FRBR_{OO}'s authors give us, it is clear that the R5 relation's purpose is much the same as FRBR's *bibliographic-part-of* relation. The R5 relationship groups together expressions into composite expressions where the author of a work has intended that they be grouped together. Since the R5 relationship behaves in a similar manner to FRBR's *bibliographic-part-of* relation, it is extremely likely that it evinces all of the same metaproperties that FRBR's *bibliographic-part-of* possesses. To avoid confusion between FRBR_{OO}'s R5 relationship and our previously established formalization of the *has-component* relation, we will coin a new label for FRBR_{OO}'s R5 relation—*has-expression-component*.⁵⁷

Just by examining the example of Dante's *Divine Commedia*, we can see that R5 is both irreflexive and asymmetric, as it is not the case that Dante's *Divine Commedia* is a component of itself, nor is it the case that Dante's *Inferno* *has-expression-component* Dante's *Divine Commedia* (as it is the opposite case that would be true). Hence the R5 relationship is both *irreflexive* and *asymmetric*.

⁵⁷ Note though, we are only continuing to use the "component" label because FRBR_{OO} uses it. The relationship itself is quite probably a version of Winston, Chaffin, and Herrmann's *member-of* relationship and not the *component-of* relationship at all.

Dante’s *Inferno* is itself divided into 34 cantos. Since each canto is itself a unique block of content, we might refer to them directly as expression-components of Dante’s *Divine Commedia* (e.g., Dante’s *Divine Commedia has-expression-component Dante’s Inferno Canto X*). As in our previous cases, we have no reason to believe that R5 is not **transitive**. Since there is a readily identifiable entity hierarchy in the aggregation, R5 is also **discrete**. We can also apply arbitrary groupings (as has been done in the Wikipedia article describing the Dante’s *Inferno*),⁵⁸ and so R5 is also **dense**. As we noted, we can easily separate the various parts of the *Divine Commedia* from one another without damaging their individual existence or the existence of Dante’s *Divine Commedia* itself, and so R5 is **separable**. Finally, this entire discussion would make little sense if there were not at least 2 portions of the *Divine Commedia* for us to link together with R5, thereby we should treat R5 as **NSIP**.

R5 is also clearly **homogeneous** since it is limited in scope with regard to its domain and range. In fact, the R5 relationship possesses a narrower sense of homogeneity than FRBR’s *bibliographic-part-of* relation. We can express this narrower sense of homogeneity through the following axiom.

A.35 (Domain and Range of R5): $\forall x \forall y \left(hasExpressionComponent(x, y) \rightarrow (Expression(x) \wedge Expression(y)) \right)$

Axiom A.35 states that for all entities x and y , if entity x *has-expression-component* entity y , then it is also the case that both entities x and y are *expressions* (in the FRBR sense).

⁵⁸ [https://en.wikipedia.org/wiki/Inferno_\(Dante\)](https://en.wikipedia.org/wiki/Inferno_(Dante))

Because axiom A.35 is a narrower version of axiom A.33, and because FRBR_{oo}'s *has-expression-component* relationship possesses all of the same metaproperties as FRBR's *bibliographic-part-of*, it follows that *has-expression-component* is actually a narrower sense of *bibliographic-part-of*. In turn, because *bibliographic-part-of* is a narrower sense of *is-gathered-into*, it trivially follows from A.32, that *has-expression-component* is also a narrower sense of *is-gathered-into*. We can express these outcomes through the following theorems.

T.2 (Expression Componenthood as Bibliographic Parthood):

$$\forall x \forall y (hasExpressionComponent(x, y) \rightarrow bPartOf(y, x))$$

Theorem T.2 trivially follows from the combination of axioms A.32 and A.35. Theorem T.2 states that for all entities x and y , if entity x *has-expression-component* entity y , then it is also the case that entity y is *bibliographically-part-of* entity x .

T.3 (Expression Componenthood as Collection Membership):

$$\forall x \forall y (hasExpressionComponent(x, y) \rightarrow isGatheredInto(y, x))$$

Theorem T.3 trivially follows from the combination of axiom A.32 and theorem T.2. Theorem T.3 states that for all entities x and y , if entity x *has-expression-component* entity y , then it is also the case that entity y is *gathered-into* entity x .

If we accept that axioms A.34 and A.35 and theorems T.2 and T.3 are true, then it follows that one view of FRBR_{oo}'s sense of bibliographic aggregates might be formalized via the following definition.

D.15-EC (Bibliographic Aggregates defined as Aggregate Expressions):

$$\forall x (bAggregate(x) =_{df} \exists y (hasExpressionComponent(x, y) \wedge \sim(x = y)))$$

Definition D.15-EC states that for all entities x , it is the case that entity x is a *bibliographic aggregate*, if and only if there exists some entity y such that entity x *has-expression-component* entity y and entities x and y are not the same entity.

Definition D.15-EC suffices to account for *aggregate expressions* in the FRBR₀₀ sense. However, additional candidate aggregate entities exist at FRBR's *work* level. The next one for this analysis is FRBR₀₀'s Complex Work (F15) entity.

The Complex Work entity is the complement to the aggregate expression that has just been described. It links an aggregate work to other works that comprise it. It links very closely to the kind of aggregate expression that is defined in definition D.15-EC in a complementary way. Here again we are given the example of Dante's *Inferno*. In Bekiari et al.'s own words, "Dante's textual work entitled 'Divina Commedia' (F15) R10 has member Dante's textual work entitled 'Inferno' (F15)" – pp 76-7. The implication here is that another specialized relationship is used to aggregate multi-part works (potentially unrealized) together.

It would be necessary to do some conceptual gymnastics to fully align this relationship with the earlier expression-level use case but, as Bekiari et al. tell us, "This property [R10] associates an F2 Expression X with a structural component Y that conveys in itself the complete concept of a work that is member of (R10) the overall work realized by X" – p 74. It is important to note that *hasMember* (R10) does not directly associate an Expression (F2) with a Complex Work (F15). Rather the *hasMember* (R10) relationship links a Work (F1) that is realized by an Expression (F2) with the Complex Work (F15) that is realized by an "*aggregate*" Expression (F2).

Here FRBR_{oo} seems to be ascribing to the formal definition set forth by Wickett and Renear in definition D.31' and thereby establishes a model for two things:

- Works realized by corresponding expressions (aggregate and otherwise) and,
- Unrealized aggregate works.

For our formalization of the FRBR_{oo} account to be complete, we must extend definition D.35 to encompass both aggregate expressions and aggregate works.

Bekiari et al.'s use of Dante's *Inferno* indicates that FRBR_{oo}'s R10 relationship serves the same purpose as the R5 relationship (*has-expression-component*). Even though Bekiari et al. use language that describe the member *works* as "structural components" for a complex work, the label used—*has-member*—seems more correct (since the *is-gathered-into* relationship was shown to be the same as Winston, Chaffin, and Herrmann's *member-of* relationship in the first section of this chapter). To differentiate the R10 relationship from our existing (Winstonian) *membership* relation, we will apply the following label to R10—*has-work-member*.

Since, as we noted, R10 serves the exact same purpose as R5, using the exact same examples, we have no reason not to accept that R10 is also *irreflexive, asymmetric, transitive, discrete, dense, separable, NSIP*, and *homogeneous*. It also possesses all of the same metaproperties as FRBR's *bibliographic-part-of* relation. The primary difference between FRBR_{oo}'s R5 and R10 relationships is with regard to their domain and range. Whereas the R5 relationship is a

homogeneous relationship with respect to *expressions*, the R10 relationship is homogenous with respect to *works*. We can use the following axiom and theorems to showcase this.

A.37 (Domain and Range of R10): $\forall x \forall y (hasWorkMember(x, y) \rightarrow (ComplexWork(x) \wedge Work(y)))$

Axiom A.37 states that for all entities x and y , if entity x *has-work-member* entity y , then it is the case that entity x is a *complex-work* and entity y is a *work*.

As was the case with the R5 relationship (*has-expression-component*), the R10 (*has-work-member*) relationship is not a kind of *componenthood* relation. We can demonstrate that this is the case when we consider that characters, settings, and climaxes are all (functional or structural) *components* of a *work* but are not typically considered to be *works* in and of themselves. We should note, though, that axiom A.37 does not prevent entity y from being a *complex-work* itself.

Like the R5 relation, the R10 relationship is actually a specialization of FRBR's *bibliographic-part-of* relation. In turn, it is also trivially the case that the R10 relationship is also a specialization of the *is-gathered-into* relation. We can express this through the following theorems which follow from axioms A.32, A.33, and A.37.

T.4 (Work Membership as Bibliographic Parthood):
 $\forall x \forall y (hasWorkMember(x, y) \rightarrow bPartOf(y, x))$

Theorem T.4 is a trivial consequence of the combination of axioms A.33 and A.37. Theorem T.4 states that for all entities x and y , if entity x *has-work-member* entity y , then it is also the case that entity y is *bibliographically-part-of* entity x .

T.5 (Work Membership as Collection Membership):

$$\forall x \forall y (hasWorkMember(x, y) \rightarrow isGatheredInto(y, x))$$

Theorem T.5 is a trivial consequence of the combination of axiom A.32 with theorem T.4.

Theorem T.5 states that for all entities x and y , if entity x *has-work-member* entity y , then it is also the case that entity y is *gathered-into* entity x .

From axiom A.37 and theorems T.4 and T.5, we can now craft an initial formalization to define Complex Works through means of the *hasWorkMember* relationship.

D.16 (Complex Works Defined):

$$\forall x (ComplexWork(x) =_{df} \exists y (hasWorkMember(x, y) \wedge \sim(x = y)))$$

Definition D.16 states that for all entities x , entity x is a *complex-work*, if and only if it is the case that there exists some entity y , such that entity x *has-work-member* entity y , and it is not the case that entities x and y are the same entity.

It seems unclear at first if this bibliographic object, *Complex Work* is, in fact, a bibliographic aggregate. It is important to recall the intention of Complex Work is to represent works that are intended to be aggregate works by their creators, e.g., Dante's 'Divine Commedia.' In this case, it seems safe to equate *Complex Works* with bibliographic aggregates. We can represent this using the following axiom.

$$\mathbf{A.38 (Complex Works as Bibliographic Aggregates):} \forall x (ComplexWork(x) \rightarrow bAggregate(x))$$

Axiom A.38 states that for all entities x , if entity x is a *complex work* then it is also the case that entity x is a *bibliographic aggregate*.

We can use axiom A.38 to rework our definition for *complex works* into one that defines *bibliographic aggregates*.

D.15-WM (Complex Works Redefined as Bibliographic Aggregates):

$$\forall x (bAggregate(x) =_{df} \exists y \exists z (hasWorkMember(x, y) \wedge \sim(x = y)))$$

The next candidate aggregate entity in FRBR₀₀ is Container Work (F16). The label that Bekiari et al. have used here would normally indicate that the entity is an aggregate of some kind in ordinary English. However, the text of FRBR₀₀ makes it clear that our ordinary everyday concept of “container” is not what the FRBR₀₀ authors intend. Instead they give us the rejoinder that:

“This class is an ‘abstract class,’ in that it only serves as an umbrella for its three subclasses. As a consequence, it can only be instantiated by instances of any of its subclasses: nothing can be an instance of it, unless it is an instance of either F17 Aggregation Work, F19 Publication Work, or F20 Performance Work.” – Bekiari et al, p 54.

The implication here is that Container Work (F16) is an abstraction which serves only to conceptually link Aggregation Works (F17), Publication Works (F19), and Performance Works (F20) together into a class/subclass hierarchy, i.e., it isn’t actually a container and so the *containment* relationship will not figure into our discussion.

As such, an instance of Container Work (F16) can only be viewed as an aggregate entity if an instance of any of its subclasses is an aggregate entity. Considering each of Container Work's (F16) subclasses in turn, we will see that none of them could be interpreted as bibliographic aggregates in and of themselves.

Bekiari et al. define Aggregation Work (F17) as, a work that, “comprises works whose essence is the selection and/or arrangement of expressions of one or more other works” – p 54. They go on to say that, “This does not make the contents of the aggregated expressions part of this work, but only part of the resulting expression” – p 54. The implication is that F17 captures only the intellectual effort of the act of aggregating and therefore is not an aggregate entity of its own accord.

Bekiari et al. define Publication Work (F19) as a work that “comprises works that have been planned to result in a manifestation product type or an electronic publishing service and that pertain to the rendering of expressions from other works” – p 55. At first it is not clear how this is to be interpreted but the examples tellingly all begin with the noun, “concept.” The implication here too, is that F19 only captures the intellectual effort of designing a manifestation of one or more works.

Bekiari et al. define Performance Work (F20) as a work that “comprises the sets of concepts for rendering a particular or a series of like performances” – p 55. The talk of sets makes this entity sounds suspiciously aggregate-like, but, Bekiari et al. provide a detailed account of the implications of Performance Work (F20) being a subclass of Container Work (F16).

“F20 Performance Work is declared as a subclass of F16 Container Work. This implies that the incorporated expressions (such as the text of the staged play, the text of the argument for the ballet, the recorded music to be used for the ballet, or the content of the musical score to be used for a concert, etc.) are not by themselves a part of the expression of this F1 Work. Rather, an expression (F25 Performance Plan) of the instructions the stage production, choreography or musical performance consists of incorporates (P165) that textual or musical content. In other words, the text of ‘Hamlet’ is not a component of the concepts that underlie a given mise-en-scène of ‘Hamlet,’ but any staging directions (F25 Performance Plan) that convey a given director’s vision of ‘Hamlet’ must necessarily incorporate the text of ‘Hamlet.’” – Bekiari et al. pp 55-6.

From this paragraph it is clear that the set of concepts is to be treated as a whole work (i.e., a special kind of work) which is realized by a special kind of expression (F25 Performance Plan). Since FRBR_{oo}’s Performance Plans (F25) are a subclass of FRBR_{oo}’s Self-Contained Expressions (F22), which in turn are a subclass of FRBR_{oo}’s Expressions (F2), we can safely say that definition D.35 (which we previously defined) accounts for this very specific case. However, we must also note that the relationship through which aggregation occurs with respect to Performance Plans (F25) is labeled—*incorporates* (P165). If we accept that definition D.35 accounts for this usage of P165, then we should accept that P165 is also *irreflexive, asymmetric, transitive, discrete, dense, separable, NSIP*, and *homogeneous* like relationships R5 and R10.

This alternate account of bibliographic aggregation can be modeled by developing an additional axiom (below).

A.39 (Domain and Range of Incorporation): $\forall x \forall y \left(\text{incorporates}(x, y) \rightarrow \left(\text{PerformancePlan}(x) \wedge \text{Expression}(y) \right) \right)$

Axiom A.39 states that for all entities x and y , if entity x *incorporates* entity y , then it is also the case that entity x is a *performance plan* and entity y is an *expression*.

What axiom A.39 is trying to communicate is that the Performance Plan (F25), a special kind of *expression*, incorporates the contents of other Expressions (F2). However, Bekiari et al. also use the *incorporates* (P165) relationship a second means to aggregate Expressions (F2) and Expression Fragments (F23) into Self-Contained Expressions (F22). To accommodate these other use cases for the *incorporates* relationship, axiom A.81 must be broadened to include them.

A.39' (Expanded Domain and Range of Incorporation):
 $\forall x \forall y \left(\text{incorporates}(x, y) \rightarrow \left(\left(\text{PerformancePlan}(x) \wedge \text{Expression}(y) \right) \vee \left(\text{SelfContainedExpression}(x) \wedge \left(\text{Expression}(y) \vee \text{ExpressionFragment}(y) \right) \right) \right) \right)$

Axiom A.39' states that for all entities x and y , if entity x *incorporates* entity y , then it is the case that either entity x is a *performance plan* and entity y is an *expression*, or entity x is a *self-contained expression* and entity y is either an *expression* or an *expression fragment*.

Axiom A.39' can be simplified if one resorts to FRBR_{oo}'s class/subclass hierarchy which is expressed via the following axioms.

A.40 (Expression Fragments as Expressions): $\forall x(\text{ExpressionFragment}(x) \rightarrow \text{Expression}(x))$

Axiom A.40 states that for all entities x , if entity x is an *expression fragment*, then it is also the case that entity x is an *expression*.

A.41 (Performance Plans as Self-Contained Expressions):
 $\forall x(\text{PerformancePlan}(x) \rightarrow \text{SelfContainedExpression}(x))$

Axiom A.41 states that for all entities x , if entity x is a *performance plan*, then it is also the case that entity x is a *self-contained expression*.

Axioms A.40 and A.41 allow a further revision to A.39', distilling it down to the simpler form expressed below.

A.39'' (Refined Domain and Range of Incorporation): $\forall x\forall y(\text{incorporates}(x, y) \rightarrow \text{SelfContainedExpression}(x) \wedge \text{Expression}(y))$

Axiom A.39'' states that for all entities x and y , if entity x *incorporates* entity y , then it is also the case that entity x is a *self-contained expression* and entity y is an *expression*.

Once again though, we can see from the examples that have been used, that FRBR_{oo}'s *incorporation* relationship matches the R5 relation's usage. The only substantial difference is a yet narrower domain and range. Thereby, we can state that FRBR_{oo}'s *incorporation* relationship is actually a narrower sense of FRBR_{oo}'s R5 (*hasExpressionComponent*) relation. It trivially follows

that *incorporation* is also a narrower sense of the *gathered-into* relation. We can express this through the following theorems.

T.6 (Incorporation as Expression Componenthood): $\forall x\forall y(\text{incorporates}(x,y) \rightarrow \text{hasExpressionComponent}(x,y))$

Theorem T.6 trivially follows from the combination of axioms A.33 and A.39''. Theorem T.6 states that for all entities x and y , if entity x *incorporates* entity y , then it is also the case that entity x *has-expression-component* entity y .

T.7 (Incorporation as Collection Membership): $\forall x\forall y(\text{incorporates}(x,y) \rightarrow \text{isGatheredInto}(x,y))$

Theorem T.7 trivially follows from the combination of axiom A.32 with theorem T.6. Theorem T.7 states that for all entities x and y , if entity x *incorporates* entity y , then it is also the case that entity y is *gathered-into* entity x .

One potential formalization of this account of bibliographic aggregates might appear as an alternate definition for bibliographic aggregates as aggregate expressions.

D.15-I (Alternative Definition of Bibliographic Aggregates as Aggregate Expressions): $\forall x \left(\text{bAggregate}(x) =_{df} \exists y(\text{incorporates}(x,y) \wedge \sim(x = y)) \right)$

Definition D.15-I states that for all entities x , entity x is a *bibliographic aggregate*, if and only if, there exists some entity y , such that entity x *incorporates* entity y and it is not the case that entities x and y are the same entity.

Given theorem T.6 though, we can see that definition D.15-EC subsumes the conceptual space defined by definition D.15-I, and so, definition D.15-I has only been provided here for the sake of completeness.

One final candidate aggregate entity remains to be accounted for: FRBR_{OO}'s Serial Work (F18). Bekiari et al. state that:

“This class comprises works that are, or have been, planned to result in sequences of Expressions or Manifestations with common features. Whereas a work can acquire new members during the time it evolves, Expressions and Manifestations are identified with a certain state achieved at a particular point in time. Therefore[,] there is in general no single Expression or Manifestation representing a complete serial work, unless the serial work has ended.” – pp 54-5.

It seems clear that that Serial Works are to be understood as a special kind of Complex Work that, “may or may not have a plan for an overall expression” – Bekiari et al, p 55. Since plans for expressions are captured through FRBR_{OO}'s Publication Works (F19), the following axiom can be used to situate Serial Works (F18) within the formalization that has already been constructed.

A.42 (Serial Works as Complex Works): $\forall x \left(\text{SerialWork}(x) \rightarrow \left(\text{ComplexWork}(x) \wedge \left(\text{PublicationWork}(x) \vee \sim \text{PublicationWork}(x) \right) \right) \right)$

Axiom A.42 states that for all entities x , if entity x is a *serial work*, then it is also the case that entity x is a *complex work* and it may also be the case that entity x is or is not a *publication work*.

Since the disjunctive part of axiom A.42 results in a tautology, i.e., it is always the case that x is a Publication Work (F19) or not a Publication Work (F19), axiom A.42 can be reduced to the following implication:

A.42' (Serial Works as Complex Works [simplified]): $\forall x(\text{SerialWork}(x) \rightarrow \text{ComplexWork}(x))$

Axiom A.42' states that for all entities x , if entity x is a *serial work*, then it is also the case that entity x is a *complex work*.

As Complex Works have already been described (through axioms A.37, theorems T.4 and T.5, and definition D.15-WM), it appears to be the case the Definition D.15-WM adequately accounts for FRBR_{OO}'s Serial Works (F18). The three FRBR_{OO} definitions can now be merged together to provide a singular account of bibliographic aggregates from the FRBR_{OO} perspective.

D.15-FBA (FRBR_{OO}'s Bibliographic Aggregates Defined):

$$\forall x \left(bAggregate(x) =_{df} \exists y \left(\left(hasExpressionComponent(x, y) \vee hasWorkMember(x, y) \vee incorporates(x, y) \right) \wedge \sim(x = y) \right) \right)$$

Taken altogether, we can now produce a table (Table 4.4 below) that showcases how these three FRBR_{OO} relationships are situated among their sibling relations.

Once again though, the crux of the analysis rests upon our acceptance of theorems like T.3, T.5, and T.7. However, as in FRBR, FRBR_{OO} is extremely focused on providing a model for just the following kinds of aggregates: collections, anthologies, journal issues, monograph series and similar bibliographic aggregates. Since the scope and intent is the same we can see no compelling

reason to reject theorems T.3, T.5, and T.7 and find that like FRBR's *bibliographic-part-of* relationship, FRBR_{OO}'s *has-expression-component*, *has-work-member*, and *incorporates* relationships are all specialized versions of Winston, Chaffin, and Herrmann's *member-of* relationship. Which is to say, they are specialized versions of Wickett, Renear, and Furner's *is-gathered-into* relationship, despite the rather misleading labels that FRBR_{OO}'s authors have given them.

4.5.3. Initial Findings

Examining Table 4.4 (below) one thing becomes apparent—our intuition that *mereological aggregate models* are independent of concerns regarding specific definitions for bibliographic entities that are *bibliographic aggregates* is mostly correct. We have to say mostly here because the FRBR_{OO} model lacks support for two of the kinds of things in FRBR's WEMI model (i.e., *manifestations* and *items*). It does not appear that Bekiari et al.'s effort to harmonize the FRBR and CIDOC-CRM standards was fully successful.

This is of some concern because FRBR_{OO} does not support aggregates as the artifact level (i.e., for WEMI's *items*). Such an outcome is troubling because it would appear to miss half of the conceptual conversation that Verona (1959), Wilson (1968), Lubetzky (1969), Svenonius (2000), and Smiraglia (2001) were having. Support for bibliographic entities at what WEMI calls the *item-level* is a must. The point these authors were all trying to make is that we need a balance between describing the abstract (content) and the concrete (artifacts).

Table 4.4: FRBRoo Bibliographic Parthood (*bPartOf*) Relationship Metaproperties

Metaproperty	<i>is-gathered-into</i>	<i>bibliographically-part-of</i>	<i>has-expression-component</i>	<i>incorporates</i>	<i>has-work-member</i>
<i>Reflexive</i>	-	-	-	-	-
<i>Irreflexive</i>	+	+	+	+	+
<i>Symmetrical</i>	-	-	-	-	-
<i>Asymmetrical</i>	+	+	+	+	+
<i>Antisymmetrical</i>	+	+	+	+	+
<i>Transitive</i>	+	+	+	+	+
<i>Dense</i>	+	+	+	+	+
<i>Discrete</i>	+	+	+	+	+
<i>NSIP</i>	+	+	+	+	+
<i>SIS</i>	-	-	-	-	-
<i>Separable</i>	+	+	+	+	+
<i>Homeogeneous</i>	+	+	+	+	+
<i>WEMI-Homogeneous</i>	-	+	-	-	-
<i>W-Homogeneous</i>	-	-	-	-	+
<i>E-Homogeneous</i>	-	-	+	-	-
<i>Special-E-Homogeneous</i>	-	-	-	+	-

We also must note that when we move from the minimal metadata necessary to describe *bibliographic aggregates* in general to the minimal metadata needed to articulate the more specific differences among particular kinds of *bibliographic aggregates*, then we fully expect the description process to break down. This is a fatal flaw in the WEMI model and not caused by the *mereological aggregate model* itself.

4.6. Findings on Mereological Aggregate Models

Despite the incompleteness of the bibliographic aggregate model set forth in FRBR₀₀, we can conclude several things which definitively answer research question one, “How comparable are the *mereological aggregate models* used by the standards in our case studies?” They are completely comparable. We further note that:

1. All the *mereological aggregate models* are, in fact, all the same aggregate model. They all conform to Winston, Chaffin, and Herrmann’s notion of *membership* (i.e., they are all examples of the *member-of* relationship).
2. So long as practioners stick to the general case of *bibliographic aggregates*, then adequate basic metadata allowing users to perform FRBR’s find, identify, select, and obtain tasks should be achievable, at least within the limitations of the particular bibliographic standard they are employing. Where basic metadata is found inadequate to meet these needs, then the focus should be on extending the model and not rewriting it to better match intuitions. In the case of FRBR and FRBR₀₀, extension will be a difficult process because, as has been shown, the WEMI model breaks down when we move from general features of

bibliographic entities to more specific features possessed by specific kinds of bibliographic entities like video games, serials, series, etc.

3. The FRBR_{oo} standard is also clearly deficient in being capable of providing minimal metadata for the general case as it only succeeds in harmonizing one half of FRBR's original aggregate model. However, FRBR_{oo} has several benefits that can be realized through its much clearer semantics and easier to implement ontology,⁵⁹ which may indicate that its detrimental aspects can be overlooked in certain cases where its overall benefits outweigh its disadvantage for adequately describing *bibliographic aggregates*. It remains unclear however, if it is possible to manage the flaws in the WEMI model regarding the minimal metadata needed for more specific kinds of bibliographic entities.
4. The specifics of what collections, serials, series, compounding entities, or even “superworks” (or anthologies, journal issues, music compilations, etc.) are not necessarily pertinent for the application of any of these standards as their aggregate models are general enough to accommodate all of these bibliographic entities. Even the WEMI model provides sufficient minimal metadata for bibliographic entities in general.

The upshot of all of this is that we do not seem to need to come to definitive answers to questions that Lagoze and Fielding (1998); Lee (2000, 2005); Antelman (2004); Currall, Moss, and Stuart (2004); Gonçalves et al. (2004); Palmer (2004); Palmer et al. (2006); Krier (2008); Galton (2010); Yeo (2012); Lee, Jett, and Perti (2015); Jett et al. (2017); Kiryakos and Sugimoto (2018); Sugimoto

⁵⁹ FRBR is a relatively informal and vague standards document, even by LIS practices.

et al. (2018); Lee et al. (2018) are all trying to answer.⁶⁰ So long as we treat bibliographic aggregates as our bibliographic standards suggest, at least for those using the *mereological aggregate model*, then the essential metadata required for users to accomplish FRBR's user tasks can be assured to be included in our IR systems.

Starting from that point, we can ask and consider answers for the questions regarding how each specific kind of *bibliographic aggregate* is distinct from the other, and thereby suggest extensions for our standards with additional metadata at various levels along the content-artifact continuum that will help users accomplish FRBR's user tasks better than they already are. However, as we mentioned in Chapter 2, the WEMI model almost immediately begins to break down as soon as we begin to try extending it with attributes intended to highlight the finer-grained differences among different kinds of bibliographic entities. This is also true for specific kinds of *bibliographic aggregates* and, to a great extent, it is the *casus belli*, the *raison d'être* for why we even possess an alternative to the *mereological aggregate model*. So, in the case of FRBR and FRBR_{oo}, extension from the general case to encompass more specific cases necessitates addressing the WEMI model's inherent flaws.

4.7. Chapter Summary

In this chapter, we closely examined the *mereological aggregate models* employed by three high-level conceptual standards: DC-CAP, FRBR, and FRBR_{oo}. Through our analysis we were able to show that all of the *mereological aggregate models* are the same model by virtue of Winston, Chaffin, and Herrmann's *member-of* relationship, or a sub-property of it, being the core part/whole

⁶⁰ At least from the point of view of producing a general overarching models for bibliographic entities and bibliographic aggregates.

relationship being employed by each of the conceptual standards' aggregate models. However, we also found that the aggregate model employed in the FRBR₀₀ standard is an incomplete one since it lacks the ability to articulate part/whole differences for artifacts (i.e., the WEMI *item*-level).

We also noted that while this aggregate model is sufficient for the creation of minimum-level bibliographic metadata describing *bibliographic aggregates* in general, it is also too coarse to capture the particular differences between different kinds of *bibliographic aggregates* (e.g., anthologies, digital collections, journal issues, monograph series, etc.). To some extent, a *content-artifact model* like the WEMI model should be expected to help make some of these distinctions clear; however, in practice the WEMI model has experienced problems when moving from the coarse-grained descriptive needs of bibliographic entities in general to the finer-grained distinctions required to distinguish particular kinds of bibliographic entities (e.g., video games, novels, movies, music albums, etc.) from one another. This situation has led to the development of a new kind of aggregate model, one that focuses not on part/whole differences but on content-artifact differences.

In the next chapter, we closely examine this new aggregate model, which we call the *content-artifact aggregate model*, within the context of IFLA's LRM conceptual standard. LRM is a new standard (not yet even two years old at the time this dissertation was written), and it makes a series of refinements to the WEMI model that have serious implications for the modeling of *bibliographic aggregates*.

5. Analysis of the Content-Artifact Aggregate Model

5.1. Chapter Overview

In this chapter, we are going to closely examine an alternate aggregate model to the one we examined in the previous chapter. This alternate aggregate model specifically exploits features of the WEMI model, which we have previously noted is a model that describes information objects on a continuum from content to artifact, i.e., it is a *content-artifact model*. Since the alternate aggregate model analyzed in this chapter employs features particular to *content-artifact models*, we call this alternate aggregate model the *content-artifact aggregate model*.

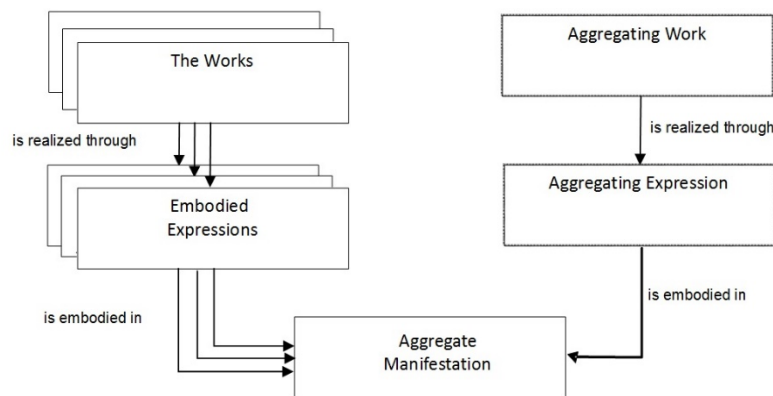


Figure 5.1: LRM Aggregate Model

As we noted previously, this model focuses on differences along the content-artifact continuum of aspects describing bibliographic entities, i.e., it requires an attendant content-artifact model in order for it to be understood. We begin then by discussing how these two models, aggregate model and content-artifact model, relate to one another. In this case the primary content-artifact model is the WEMI model. However, as we noted in the previous chapter, LRM's WEMI model is quite different than the one originally published with FRBR. Hence, LRM's WEMI model is also analyzed in order to make better sense of LRM's aggregate model. Through the analysis we are

able to argue that the changes to WEMI are significant enough to both invalidate the figure typically used to illustrate it (i.e., the Group 1 Entities figure [Figure 5.2 below]) and to cause serious deficiencies in LRM’s aggregate model.

5.2. IFLA’s Library Reference Model

5.2.1. The Content-Artifact Aggregate Model

The *content-artifact aggregate model* situates *bibliographic aggregates* as a particular conceptual aspect of a bibliographic entity. Specifically, *bibliographic aggregates* are equated with WEMI’s notion of *manifestations* (see Figure 5.1 above). A very early form of this kind of model appears in Barbara Tillett’s dissertation (1987) in the form of her “extractive” part/whole relationship.⁶¹

The earliest forms of the specific version of this model that we are analyzing were published by the FRBR-WGA (2011) and Žumer and O’Neill (2012) and specifically defined *bibliographic aggregates* as kinds of WEMI *manifestations*. Recently though, this model has been codified in LRM, which defines an all-new WEMI model, which we call WEMI2 here. As we will see through this analysis, WEMI2 has a profound effect when *bibliographic aggregates* are defined as WEMI2 *manifestations*.⁶²

⁶¹ In Tillett’s case, unlike the one we will closely examine, the aggregate whole is the *work* and the part is the *item*. This is an interesting model because one thing that it immediately implies is that when a part, like a short story, an article, a photograph from a collection, etc., is separated from the whole, it loses some (or perhaps all) of its meaning. This is rather a different state of affairs from the *content-artifact aggregate model* we are about to closely examine. However, we should note, outside of some anecdotal considerations for how serials are represented in MARC records (and on catalog cards) this particular conceptualization does not appear in any of our high-level conceptual models (such as FRBR, FRBR_{OO}, LRM, etc.).

⁶² We will also see that there are some interpretations of WEMI2 that are going to find it much more robust than the original WEMI model when it comes time for extending LRM to articulate minimal metadata for particular kinds of bibliographic entities like video games, serials, series, etc.

While we will still be proposing various axioms, definitions, and theorems to help with our analysis, much of the analysis is going to be focused on WEMI2 directly and the illustrative diagrams—ER-diagrams—that help us understand it. Thus, the analysis is also going to call on the common symbol system employed in ER-diagrams (Chen 1976) to aid in the analysis. However, in the end, we will still be examining the roles and participation constraints of the entities and relationships in the model. We will see as a result of this analysis that the *content-artifact aggregate model* is extremely sensitive to variations in the definition of WEMI’s *manifestation* and that in the LRM case, the *manifestation* entity is semantically overloaded.

5.2.2. Bibliographic Entities and FRBR

The FRBR approach to describing bibliographic entities (IFLA 1998) has become a widely accepted one among libraries as evidenced by the general “FRBRization” of library catalogs worldwide (Salaba and Zhang 2007) and more recently, the advent of OCLC’s WorldCat Work Descriptions (OCLC 2015). It has also made some inroads into conceptual models for entities within archives and museums (Nimer and Daines 2013) through conceptual model harmonization efforts like FRBR_{oo} (Le Bœuf 2012, Bekiari et al. 2015).

Its central model—WEMI—has brought the conceptual tension between content and artifacts (Verona 1959, Wilson 1968, Lubetzky 1969, Svenonius 2000, Smiraglia 2001) to the fore of considerations for metadata. Unfortunately, as we saw in the analysis of FRBR in the previous chapter, the original WEMI entities are poorly defined in terms that are easy to understand (Taniguchi 2002, 2003; Renear and Choi 2006) and the relationships between them are not defined at all. This last issue put Wickett and Renear (2009) in the position of considering WEMI’s

relationships (*realizes, embodies, exemplifies*) to be primitive for their analysis of the WEMI model.

When the scope of an information retrieval (IR) system is primarily the *works* expressed through the mediums of books or journal articles, and where the focus is equivalence relationships, then FRBR's WEMI model works very well. However, as we previously mentioned, several studies have demonstrated that the FRBR approach begins to break down when confronted by *works* that are expressed through kinds of media other than **text**. Specifically, conceptually adequate mappings for entities running the gamut from medieval (and older) manuscripts to video games to live performances of various kinds seem to elude us at every turn (Jonsson 2005, Miller and Le Bœuf 2005, Nicolas 2005, Baca and Clarke 2007, McDonough et al. 2010, Lee et al. 2012). In part, this is because the selection of attributes ascribed to WEMI's entities are a poor fit for entities that are not primarily communicated through text-bearing formats.

Recently, IFLA has deployed their LRM standard (Riva, Le Bœuf, and Žumer 2017) in an effort to respond to FRBR's critics and better inform the emerging shape of new cataloging standards like RDA. Importantly, LRM does two things which are beneficial for its adopters:

1. It provides clear, ontology-like, documentation defining and explaining all of the entities, attributes, and relationships that IR systems built in accord to the models it sets forth should be expected to support.

2. It provides an important revision to FRBR's WEMI model in the form of its own *work-expression-manifestation-item* model, which we call WEMI2 for clarity's sake.

However, as we will see in the analysis that follows, while LRM provides its adopters with a clearer picture of the bibliographic universe in general, it still contains a great deal of confounding factors. So, while it presents an important refinement to the conceptual space initially explored by the FRBR framework, a great deal of additional work remains to be done at the conceptual and related levels.

5.2.3. FRBR, LRM, and Bibliographic Aggregates

As with the cases for so many non-text-based bibliographic entities, many practitioners also found the FRBR conceptual model to be a poor fit for *bibliographic aggregates* (FRBR-WGA 2011). For this reason, a working group—the FRBR Working Group on Aggregates (FRBR-WGA)—was instituted in 2006 with the express goal of developing a conceptual model for aggregates that fit within the existing FRBR framework. This working group's effort ended in 2011 when they issued their final report. Like FRBR₀₀, that final report focuses on a particular intellectual level of abstraction. Whereas FRBR₀₀ focuses primarily on the *work* and *expression* levels, the FRBR-WGA arrived at a model that focused on the *manifestation* level.

As we previously noted, there were dissenting opinions (Tillett et al. 2014) regarding the need for an alternate aggregate model for FRBR. Nonetheless, the model suggested by the FRBR-WGA's final report (2011; Žumer and O'Neill 2012) was adopted for use in LRM (Riva, Le Bœf, and Žumer 2017) as one of the many changes it deploys to better clarify how one represents various

aspects of bibliographic entities and thereby capture a minimal amount of metadata sufficient for fulfilling FRBR's user tasks of finding, identifying, selecting, and obtaining.

The specifics of the *content-artifact aggregate model* focus on the following definition for aggregates:

“An aggregate is defined as a *manifestation* embodying multiple *expressions*” –
Riva, Le Boëf, and Žumer 2017, p 93.

There appears to be two things being communicated by this succinct definition:

1. *Bibliographic aggregations* are WEMI2 *manifestations* and
2. WEMI2 *manifestations* aggregate WEMI2 *expressions*.

Unlike the approach taken in the previous chapter, the *content-artifact aggregate model* used in LRM does not seem to take a mereological approach. Aggregates in the LRM sense are not “wholes” in the typical “whole-part” sense. Instead, they are *manifestations* that *embody* multiple *expressions*.⁶³

The model's authors attempt to clarify what they mean by noting that there are three distinct types of aggregates:

⁶³ And this is despite clear instances of part-talk in the descriptions for LRM's aggregates (e.g., “*Manifestations* may **contain** [emphasis added] multiple *expressions* as indicated by the many-to-many relationship between *expressions* and *manifestations*.” – Riva et al. 2017, p 93).

1. Aggregate Collections of Expressions
2. Aggregates Resulting from Augmentation
3. Aggregates of Parallel Expressions

Much of the focus appears to be on recording metadata about particular units of publication. However, it is somewhat unclear what the boundaries of a unit of publication is. With regard to “aggregate collections of expressions,” we are told that “Collections include selections, anthologies, monographic series, issues of serials, and other groups of resources.” – Riva, Le Boef, and Žumer 2017 p 93. Immediately afterwards we are told that “Examples include journal issues (aggregates of articles), multiple novels published together in a single volume, books with independently written chapters, complications on CD’s (aggregates of individual songs), and various collected/selected works.” – Riva, Le Boef, and Žumer 2017 p 93.

Almost immediately, we see that there is some discontinuity within the definition. A monographic series is typically published over a long period of time and through multiple individual novels. But the example given specifically notes “multiple novels published together in a single volume.” Hence, we are unsure that monographic series are actually intended for inclusion.

There are other potential problems for this model. Coyle (2016b) notes that this model has issues when one considers paratext. For instance, when paratext is considered significant enough, then there seems to be an issue of combinatorial explosion. We come back to this issue in the ultimate chapter of this dissertation. For now, we can observe that if paratext has an impact on the content of a bibliographic entity, then it is likely that a new bibliographic entity has come into existence.

Conversely, if the paratext does not have an impact on the content of the bibliographic entity, then quite possibly it is either not important enough to record any metadata about at all or its role in any particular aspect of a bibliographic entity is not as a member of a bibliographic aggregate. For example, a table of contents might be a functional component of a novel or the spine might be a structural component of a book, but neither are members of a bibliographic aggregate in the same manner in which a short story is a member in an anthology or a digital image is a member of a digital image collection.

We should also note again, as we did in the early chapters of this dissertation, that something seems to be lost in the *content-artifact aggregate model's* account of *bibliographic aggregates*. For example, we could consider a digital collection of digitized photographs. Let us reuse the specific example of a collection of digitized photographs depicting various scenes from the Meiji Era that Jett, Fenlon, and Downie (2018) use. As they note, the individual photographs have content that is specific to them individually, but when grouped together additional content—their shared context—is brought into being such that the digital collection itself has a topicality that is quite different from its individual members. We might be tempted to object that a digital collection does not seem to be the same kind of publication unit that the examples given in LRM (on p 93) seem to be. However, it is quite possible to move the files comprised by the digital collection as a singular unit. Indeed, we might store them in an off-the-shelf digital library product such as Greenstone,⁶⁴ and move the entire digital library as a singular product. And on this account, it seems as though digital collections at least are exactly the kind of production units that the *content-artifact aggregate model* is designed to describe.

⁶⁴ <http://www.greenstone.org/>

It appears that some of the information that can be captured using *mereological aggregate models* is missing from the picture that the *content-artifact aggregate model* draws. To be certain, we need to get a better sense of what the *content-artifact aggregate model* is trying to do.

5.2.4. Initial Analysis

If we take the *content-artifact model's embodies* relationship to be primitive, then an initial formalization of Riva, Le Bœf, and Žumer's account might be as follows:

D.17 (LRM Bibliographic Aggregate Definition):

$$\forall x \left(\text{Aggregate}(x) =_{df} \exists y \exists z (\text{embodies}(x, y) \wedge \text{embodies}(x, z) \wedge \sim(x = y) \wedge \sim(x = z) \wedge \sim(y = z)) \right)$$

Definition D.17 states that for all entities x , entity x is a [bibliographic] aggregate, if and only if there exists some entities y and z such that entity x embodies entity y , entity x also embodies entity z , and it is not the case that entities x , y , and z are the same entity.

If we consider the examples given, i.e., selections, anthologies, monographic series, issues of serials, supplementary dependent works, and parallel expressions, they conform to the examples used in the preceding chapter. As such, we would normally apply the metaproperties we have already developed and situate the embodiment relationship among our existing meronymic relationships. However, those metaproperties are specific and appropriate for well-defined part/whole relationships. Unfortunately, FRBR's *embodies* relationship is not well-defined (Wickett & Renear 2009), and as we shall see the version of the *embodies* relationship employed by LRM is not much better defined. Further, there are aspects of the LRM document that imply that we are not meant to interpret WEMI relationships like *embodies* as meronymic relationships.

With regard to serials⁶⁵ Riva, Le Boëf, and Žumer tell us that, “Serials are complex constructs that combine whole/part relationships and aggregation relationships” – p 94. So, from the perspective of the editors of LRM, the WEMI relationships (i.e., aggregation relationships) are apparently not the same as whole/part relationships. Therefore, we will need to develop a different explanation for the WEMI relationships regarding what they are attempting to communicate.

However, the LRM WEMI model (WEMI2) is actually quite different than the one set forth in FRBR and so we may not be able to take the *embodies* relationship to be primitive as Wickett and Renear (2009) did in their analysis.

5.3. Formal Analysis of the LRM WEMI Model—WEMI2

Since LRM sets forth a new model of the WEMI entities, we may not be able to use the definitions from the formal analysis of FRBR’s WEMI entities set forth by Wickett and Renear (2009) as the basis for our analysis. Instead we will have to begin anew.

Like Wickett and Renear, we are concerned with producing a first order theory that will help us better understand, in this case, what the *content-artifact aggregate model* is trying to accomplish. Before we can attempt to unpack what is happening with the *content-artifact aggregate model*, we need to understand what the WEMI2 (i.e., LRM’s central *content-artifact model*) is trying to explain.

⁶⁵ Which are oddly excepted from IFLA-LRM’s aggregate model.

Examining the LRM text, we find that the natural language definition of a WEMI2 *work* is similar to what was in FRBR—a work is defined to be:

“The intellectual or artistic content of a distinct creation.” – Riva, Le Bœf, and Žumer 2017, p 21.

Similarly, Riva, Le Bœf, and Žumer tell us that:

- An *expression* is, “A distinct combination of signs conveying intellectual content” – Riva, Le Bœf, and Žumer 2017, p 23.
- A *manifestation* is, “A set of all carriers that are assumed to share the same characteristics as to intellectual or artistic content and aspects of physical form. The set is defined by both the overall content and the production plan for its carrier or carriers” – Riva, Le Bœf, and Žumer 2017, p 26.
- An *item* is, “An object or objects carrying signs intended to convey intellectual or artistic content” – Riva, Le Bœf, and Žumer 2017, p 27.

These natural language definitions are substantially different from those Wickett and Renear were working with from the original FRBR WEMI model. We reprise the original definitions below:

- A FRBR *work* is “a distinct intellectual or artistic creation.” – IFLA 1998, p 17.

- A FRBR *expression* is “the intellectual or artistic realization of a *work* in the form of alphanumeric, musical, or choreographic notation, sound, image, object, movement, etc., or any combination of such forms.” – IFLA 1998, p 19.
- A FRBR *manifestation* is “the physical embodiment of an *expression* of a *work*.” – IFLA 1998, p 21.
- A FRBR *item* is “a single exemplar of a *manifestation*.” – IFLA 1998, p 24.

Comparing these two groups of definitions we can see that the WEMI2 entity definitions are actually considerably different from those used for the original WEMI model. The rather specific lack of the verbs “*realizes*,” “*embodies*,” and “*exemplifies*” in the definitions suggests that we might not be able leverage the Group 1 ER-diagram (Figure 5.2 below) as Wickett and Renear did to give us clues as to which things might be taken as primitive and thereby help us kick start our formalization of the WEMI2 model.

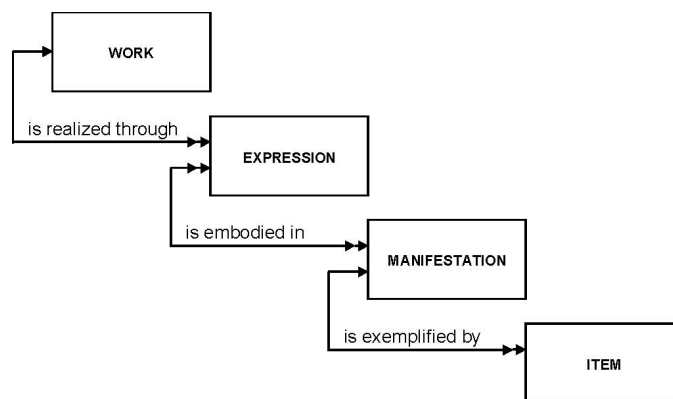


Figure 5.2: WEMI/WEMI2 ER-Diagram

We can confirm this by examining definitions given for the WEMI2 relationships:

- The *realizes* relationship “links a *work* with any of the *expressions* which convey the same intellectual or artistic content.” – Riva, Le Bœf, and Žumer p 65.
- The *embodies* relationship “links an *expression* with a *manifestation* in which the *expression* appears.” – Riva, Le Bœf, and Žumer p 65.
- The *exemplifies* relationship “connects a *manifestation* with any *item* that reflects the characteristics of that *manifestation*.” – Riva, Le Bœf, and Žumer p 66.

The linking relationships of *realizes*, *embodies*, and *exemplifies* cannot be interpreted as primitives since they are defined in terms of what role they play. Since WEMI2 *works* are defined as intellectual or artistic content, we should be relatively safe using *works* as a primitive. However, we can see from the quotations that WEMI2 *expressions* “convey” WEMI2 *works* and so the “is realized through” relationship in the figure does not seem to play a direct role in the *work-expression* relationship (unless it is identical with the “conveys” relation).

More specifically, the definition for the “is realized through” relationship tells us that “This relationship links a *work* with any of the *expressions* which convey the same intellectual or artistic content” – Riva, Le Bœf, and Žumer 2017, p 65, and does not clarify the role that the “is realized through” relationship plays in the WEMI2 model. *Linking* is rather a different conceptual relationship than *conveying* and so it does not seem that Figure 5.2 has a role in telling us exactly how we might formalize the entities such that they can fulfill the basic LRM user tasks.

We might push on and try anyway using the following as an initial definition and supporting axiom to provide an initial formal account of WEMI2 *expressions* and the WEMI2 *realizes* relationship.

We will take the concepts of WEMI2's *work* and *conveys* to be primitive.

D.18 (WEMI2 Expressions Defined): $\forall x \left(\text{Expression}(x) =_{df} \exists y (\text{conveys}(x, y) \wedge \text{Work}(y)) \right)$

A.43 (Realization of Works): $\forall x \forall y \left(\text{realizes}(x, y) \rightarrow (\text{conveys}(x, y) \wedge \text{Work}(y)) \right)$

Examining the scope note on the same page, we find additional information clarifying what is meant by “content”—“A *work* is an abstract entity that permits the grouping of *expressions* that are considered functional equivalents or near equivalents.” – Riva, Le Boef, and Žumer 2017, p 21.

It seems like we might be on the right track as the scope notes indicate that WEMI2 *works* are intended to inform criteria through which WEMI2 *expressions* are grouped.

Rather importantly, we are also told:

“A *work* comes into existence simultaneously with the creation of its first *expression*, no *work* can exist without there being (or there having been at some point in the past) at least one *expression* of the *work*.” – Riva, Le Boef, and Žumer 2017, p 21.

Now we have a worry. It might be the case that WEMI *works* are not primitives as they are burdened with constraints. We might formalize this constraint of a WEMI2 *work* through the following axiom.

A.44 (LRM Work Existential Constraint): $\forall x (Work(x) \rightarrow \exists y (conveys(y, x) \wedge Expression(y)))$

Compare this account to the one for FRBR that Wickett and Renear provide for Expressions (definition D.12 below; which takes *works* to be primitive).

D.12 (FRBR Expressions Defined): $\forall x (Expression(x) =_{df} \exists y (realizes(x, y) \wedge Work(y)))$

Two things immediately become apparent. First, where FRBR was vague enough that Wickett and Renear could take WEMI *works* to be primitive entities, LRM seems to close that door by requiring WEMI2 *works* to depend upon WEMI2 *expressions* for their existence. This implies the second thing—under the WEMI2 regime, there can be no *unconveyed works* (or any such works are outside of the model’s scope).

While this seems rather trivial at first, it is clear that WEMI2 is fundamentally different in contextual scope from FRBR’s WEMI. In the latter case, we would not be wrong to make an open-world assumption regarding the model’s context. Here, while there might be unrealized *works* somewhere out in the open world, such works are outside of the WEMI2 conceptual model’s scope. Only a closed world of *works* that have actually been *realized* will be accounted for by the WEMI2 model. These closed world constraints extend to the other aspects of bibliographic entities.

For instance, examining the scope notes for WEMI2 *expressions*, we are told that, like WEMI2 *works*:

“An *expression* comes into existence simultaneously with the creation of its first *manifestation*, no *expression* can exist without there being (or there having been at some point in the past) at least one *manifestation*.” – Riva, Le Bœf, and Žumer 2017, p 23.

There seems to be an additional existence constraint on WEMI2 *expressions* but, notice the lack of any references to verbs like “embodies.” We might formalize this constraint for WEMI2 *expressions* using the following axiom.

A.45 (LRM Expression Existential Constraint): $\forall x(Expression(x) \rightarrow \exists y Manifestation(y))$

Examining the definition for the “is embodied in” relationship from Figure 5.1, we see that Riva, Le Bœf, and Žumer tell us, “This relationship links an *expression* with a *manifestation* in which the *expression* appears” – p 65. Rather like the “is realized through” relationship, the “is embodied in” relationship would seem to be a conceptually weaker role in the model (i.e., “linking”) than that suggested by the definition’s other verb—“appears.” However, as we did with the *realizes* relationship, we can suggest an initial formalization of the *embodies* relationship through the following axiom.

A.46 (Embodiment of Expressions): $\forall x \forall y (embodies(x, y) \rightarrow (appearsIn(y, x) \wedge Expression(y)))$

At this point, it is tempting to replace the consequent in axiom A.45 with the antecedent in the above axiom. But since we cannot take the relationship named by that antecedent, it will likely be more informative if we reuse the consequent in the above axiom as the consequent of axiom A.45 so that the role that WEMI2 *manifestations* play with regard to WEMI2 *expressions* is much clearer.

A.45' (LRM Expression Existential Constraint): $\forall x (Expression(x) \rightarrow \exists y (appearsIn(x, y) \wedge Manifestation(y)))$

However, the fact that we needed to make the clarification at all is signal that our approach is not going to be as smooth as the one take by Wickett and Renear (2009). The problems for us coalesce when we consider how we might formalize the WEMI2 account of *manifestations*.

We are told that a WEMI2 *manifestation* is:

“A set of all carriers that are assumed to share the same characteristics as to intellectual or artistic content and aspects of physical form. That set is defined by both the overall content and the production plan for its carrier or carriers[.]” – Riva, Le Bœf, and Žumer 2017, p 25.

This definition is difficult to formalize in the same manner that we took for WEMI2 *expressions*. One of the problems is scope. While *works* and *expressions* are constrained through additional axioms, WEMI2 *manifestations* do not seem to be and sets certainly cannot be. And so, any formalization making a strong use of a bidirectional connective

here would seem to be false, unless the domain of discourse described by WEMI2 is limited to such an extent that when the term “set” is invoked, Riva, Le Bœf, and Žumer really mean “manifestation.” It is doubtful that this is actually the case and so we will propose softer axioms instead.

We might then formalize this account of WEMI2 *manifestations* through the following axiom.

A.47 (LRM Manifestations Semi-Defined): $\forall x(\textit{Manifestation}(x) \rightarrow \textit{Set}(x))$

Unfortunately, the above axiom does not account for the part of natural language definition that tells us that *manifestations* are not just sets, they are curated sets, grouped together through specific criteria. Unlike the collectives Winston, Chaffin, and Herrmann describe, or the bibliographic aggregates we discussed in the previous chapter, the parts of a *manifestation* are grouped together through the ZFC *set-member-of* relation, which does not imply that the set members are all homogenous according to some curatorial criteria. We will need to modify axiom A.47 to account for this by making the following formalization.

A.47’ (LRM Manifestations Semi-Defined): $\forall x(\textit{Manifestation}(x) \rightarrow \textit{CuratedSet}(x))$

From here, things get quite complex. We might try to gain further traction by examining the “is exemplified by” relationship, which Riva, Le Bœf, and Žumer define as one that connects a *manifestation* with any *item* that reflects the characteristics of that *manifestation*” – Riva, Le Bœf, and Žumer 2017, p 66. This natural language definition seems extremely peculiar since it is highly

focused on something called *characteristics*. We can turn to the scope notes for WEMI2 *manifestations* for more information. One of them tells us:

“A *manifestation* results from the capture of one or more *expressions* onto a carrier or a set of carriers. As an entity, *manifestation* represents the common characteristics shared by those carriers, in respect to both intellectual content and physical form.” – Riva, Le Bœf, and Žumer 2017, p 25.

Similarly, the scope notes for WEMI2 *items* tell us that:

“In terms of intellectual or artistic content and physical form, an *item* exemplifying a *manifestation* normally reflects all the characteristics that define the *manifestation* itself.”

So, it would seem that WEMI2 *manifestations*, a kind of **set**, represent *characteristics* which are *shared* among the *carriers* that are elements of the *manifestation* and that, when in the role of *exemplifying* the *manifestation*, *reflect* those *characteristics*. Unfortunately, here Riva, Le Bœf, and Žumer chose to introduce a new term, “carriers.” Before proceeding further, we need to equate this term to WEMI2 *items*, which are the entities that Riva, Le Bœf, and Žumer mean to invoke. We can do this by examining the definition provided for WEMI2 *items*. Recall that Riva, Le Bœf, and Žumer tell us that WEMI2 *items* are “objects carrying signs intended to convey intellectual or artistic content” – p 27. So, an initial formalization for WEMI2 *items* might look like the following definition:

D.19 (LRM Items Defined): $\forall x (Item(x) =_{df} \exists y (carries(x, y) \wedge Expression(y)))$

Quite literally, WEMI2 *items* are carriers that carry WEMI2 *expressions*. The scope notes discussing *items* provide important additional information, telling us that WEMI2 *items* are specifically physical objects that carry signs which are intended to convey content. Having clarified what is meant by the term “carriers,” we can now attempt to formalize the WEMI2 account of *manifestations* through the following axioms.

A.48 (Manifestations as Representations): $\forall x (Manifestation(x) \rightarrow \exists y (represents(x, y) \wedge Characteristics(y)))$

A.49 (Items as Reflections): $\forall x (Item(x) \rightarrow \exists y (reflects(x, y) \wedge Characteristics(y)))$

A.50 (Exemplification of Manifestations): $\forall x \forall y (exemplifies(x, y) \rightarrow \exists z (reflects(x, z) \wedge represents(y, z) \wedge Manifestation(y) \wedge Characteristics(z)))$

There is a further constraint when *manifestations* have more than two members. From the first scope note, we see that the members are said to share characteristics. To properly formalize this sharing relationship, we would normally want to employ a three-place predicate (e.g., a shares b with c). However, we can employ axiom A.47’ to produce a serviceable, if cumbersome work around.

A.51 (Manifestation Members Share Characteristics):
 $\forall x \forall y \forall z ((setMemberOf(x, z) \wedge setMemberOf(y, z) \wedge Manifestation(z)) \rightarrow \exists w (reflects(x, w) \wedge reflects(y, w) \wedge Characteristics(w)))$

Despite this further complication, we can propose a formal definition for Manifestations through the following biconditional.

D.20 (Manifestations Defined):

$$\forall x \left(\text{Manifestation}(x) =_{df} \exists y \exists z \left(\text{represents}(x, y) \wedge \text{reflects}(z, y) \wedge \text{setMemberOf}(z, x) \wedge \text{Characteristics}(y) \wedge \text{Item}(z) \right) \right)$$

Now though we have a further question regarding how WEMI2 *manifestations* are to be defined. Specifically, we know that WEMI2 *expressions* play a particular role with regard to WEMI2 *manifestations*. We likely need to account for this role by modifying definition D.20.

D.20' (Manifestations Redefined):

$$\forall x \left(\text{Manifestation}(x) =_{df} \exists y \exists z \exists w \left(\text{appearsIn}(w, x) \wedge \text{represents}(x, y) \wedge \text{reflects}(z, y) \wedge \text{setMemberOf}(z, x) \wedge \text{Expression}(w) \wedge \text{Characteristics}(y) \wedge \text{Item}(z) \right) \right)$$

We also now have a worry about our definition (D.19) for WEMI2 *items*. Specifically, how important are their characteristics? With regard to their role in WEMI2 *manifestations* (beyond the *set-member-of* relationship), the fact that WEMI2 *items reflect characteristics* which are *represented* by WEMI2 *manifestations* provides the curatorial glue that distinguishes WEMI2 *manifestations* from any other arbitrary set. However, the existence of *characteristics* themselves does not seem to be necessary for a formal definition of WEMI2 *items* because the natural language account never invokes WEMI2 *manifestations*. Since we already have axiom A.48, and now definition D.20', the role that WEMI2 *items* play regarding WEMI2 *manifestations* is accounted for. And, at this point we seem to have accounted for all of peculiarities in LRM's WEMI2 model.

In the next section, we discuss the implications these peculiarities have for metadata infrastructure. However, before we proceed, we will produce an illustration (Figure 5.3 below) that will allow us to take stock of all of the new entities and properties that we have introduced. This illustration uses Chen’s method for ER-diagrams with some notable exceptions. As in Figures 5.1 and 5.2, we are neglecting to use the proper (diamond) symbol to represent relationships. Furthermore, we are electing to, for the time being represent concepts like *manifestation*, *curated sets*, and *characteristics* as though they are (first-class) entities. However, we are also going to shade them in as a means to indicate that all three entities seem to be doing something relationship like among the other entities.

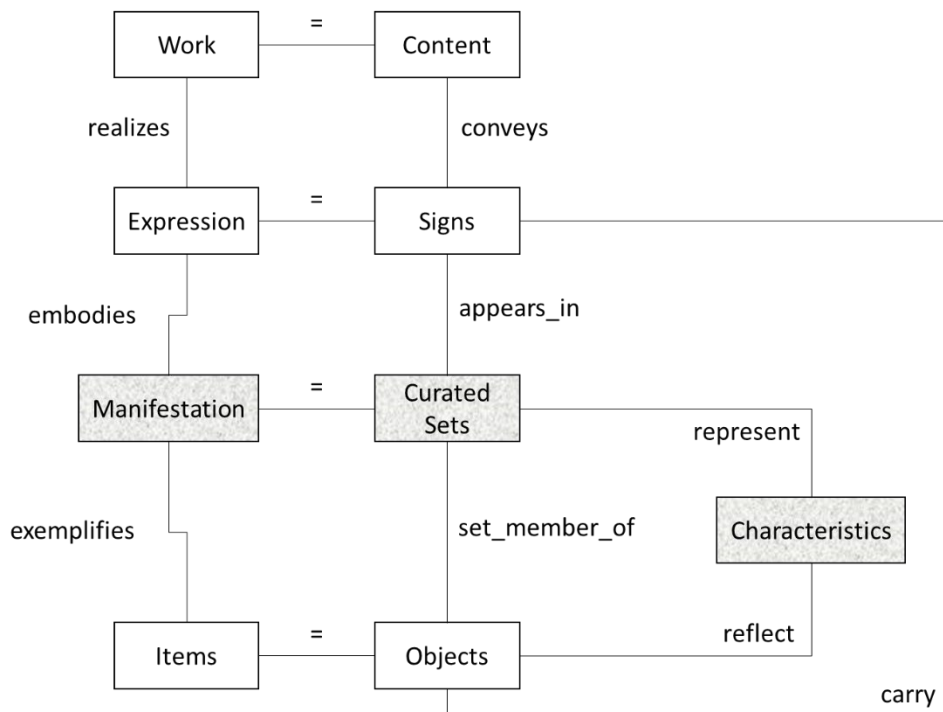


Figure 5.3: Illustration of Entities and Relationships Involved in WEMI2 Model

As we can see from Figure 5.3 (above), we are dealing with a model that has both more and different primitives than the initial WEMI model that appeared in FRBR. Relationships such as

realizes, embodies, and exemplifies are can no longer be taken as primitive when other primitive relationships, that more clearly distinguish the roles of the entities involved, appear in the model. And so, with regard to relationships, we take the following ones to be primitive: *conveys, appears-in, set-member-of, represents, reflects, and carries*.

Concerning entities, we must take the *characteristics* entity to be a primitive one for now, as other entities rely on it as a partial explanation for their nature. However, we have a choice regarding this model's other primitive. We could take the *work* (or *content* since it is equivalent) entity to be primitive, as Wickett and Renear (2009) do. Or, we could take the *item* (or *object* since it is equivalent) entity to be primitive.

This is a choice which affects the kinds of implications and overall world view orientation of the WEMI2 model, i.e., it is a sweeping interpretation. If we take *items* to be primitive, we will essentially be making an ontological commitment to an anti-realist position where entities like *works* are in fact *roles* played by *objects*. This position is actually consistent with reductionist models like Svenonius that liken *works* to *sets of documents* and the standard approach to cataloging which focuses on artifacts in hand. Several of the constraints, such as the injunction against unrealized *works*, also support this interpretation. Implications that might occur from this ontological position include *manifestations* that *embody no expression* and *expressions* which realize no *work*.⁶⁶

⁶⁶ Note that because of the intrusion of set theoretics through the entity of WEMI2 *manifestation*, there can be no *items* that do not *exemplify* some *manifestation*. This is due to the fact that *manifestations* are sets and that *items* are elements of them. We discuss these implications further in the sections below.

However, we have already produced definitions that take certain relationships into account that make the ontological commitment to the primitiveness of *works* and thereby a realist account of the bibliographic universe. This position is supported by the natural language definitions themselves. Of the definitions in LRM, only the definition of *works* references a singular entity, *content*, directly. All the other definitions are combinations of entities and relationships, e.g., an *expression* is “signs conveying content.” So, we will add a second primitive entity (beyond *characteristics*) to our account, and it is *work*.

5.4. Discussion of the LRM Aggregate Model

Now that we have established an understanding of the *content-artifact model* (WEMI2) set forth in LRM, we can finally explore its impact on the *content-artifact aggregate model* that LRM employs. One thing we should note is that the *content-artifact aggregate model* we will be discussing is not different from the one set forth by the FRBR-WGA (2011) and Žumer and O’Neill (2012), in the sense that both the one we will be discussing and the one set forth by these author groups both define *bibliographic aggregates* as *manifestations*. However, what is meant by *manifestation* is altogether different. In the case of the FRBR-WGA and Žumer and O’Neill, *manifestation* refers to WEMI *manifestation* as set forth in FRBR. In the case we are examining here, the LRM case, *manifestation* refers to WEMI2 *manifestation*, i.e., a curated set.

The first and most immediate implication is that *bibliographic aggregates* in this case are exactly curated sets (and not merely sets in a particular role at a particular time). With regard to what kinds of things are allowed to members of a *bibliographic aggregate*, there are distinct limitations, and as we show presently, this causes this particular version of the *content-artifact aggregate model* to break down immediately.

We will take as an example an issue of a journal. The issue of a journal is a mass-produced object such that it has certain physical characteristics like the number of pages in it, the number of articles in it, and the like. In so far as the WEMI2 model is concerned, the *manifestation* of this particular journal issue is the set of all copies printed, i.e., the set of all *objects sharing* the same *characteristics*. This all seems quite reasonable. So far, so good.

However, based on definition D.17 which tells us that *bibliographic aggregates embody* the individual *expressions* of what is, in this case, the articles in the journal issue. This would indicate that the *items* then are actually the articles in the journal issue. This would seem to follow from definition D.19 which tells us that *items carry* expressions and parts of definition D.20' which tells us that *items reflect characteristics* that *manifestations represent*. Unfortunately, journal issues and articles do not *share characteristics*, and so axiom A.51 which articulates a constraint that the members of a *manifestation* must *share characteristics*, would seem to be violated.

Now we have quite a conundrum, as it does not seem as though it is actually possible for our *bibliographic aggregate* to *embody* the individual articles in the journal issue on account that the *characteristics reflected* by the journal issue *items* are different than those *reflected* by the article *items* and thereby definition D.20' is violated.

We might try an easy repair by rejecting axiom A.51. This works, but it also defeats the purpose of the WEMI2 model, as we would no longer be able to distinguish *items* or *manifestations* based on *characteristics*. More pragmatically, we can confidently say that rejection of axiom A.51 makes the separation of editions whose differences are at the *item*-level, e.g., a mass-market paperback

edition versus a trade-paperback edition of the same *expression*, impossible to distinguish from one another—they all become members of the same *manifestation*.

We might try a different repair by rejecting definition D.19 which tells us that *items* are *carriers*. However, this defeats the purpose of the *content-artifact aggregate model*, because now we only have *items* that correspond to the journal issues, but we have no way to relate them to the *expressions* of the articles because the articles no longer have *items reflecting characteristics* unique to them. This solution does not appear to be desirable either.

We might try another repair by introducing a *part-of* (i.e., a mereological) relationship to the model. In fact, we can see upon further examination that Riva, Le Bœf, and Žumer have already done this for the *manifestation* entity. The domain and range of this *part-of* relationship is quite specifically limited to WEMI2 *manifestations*.⁶⁷ Since WEMI2 *manifestations* are curated sets, we would expect that this particular *part-of* relation, which we will hereafter refer to as *[manifestation]-part-of*, is the same as the ZFC *subset-of* relation.

Riva, Le Bœf, and Žumer define *[manifestation]-part-of*, by telling us:

“This is a relationship between two *manifestations* where one is a component of the other” – Riva, Le Bœf, and Žumer, p 75.

⁶⁷ We should note here that LRM also defines *part-of* relationships whose collective domains and ranges are constrained to *works* and *expressions* respectively, i.e., there are *[work]-part-of* and *[expression]-part-of* relationships. We discuss these further in section 5.7 below.

So far, interpreting *[manifestation]-part-of* as a relationship equivalent to ZFC *subset-of* seems within the realm of possible and charitable interpretations. However, in the scope notes Riva, Le Boef, and Žumer say:

“In some cases[,] the components of a *manifestation* are based on physical considerations relating to the carrier in which the *manifestation* is intended to be issued (for example, a recording is too long to fit on a single disc and is issued in a two-disc boxed set). An alternate *manifestation* on another carrier may not display the same components.” – p 75.

The scope note does not seem to be talking about subsets. In fact, it seems to be talking about aggregate *items* (or *items* that are sets). This is an odd state of affairs. We were expecting something to do with subsets. So, this existing *[manifestation]-part-of* relationship does not help us since it actually captures situations where a journal issue itself is a multi-*item* entity, such as the case of a journal issue and a supplementary issue specific to it. In this case the *manifestation* would be the set of all issue/supplement pairs.

We could try a fourth repair by introducing an *[item]-part-of* relationship which LRM does not actually possess. This allows us to draw direct links between the characteristics of the articles and those of the journal issues, but it does not seem to overcome the problem that these are still different characteristics.

At this point, we could go on and try a fifth repair by resorting to set theoretic rules directly and supposing that the *item-sets representing the characteristics* of the articles are actually subsets of the journal issue *manifestation*. Unfortunately, this still does not work because the characteristics *reflected* by subsets are not the same as the characteristics *represented* by the *manifestation* of the journal issue. While our analysis here is not exhaustive, after five failed attempts to find a satisfactory repair for this issue, we feel that we have demonstrated the severity of this modeling issue. This is not a minor flaw that is easy to fix, but rather is indicative of a fundamental flaw that occurs through the pairing of the *content-artifact aggregate model* used in LRM with the particular WEMI entity definitions used in LRM.

We should note that quite importantly, since the *content-artifact aggregate model* breaks down in the case of LRM precisely because of the manner in which *manifestations* are defined in LRM's WEMI2 *content-artifact model*, we fully expect that application of the same *content-artifact aggregate model* in the context of a different *content-artifact model*, e.g., FRBR's original WEMI model, does not experience this break down. However, there are still unresolved issues for *content-artifact aggregate model*, regarding contextual mass and content (Jett, Fenlon, and Downie 2018) and for paratext (Coyle 2016b). At least the first of these two issues can be overcome by using FRBR's existing *mereological aggregate model*, which we analyzed in the previous chapter.

5.5. The WEMI2 Content-Artifact Model

5.5.1. Developing an Alternate Illustration

As Figure 5.3 implies, the actual *content-artifact model* set forth in the pages of LRM is nothing like the *content-artifact model* suggested by Figure 5.2. In this section, we will endeavor to develop a refined illustration, employing as few non-primitive entities and relationships as possible. This

will be useful because it will further illustrate some of WEMI2's idiosyncrasies in comparison to other *content-artifact models*, like the original WEMI. It will also highlight several of the idiosyncrasies of the original WEMI model.

Since the *realizes*, *embodies*, and *exemplifies* relationships are not primitive ones, we will at first simplify the pseudo-model illustrated in Figure 5.3 by eliminating them in favor of primitive relationships. In the case of the *realizes* relationship, we can substitute the primitive *conveys* relationship. In the case of the *embodies* relationship, we will substitute the no less mystifying *appears-in* relationship.⁶⁸ We also want to replace the *exemplifies* relationship, but because *manifestations* are sets whose members are *items*, this task is more onerous. It requires a decomposition of the *exemplifies* relationship into the primitive entity of *characteristics* and the primitive relationships of *set-member-of*, *represents*, and *reflects*. Finally, we will drop the equivalent entities of *content*, *signs*, *curated sets*, and *objects* so that we can maintain as much core terminology (i.e., *works*, *expressions*, *manifestations*, and *items*) as possible.

5.5.2. Discussion of the “conveys” Relationship

Definition D.18 (WEMI2 *expressions convey* WEMI2 *works*) and axiom A. (a WEMI2 *work* only exists if there exists a WEMI2 *expression* that *conveys* it) tell us how WEMI2's *works* and *expressions* relate to one another. We know from Figure 5.2 that the “is realized through” relationship is intended also link WEMI2 *works* and *expressions*. The *conveys* relationship is more

⁶⁸ Here though, it is tempting to claim that we still need the *embodies* relationship on account of the *content-artifact aggregate model*. However, we amply demonstrated that this model does not work in the context of WEMI2 and so we are safe (at least until someone offers a repair by means of strong counter-example) to exclude it altogether. In fact, for the moment we will table any further discussion of *bibliographic aggregates* until we arrive at a refined illustration for the WEMI2 model.

informative though because we know from the natural language definitions what it means in the context of WEMI2. More specifically, it tells us that the “signs” that are WEMI2 *expressions* have the role of *conveying* the “content” that are WEMI2 *works*.

An open question is whether or not the *conveys* relationship should be interpreted as having the same cardinality that Figure 5.2 illustrates the “is realized through” relationship as having. Axiom A.43 takes a neutral stance by using the material implication connective rather than the equivalence biconnective. There is a good reason for Axiom A.43 to take a neutral stance. Specifically, intuitions that the *realizes* relationship should be restricted to many-to-one *expression-to-work* cardinality may be mistaken. Jett and Dubin (2018), examining the issue of whether expressions *realize* more than one work in the case of dependent works, provide an interesting discussion of one particular case.

In their paper, “How are dependent works realized?”, Jett and Dubin consider seriously how the editorial contributions made by a humanities scholar transcribing and adding TEI markup to text relate to the text being transcribed. The specimen examined is Molly O’Hagan Hardy’s TEI transcription⁶⁹ of Absalom Jones and Richard Allen’s 1794 *A Narrative of the Proceedings of the Black People, During the Late Awful Calamity in Philadelphia, in the Year 1793*. As Jett and Dubin point out, there is at least one account wherein it is possible to understand the text of the transcribed document playing two distinct *expression* roles. In one of these roles, the transcribed text *realizes* Jones and Allen’s work. In the other role, the transcribed text *realizes* Hardy’s scholarly *work* which makes claims about Jones and Allen’s *work*.

⁶⁹ <http://tapasproject.org/proceedingsofblackpeople/files/narrative-proceedings-black-people-during-late-awful-calamity>

A potential problem for Jett and Dubin might be that these different roles are actually different instances of the *realizes* relationship. However, the cardinality constraint on the *realizes* relationship entails that while a *work* might have the property of being the object of many different instances of the *realizes* relationship, an *expression* can only ever be the subject of a single instance of the *realizes* relationship. The entire point that Jett and Dubin are making is that, despite this limitation, it seems as though there are situations where *expressions* should participate in more than one *realizes* role. Hence, there seems to be a problem with this cardinality constraint.

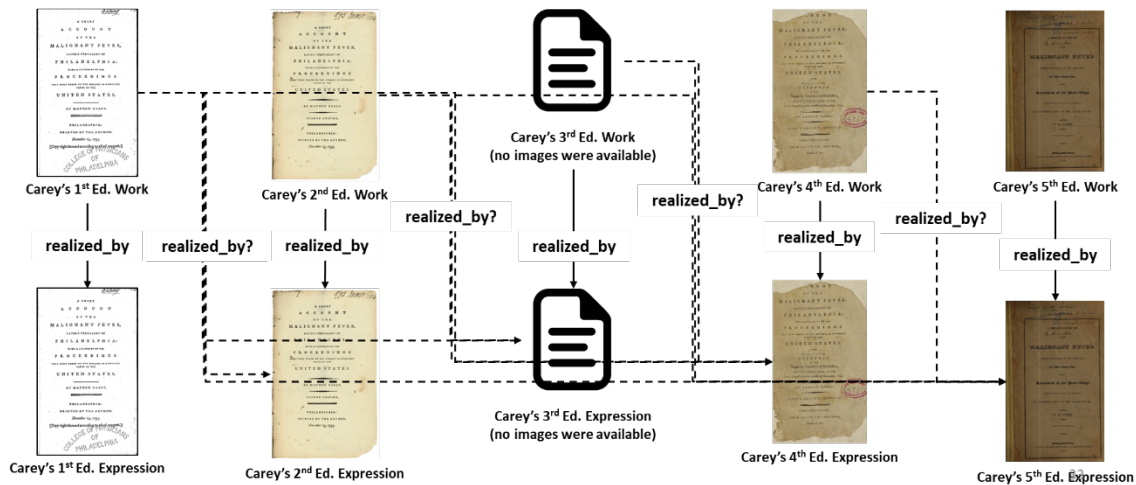


Figure 5.4: Realizing Different Editions of Carey

This can be further illustrated when we note that Jones and Allen's *work* is itself something of a dependent *work*, as it was created in response to a pamphlet published by Mathew Carey, his 1793 *Account of the Malignant Fever Lately Prevalent in Philadelphia*. During the months when the epidemic was occurring, and for some months afterwards, Carey published five editions of his pamphlet. What is of interest here is that each edition possessed expanded content, growing from some ~90 pages in the 1st edition to more than 160 pages in the 4th edition. It is obvious that we can account for the differences in intellectual content by claiming that the *work* that is each edition

is derived from the *work* of the previous edition. But it is also equally clear that the expression of each edition likely *realizes* not only the *work* of its edition but also the *works* that are the preceding editions. Figure 5.4 (above) illustrates this conundrum.

As Jett and Dubin note with the case of Carey’s *Account of the Malignant Fever Lately Prevalent in Philadelphia*, enlargement of a document can entail an increasing of its intellectual content by 100% or more. Can we truly say that Carey’s 1st edition shares “essentially the same information” as his 4th edition? Are they truly the same *work*?

To some extent, traditional cataloging practice renders much of this discussion moot, since Carey’s *Account of the Malignant Fever Lately Prevalent in Philadelphia* would be treated as a monograph case, and thereby, each of its editions would receive a distinct catalog record. However, such is not the case for one of the cataloging traditions greatest exceptions to its own rules—serial publications.

Serial publications are generally considered to be singular, ever-expanding (for so long as they are being published) *works*. The general theory made for this style of representation is that the nature of their content does not significantly vary over time, i.e., all of their issues, volumes, or editions share “essentially the same information.” In reality though, nothing could be farther from the truth, and Figure 5.4 is merely one example that showcases the problem.

However, if we accept that our primitive *conveys* relationship possesses many-to-many cardinality, then this issue can be sidestepped altogether. Since we already have established that *conveys* and “is realized through” are not the same relationship, there seems to be no apparent problem if we

conjecture that it has many-to-many cardinality and thereby avoid the problem that Jett and Dubin showcase in the example above. We can illustrate this through the following ER-diagram.




Figure 5.5: ER-Diagram Illustrating the *conveys* Relationship

5.5.3. Discussion of the “*appears-in*” Relationship

Axioms A.45’ (a WEMI2 *expression* only exists if there exists a WEMI2 *manifestation* that it *appears-in*) and axiom A.46 (something *embodies* something else when that something else is a WEMI2 *expression* that *appears-in* it [the first something]) are the only significant information that tells us how WEMI2’s *expressions* and *manifestations* relate to one another. We know from Figure 5.2 that the “is embodied in” relationship is intended also link WEMI2 *expressions* and *manifestations*. However, the *appears-in* relationship is somewhat more informative, because we know from the natural language definitions that it means more than just linking in the context of WEMI2. More specifically, it tells us that the “signs” that are WEMI2 *expressions* have the role of *appearing-in* the “curated sets” that are WEMI2 *manifestations*.

Here, an idiosyncrasy from Figure 5.2 is helpful. We see that the “is embodied by” relationship has many-to-many cardinality. Like the *conveys* relationship, we have good reason to believe that the *appears-in* relationship has many-to-many cardinality. However, because the “is embodied by” relationship has many-to-many cardinality, we do not need to argue the *appears-in* relationship also has many-to-many cardinality. Having such cardinality would be consistent with the account of other relationships (i.e., “is embodied by”) that link WEMI2 *expressions* with WEMI2 *manifestations*. We can illustrate this through the following ER-diagram.




Figure 5.6: ER-Diagram Illustrating the *appears-in* Relationship

5.5.4. Discussion of the “*carries*” Relationship

Definition D.19 (WEMI2 *items carry* WEMI2 *expressions*) tells us how WEMI2 *items* relate to WEMI2 *expressions*. At first glance, this seems like an odd definition for us to have proposed in the first place since the original WEMI model has nothing like it. However, the natural language definition for WEMI2 *items* (“An object or objects carrying signs intended to convey intellectual or artistic content” – Riva, Le Boef, and Žumer 2017, p 27) mentions WEMI2 *expressions* directly (i.e., signs conveying content) and WEMI2 *manifestations* not at all.

The question that faces us, once again, is what kind of cardinality should the *carries* relationship possess? Is it possible for an inscription to simultaneously *carry* text from more than one *expression*?

Here, a commonplace example of after-market binding of journal issues into journal volumes (or partial journal volumes or occasionally journal multi-volumes) is helpful. We can observe that if the *content-artifact aggregate model* is correct, then there is no corresponding *expression* for the journal volume (indeed there is no corresponding *expressions* for the journal issues bound into the volume either) and so it must be the case that the bound journal volume is acting in the role of *carrying* the signs which are the WEMI2 *expressions* that correspond to the individual articles in the issues that have been bound into the volume. It seems to be the case that we have a single WEMI2 *item* which *carries* multiple WEMI2 *expressions* that correspond to the articles comprised by the bound journal volume.

From this example, we can see that our *carries* relationship should also possess many-to-many cardinality so that our model can accommodate the example we just showcased. We illustrate this through the following ER-diagram.




Figure 5.7: ER-Diagram Illustrating the *carries* Relationship

5.5.5. Discussion of “Manifestations” and Their Particulars

5.5.5.1. Set Membership

We know from Axiom A.47’ that WEMI2 *manifestations* are by definition sets (curated sets but sets nonetheless). We also know that their members are WEMI2 *items*. Therefore, there must be a *set-member-of* relationship that links them. Furthermore, we know from set theoretics that this *set-member-of* relationship must possess many-to-many cardinality. We can illustrate this with the following ER-diagram.




Figure 5.8: ER-Diagram Illustrating the *set-member-of* Relationship

5.5.5.2. Discussion of the “Characteristics” Entity and its Linking Relationships

From axiom A.48, we know that WEMI2 *manifestations* represent something called *characteristics*. Similarly, we know from axiom A.49 that WEMI2 *items* reflect the same *characteristics* that are represented by WEMI2 *manifestations* that they are in a *set-member-of* relationship to. It seems clear that many *items* might reflect many different *characteristics*, but are there some *characteristics* which are only represented by a single *manifestation*? In other words, does the combinatorial of the *characteristics* entity and the *represents* and *reflects* relationships possess the many-to-one (*items-to-manifestations*) cardinality that the *exemplifies* relationship possesses?

Here, we will reintroduce the *content-artifact aggregate model*. We will ignore for the moment the problem of sets for this model to examine another peculiar problem. Let us consider author Terry Brooks' long-running fantasy series *Shannara* as an example. Figure 5.9 below illustrates most, but not all, of the series.

As we can see from the figure, Brooks' *Shannara* series is itself broken into smaller sub-series. The question that begs to be answered is, if we are a library patron and want an item that is part of one of these series, e.g., Brooks' *Running with the Demon* (because say, we have used its ISBN as our search criteria), what kind of search results will we receive?

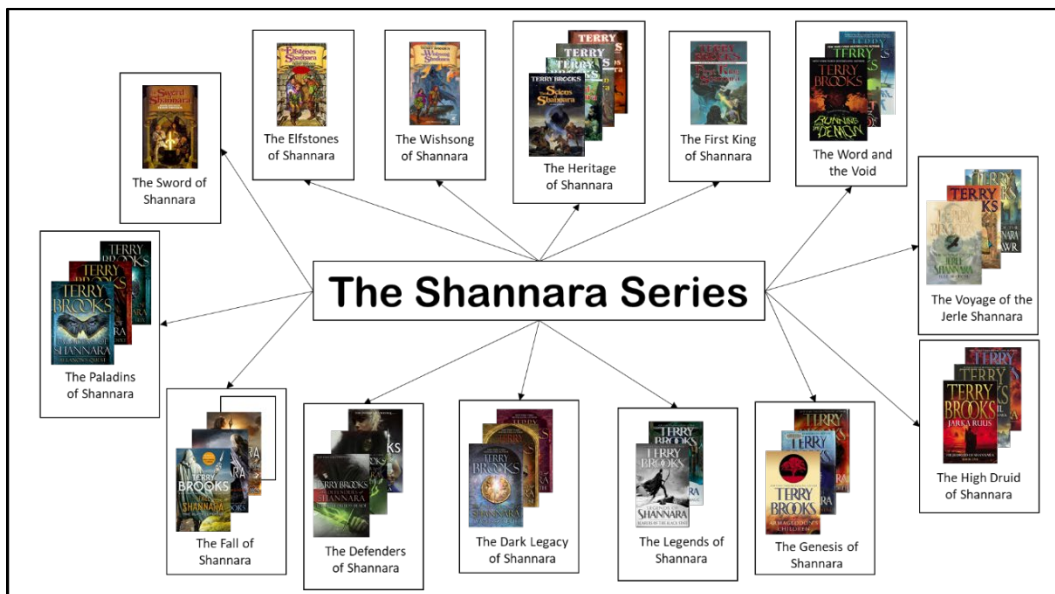


Figure 5.9: Brooks' *Shannara* Series and Most of Its Members

A faithful implementation of the *content-artifact aggregate model* should only display just one of the following bits of information to the end user (because we are using the “is exemplified by” relationship to manage the search results):

1. All the copies of the particular edition of *Running with the Demon* that we asked for, i.e., the *manifestation*-level representation of *Running with the Demon* that corresponds to our particular desired *item*.
2. Metadata describing Brooks' series, *The Word and the Void*, for which our particular desired copy of *Running with the Demon* corresponds to a particular *expression* of Brooks' *Running with the Demon* that has been aggregated (i.e., *embodied*) into the *manifestation*, Brooks' *The Word and the Void*.
3. Metadata describing Brooks' *Shannara* series, for which our particular desired copy of *Running with the Demon* corresponds to a particular *expression* of Brooks' *Running with the Demon* that has been aggregated (i.e., *embodied*) into the *manifestation*, Brooks' *Shannara* series.

In the vast majority of cases, we would expect that we will only see result number one of the three possibilities listed above. We should expect this, because this is the exact use case that WEMI, WEMI2, and the “is exemplified by” relationship are designed to accommodate.

In an ideal world, we would expect to get all three results. We can also observe that such links are made at least textually, if not digitally, through notes fields (e.g., in MARC records) in our existing (pre-FRBR) metadata systems. However, we can also observe that a faithful implementation of the FRBR model already supports all three of the use cases listed above. As we discussed in the previous chapter, the FRBR model can do this because its aggregate model is separate from its

version of the WEMI model. Not so with the *content-artifact aggregate model* that LRM employs. (And, also not so for any versions of FRBR that implement the *content-aggregate model* set forth by the FRBR-WGA [2011] and Žumer and O’Neill [2012]—the *content-artifact aggregate model*, even though not officially adopted by FRBR, was originally intended for FRBR, after all.)

The above problem can be compounded if we consider one of the members of Brooks’ *Shannara* series that was not illustrated in Figure 5.9 (above)—Brooks’ *Indomitable*—a short story written for and published in Robert Silverberg’s (ed.) anthology, *Legends II*. Our desired item is the short story, but the corresponding *manifestation* does not describe it. Instead, it describes the anthology *Legends II* which aggregates (e.g., *embodies*) an *expression* that corresponds to the item required to read Brooks’ *Indomitable*. Here, we must admit that this search is unlikely to be successful unless there were notes listing *Indomitable* as one of the *expressions embodied-by* the *manifestation*, *Legends II* and it is also the case that we are full-text searching through our metadata records (as notes are not typically sufficient to create actionable digital links of their own accord). Once again, a faithful implementation of FRBR sidesteps this issue because the FRBR aggregate model is separate from its version of the WEMI model.

For our discussion here, we can also sidestep the issues showcased by the example simply by accepting that both the *represents* and *reflects* relationships possess many-to-many cardinality. We represent this with the following ER-diagram. Since the *characteristics* entity appears to play a mostly linking role, we will also represent it using the notation for associative entity.




Figure 5.10: ER-Diagram Illustrating the *Manifestation* Entity Complex

Unfortunately, we have one potentially vexing issue from Section 5.3 that is illustrated in Figure 5.2—*manifestations* also seem to play an associative role. This is especially true in light of the aggregate model illustrated in Figure 5.1 wherein the *manifestation* is intended to associate a group of different *expressions* with a group of *items* that correspond to those expressions. In that light, we should really use the symbol for associative entity for WEMI2 *manifestations* also. However, since the intention of LRM’s authors is to provide as much backward compatibility with WEMI and FRBR, and since they likely envision WEMI2 *manifestations* as first-class entities in their own right, much the same as we tend to envision *bibliographic aggregates*, we do not use the symbol for associative entity to represent the *manifestation* concept in Figure 5.10.

5.5.6. Simplified WEMI2 Illustration

When we assemble the various ER-diagrams from Figures 5.5, 5.6, 5.7, 5.8, and 5.10 into a single Figure (5.11 below), we can see that the WEMI2 model is very different from the depiction of it in Figure 5.2. This is not so surprising. The natural language definitions which we were supplied with were much more precise than those supplied for the original WEMI model in FRBR (IFLA 1998).

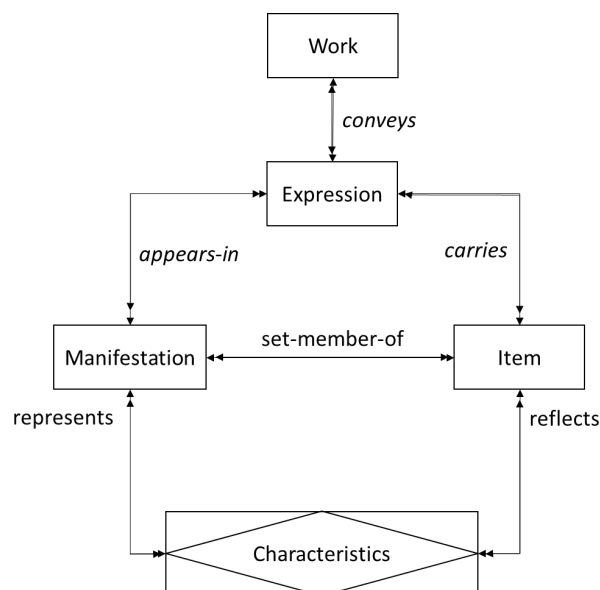


Figure 5.11: WEMI2 Model According to Natural Language Definitions

In the next two sections, we discuss the implications of this *content-artifact* model for *bibliographic entities*. We also strongly consider a simplification to it—removal of the *manifestation* complex—that sidesteps many of the issues faced by the *content-artifact aggregate model*. Sidestepping issues through simplification is possible because the overarching LRM model details a second, *mereological* aggregate model. We should note that removal of the WEMI2 *manifestation* complex directly results in removal of the *content-artifact aggregate model*. This is likely beneficial, since we no longer have competing accounts for *bibliographic aggregates*.

5.6. Initial Findings

5.6.1. Trivial Implications

There are several trivial conclusions that we can draw from the preceding two sections. For one thing, when applied to LRM’s WEMI2 *content-artifact model*, the *content-artifact aggregate model* developed by the FRBR-WGA ceases to be just a *content-artifact aggregate model*. It becomes a mixed *mereological/content-artifact aggregate model*. This occurs because **sets** are a

kind of *mereological* entity and obey their own specific *mereological* rules. By extension, this makes the *bibliographic aggregates* described in LRM *mereological entities*, at least in part.

As this is the case, then it seems that LRM's aggregate model (the *content-artifact aggregate model*) is not necessarily unjustified as WEMI2 is defined. Since WEMI2 *manifestations* are sets, i.e., they are *aggregates*, then deciding that all *bibliographic aggregates* are *manifestations*, rather than distinct bibliographic entities with distinct content (e.g., contextual mass), seems to be in accord to overall model for bibliographic entities. However, as we noted in Section 5.3, the approach breaks down because the *content-artifact aggregate model* wants to do things which **sets** are not designed to accommodate.

Moreover, as Jett, Fenlon, and Downie (2018) point out, the *content-artifact aggregate model* itself can have consequences that negatively impact the ability of users to successfully find, identify, select, and obtain *bibliographic aggregates* like anthologies, journal volumes, digital collections, and so on. After all, *bibliographic aggregates* typically seem to be kinds of bibliographic entities, equal in status to more familiar ones like books. We have already seen arguments that creating metadata records, for digital collections at least, helps digital library users assess their fitness for their own purposes (Brack et al. 2000; Sweet and Thomas 2000; Foulonneau et al. 2005; Palmer et al. 2006), and a vital component of these kinds of metadata records is description of the digital collection's overall topicality, or as Palmer (2004) and Palmer, Zavalina, and Fenlon (2010) would put it, the aggregate's contextual mass.

It is possible that some of reasons behind the FRBR-WGA's decision-making process in arriving at the *content-artifact aggregate model* is due to a misunderstanding regarding the original WEMI model that occurred among researchers, developers, and practioners alike. For some reason, it seems as though WEMI's concepts of *work-expression-manifestation-item* are confused as kinds of bibliographic entities (this is in fact a feature of Svenonius's conceptual model [2000] that she touts), when in fact they express a particular decomposition of the concept of bibliographic entity into finer-grained component-like concepts (which is more obvious in Smiraglia's treatment [2001]).

Works, expressions, manifestations, and items are not bibliographic entities in their own right. Rather, each of them captures specific qualities that a bibliographic entity (like a book) possesses. They are all aspects of bibliographic entities. However, this misunderstanding is codified through the *content-artifact aggregate model* presented by Riva, Le Boef, and Žumer in LRM.

5.6.2. Non-trivial Implications

As Figure 5.11 illustrates, the WEMI2 model is actually quite convoluted compared to the original WEMI model set forth by FRBR in 1998. If we consider the things that the model articulates, they do not necessarily seem to go together. On one hand, we have a relatively strong account of a content-artifact continuum that speaks of content, signs that convey content, and objects that carry signs. On the other hand, we have sets of objects in which signs appear and which represent characteristics which are then reflected by the objects that are elements of those sets. And finally, we have an aggregate model that also reuses signs, objects, and sets to articulate an understanding of bibliographic aggregates.

We can state more strongly that since WEMI2 *items* are directly related to WEMI2 *expressions*, either it is the case the WEMI2 *manifestations* are playing a redundant role with respect to WEMI2 *items*, or WEMI2 *manifestations* play no intermediary role in the content-artifact continuum described by the triumvirate of *work-expression-item*. In other words, WEMI2 *manifestations* have no impact on the conceptual model described by the phrase: *objects carry signs that convey content*.

This is a rather important distinction, as it speaks to much of the real value of the LRM standard. If we choose to exclude *manifestations* from the model altogether, we arrive at a succinct and powerful *content-artifact model*, through which deduplication will be superior. It is superior because we can focus our efforts for deduplication along two narrow dimensions—objects carrying the same signs, and signs conveying the same content.

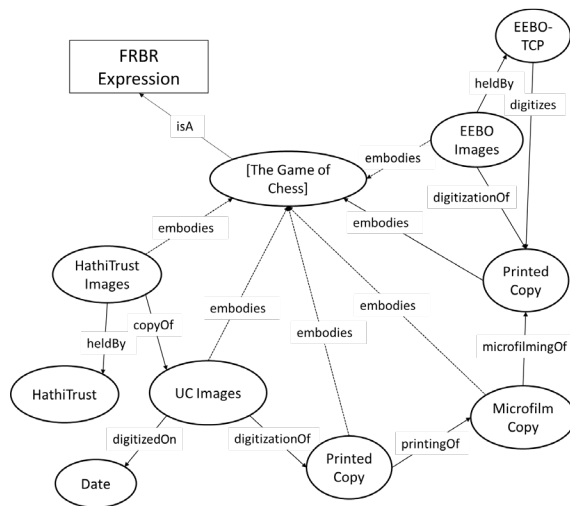


Figure 5.12: Different “manifestations” of [The Game of Chess] charting its reproduction cycle from the British Library through UC Berkeley to the HathiTrust Digital Library

We can demonstrate this by reusing an example from Chapter 2. Recall that Nurmikko-Fuller et al. (2015) discovered a curious example of a book owned by the British Library that had been

copied five times across three different formats. Nurmikko-Fuller et al. illustrated the situation through the Figure above (reproduced from Chapter 2 for convenience's sake). Figure 5.12 illustrates a view of [*The Game of Chess*] that conforms to the manner in which the original WEMI model is defined. However, we can produce an updated version of this figure (Figure 5.13 below) that conforms to what the natural language definitions of LRM's WEMI2 model tell us.

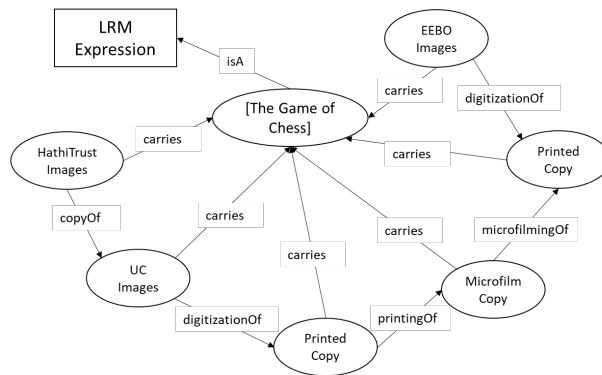


Figure 5.13: Different carriers of [*The Game of Chess*] charting its reproduction cycle from the British Library through UC Berkeley to the HathiTrust Digital Library

We have simplified Figure 5.13 (above) somewhat by removing some of the extraneous entities from the account. However, a key thing to notice is that WEMI2's *manifestation* is not present in Figure 5.13. Adding it would make the figure quite complex. We would need at least three *manifestations* in which the *expression* of [*The Game of Chess*] appears and in which the various *items* in Figure 5.13 reflect. We would also need three sets of *characteristics* upon which the *items* are aggregated through the *set-member-of* relationship. The question remains though: Do we need this kind of complexity in our conceptual model?

The answer is: no, we do not. Everything that the concept of *manifestation* accomplishes in our conceptual model is better accomplished on the implementation side, when we build the queries

that will craft the various records out of the data for the various end users and stakeholders. Databases are expressly designed for just this kind of division of labor.

In fact, a narrower account of bibliographic objects accomplishes the kind of functionality that Smiraglia, Svenonius, Tillett, and other conceptualizers of bibliographic entities had in mind all along. It suggests rather precisely what our metadata needs to describe—content, signs conveying it, and objects carrying the signs conveying it. This would also seem to confirm Taniguchi's position (2002, 2003) that our focus should be on describing the text (i.e., the *expression*) more than the variations between sets of carriers (i.e., the *manifestation*) that *items reflect* through *represented characteristics*, because it is the *text* that plays the most important linking role in Figure 5.13.

What we do not appear to need metadata about are *manifestations* (i.e., sets of objects in the role of carrying signs that convey content). This is only natural because the ability to arbitrarily group and regroup things into sets by reusing existing features and dimensions that have already been defined comes for free when we employ database technology. Indeed, the whole purpose of query languages like SQL and SPARQL are so that we do not need to define entities like *manifestation* to accomplish this task. This notion is borne out when we consider that the set of *manifestations* is a subset of the power set of *items* in the WEMI2 model (as is the set of *characteristics*); or in other words, we can dynamically derive both any arbitrary set of characteristics and any arbitrary set of manifestations just by employing queries across the power set of items.

We know from our analysis in Section 5.3 that the WEMI2 account of *manifestations* cannot be successfully employed as is to model *bibliographic aggregates*. However, verb phrases like *appears-in* and the complex linking chain of *representing characteristics reflected-by items*, taken along with the grouping property indicated by the *set-member-of* relationship, indicate that there are likely two additional entities being conflated into the *manifestation* entity. One of these entities is possibly metadata records in which *expressions* appear and which record descriptive information *representing characteristics reflected-by* the *items* being described. The other entity being modeled are sets of *items* themselves which have possess all of the properties of being a set and none of the properties that the individual elements of a set might possess.

An additional odd notion presents itself. The NULL set has particular characteristics which can be reflected by a particular *item* and *represented* by a corresponding *manifestation*. As such is the case, it would seem that the NULL set is itself a bibliographic entity. The implication here is that *works* with no content exist. This is perhaps not so perplexing as we do have signs that are expressly intended to convey this concept. Whether or not this is actually a desirable outcome of a conceptual modeling effort for bibliographic entities remains to be seen.

Whatever the case may be, WEMI2 *manifestations* being sets presents a complicating factor for search results because most IR systems will not be articulate enough to tell us when a *manifestation* is taking on the role of being a *bibliographic aggregate*, versus being a set of WEMI2 *items*, or being a metadata record that describes both WEMI2 *expressions* and WEMI2 *items*. As we have already alluded to, one potential solution would be to ignore WEMI2's *manifestation* altogether, as everything it attempts to achieve appears to be achievable through other portions of either the

WEMI2 model or the LRM standard. As we will see in the next section, LRM even has alternatives to its own aggregate model.

5.7. LRM's Other Aggregate Model

Silently and without fanfare, the pages of the LRM document contain a second model for *bibliographic aggregates*. LRM defines several *part-of* relationships. One of the two most pertinent to our closing discussion here is a *part-of* relationship whose domain and range are WEMI2 *works*. Riva, Le Bœf, and Žumer define it in the following way:

“This [part-of] is the relationship between two *works*, where the content of one is a component of the other” – Riva, Le Bœf, and Žumer 2017, p 72.

They go on to say through the relationship's scope note that:

“This applies when the component-to-whole relationship is an inherent aspect of the *works* and holds for all the *expressions* and *manifestations* of the larger *work* and of its component *works*, whether the *expression* or *manifestation* comprises the full larger *work* or just one or more (but not all) of the component *works*. Examples include movements of concertos, poems within poetry cycles, multipart novels, triptychs.” – p 72.

Finally, they give us as examples, “*A wizard of Earthsea is part of the Earthsea trilogy* by Ursula K. Le Guin,” and “*Richard Wagner's Der Ring des Nibelungen has part Richard Wagner's Götterdämmerung*,” p 72.

These seem to be the very same kinds of examples we were looking at in Chapter 4. Noticeably, something we would normally call a component of a *work*, like a fictional person, does not seem to be within the scope of either the examples or the relationship. It would seem then that Winston, Chaffin, and Herrmann's *component-of* relationship (1987) is not actually intended to be invoked by the definition that Riva, Le Boëf, and Žumer give. Instead, we are again looking a highly-constrained version of their *member-of* relation. In fact, this *part-of* relation, which we will hereafter refer to as *[work]-part-of* bears such a resemblance to FRBR₀₀'s *[work]-member-of* relation, that they are likely one and the same, and we should treat them as such. This is not as surprising as it may seem. The author-editors of LRM are actually the same group of people who produced FRBR₀₀.

Similarly, LRM gives us a WEMI2 *expression-specific part-of* relation. Riva, Le Boëf, and Žumer define thusly:

“This is a relationship between two *expressions* where one is a component of the other” – Riva, Le Boëf, and Žumer p 74.

They give us the following examples:

“The music notation of Franz Schubert's *Ave Maria* Op. 52, No. 6 is *part of* the music notation of Franz Schubert's *Sieben Gesänge aus Walter Scott's Fräulein vom See* Op. 52” – p 74.

“The audio recording of Dante Alighieri’s *La divina commedia* read by Enrico de Negri *has part* the audio record of Dante Alighieri’s *La divina commedia, Inferno* read by Enrico de Negri” – p 74.

These examples of *expressions* are actually clearer vis-à-vis *what is an expression*, but nonetheless appear to be the same as the previous examples we have examined in Chapter 4. Again, what we would take to be an actual component, e.g., a particular character or utterance, does not seem to be what is intended to be expressed by this *part-of* relationship, which we will hereafter refer to as *[expression]-part-of*. Instead, once again, Winston, Chaffin, and Herrmann’s *member-of* relationship would seem to be indicated. Unsurprisingly, this *[expression]-part-of* relationship seems to directly correspond to FRBRoo’s *[expression]-component-of* relation. Unfortunately, LRM does not define a corresponding *part-of* relationship for WEMI2 *items* (recall that we are neglecting the existence of WEMI2 *manifestations* for the time being, that there is a *[manifestation]-part-of* relationship, and that it does not help resolve the issue of *bibliographic aggregates* that we noted in Section 5.4).

5.8. Preliminary Conclusions about the Content-Artifact Aggregate Model

It is clear from an illustration of the actual WEMI2 model (Figure 5.11 above) that *manifestations* have at least three, and perhaps as many as four, semantically distinct roles in the LRM conceptual model. This is clear evidence of semantic overloading. Computer systems, especially the RDF-based, semantically-sensitive ones emerging today, require extremely clear semantics in order to operate at peak performance. With regard to RDF- or Linked-Data-based IR systems, that means metadata must be as clear, precise, and accurate as possible. As Jett et al. (2016b) point out with their exploration of mapping existing special collections metadata to the RDF-compliant

Schema.org vocabulary, in Linked-Data computing environments metadata records must be as accurate as possible with regard to what entity each assertion applies to. Metadata records that conflate things like play productions, play performances, and play scripts require significant reform to bring adequate distinctions out.

An entity like WEMI2’s *manifestation* simply tries to be too many things for it to have full utility in RDF- or Linked-Data-based IR systems. At implementation time, each of its distinct roles will need to be carefully teased apart and extensions supporting them developed for the implementation to be successful. This will doubtless be an expensive and time-consuming task.

Ultimately, as the preceding sections indicate, the *content-artifact aggregate model* is not easily separated from the WEMI2 model itself. The fact that the *embodies* relationship is defined in LRM as fulfilling a “linking” role, but that the same document defines *bibliographic aggregates* in a way that suggests that the same relationship plays an *aggregating* role in some specific situations, makes it difficult to reconcile both definitions qua formal definitions. One of them, likely definition D.17, needs to be weakened to mere implication. We can make this adjustment formally through the following axiom.

$$\mathbf{A.52 \text{ (Bibliographic Aggregates Semi-Defined): } \forall x \left(\left(\exists y \exists z \left(\left(\text{embodies}(x, y) \wedge \text{Expression}(y) \right) \wedge \left(\text{embodies}(x, z) \wedge \text{Expression}(z) \right) \wedge \neg(x = y = z) \right) \right) \rightarrow \text{bAggregate}(x) \right)}$$

Axiom A.52 has an interesting implication too. It implies is that not all *bibliographic aggregates* are accounted for by the “*embodies* multiple WEMI2 *expressions*” case. This implication may

satisfy some of the curious examples we showcased above that would have required many-to-many cardinality from the original WEMI model's *realizes* and *exemplifies* relations. The model suggested by the natural language account (and illustrated in Figure 5.11) might suffice for this purpose; however, at least some *bibliographic aggregates* will be *works* that aggregate other *works*, *expressions* that aggregate other *expressions* and *items* which aggregate other *items*. Without the definition of additional mereological relationships to satisfy these use cases, it seems doubtful that a faithful implementation of LRM will result in an IR system that can adequately provide metadata for all of the bibliographic aggregates that users are likely to desire.

Like all standards, LRM has its strengths and weaknesses. Among the former is the much clearer picture its text provides of the bibliographic universe. Unlike the original FRBR document, Riva, Le Bœf, and Žumer have taken many pains to provide adequate detail for developers and systems designers to more easily implement the standard in IR system settings. LRM's only true failing is that its account of *manifestations* seems disconnected from its account of bibliographic entities. Thereby, we are uncertain exactly what role they are meant to play in any faithful implementation of LRM.

If they are meant to specifically fulfill the role of providing a model for *bibliographic aggregates*, then they seem to fall short, as we demonstrated with the examples of *expressions conveying* more than one *work* and *items carrying* more than one *expression*. These happen because the model is not fully committed to representing *bibliographic aggregates* as the bibliographic entities they are, with their own unique content, signs, and carriers.

If *manifestations* are meant to be sets of *items*, then they become truly redundant parts of the model at implementation time because all implementations are already going to have query languages that can better fulfill users' needs for arbitrarily grouping bibliographic entities or specific facets of bibliographic entities (e.g., by content, signs conveying content, or objects carrying signs conveying content) on the fly without needlessly wasting storage space or processing cycles on a conceptual entity that does those things.

If *manifestations* are intended to represent existing metadata records, we should ask what role the rest of the model plays with respect to them. We see that *expressions* and *items* have role-specific links to *manifestations* through relationships like *reflects* and *appears-in*. But if the metadata record is to also faithfully represent a bibliographic entity's content, then the WEMI2 model would seem to be missing a relationship that explicitly draws the link from *works* to *manifestations* and, it tells us what role is being played by one with respect to the other.

Of course, providing a conceptual placeholder for metadata records in the model will also be redundant at implementation time, because we should have document models that call upon existing queries in whatever query language our IR system is using to produce these records on the fly when users need them.

Finally, if *manifestations* are intended to represent publication plans, we should ask ourselves what if, like *bibliographic aggregates*, publication plans are not actually bibliographic entities in their own right? If we write a publication plan on paper, then it becomes clear that there is an object which carries signs that convey the content of that plan. We might object that it seems to be the

case that the plan explicitly refers to some *items*, and so we might be tempted to think we need some special relationship to express this connection. However, LRM already accommodates this connection through the *has-subject* relationship defined on p 69. After all, it is not the signs or the object which are about some *items*, it is the content of the plan that enjoys topicality. And so, we do not need WEMI2 *manifestations* to arrive at a perfectly adequate model for publication plans.

At this point, we have arrived at an answer for research question number 2, how does the *content-artifact aggregate model* used by LRM differ from the *mereological aggregate models* examined in Chapter 4? For one thing, we can see that it possesses high context-sensitivity. Specifically, it is very sensitive to the manner in which WEMI's (or WEMI2's in this case) *manifestation* entity is defined. As we saw, defining *manifestation* as a kind of set provides a formidable barrier to successfully modeling bibliographic aggregates, since the characteristics of the sets and their aggregates are not the same. Further, in LRM's case, the *manifestation* entity itself is semantically overloaded, making it difficult to tease apart when any particular *manifestation* is playing the role of being a *bibliographic aggregate* or playing one of the other roles that *manifestations* play in the LRM standard.

5.9. Chapter Summary

As we have seen throughout this chapter, the WEMI model employed by the LRM conceptual standard (i.e., WEMI2) is much more clearly defined than the WEMI model first issued in the original FRBR. Overall, this is a change for the better. However, as we also showed, the WEMI2 definition for *manifestations* directly invokes set-theoretics. This causes problems for LRM's aggregate model, as it is difficult to reconcile the characteristics of the set-elements of a journal issue *manifestation* (e.g., the characteristics of the individual issues) with those of the articles

embodied by them. This exemplifies an immediate semantic overloading of the *manifestation* entity concept, as it is no longer clear when a *manifestation embodies* an article (i.e., when are the elements articles) or when it *embodies* the articles in an issue (i.e., when are the elements journal issues). Employment of language suggesting that *manifestations* also represent entities like “production plans” or play a role similar to metadata records adds additional semantic complexity. As we noted above, the kind of aggregate model employed by LRM—a *content-artifact aggregate model*—is extremely sensitive to the content-artifact model employed with it (e.g., WEMI2 vs WEMI). In the next chapter, we compare the results of the analyses in Chapters 4 and 5, summarize our findings, and suggest avenues for future research in the conceptual spaces we have been working in.

6. Findings & Next Steps

6.1. Reprise of Research Goals

Throughout the course of this dissertation, we have employed conceptual analysis as a means to reflect upon how *bibliographic aggregates* are treated in four high-level conceptual standards (DC-CAP, FRBR, FRBR_{oo}, and LRM). In many cases, we did this by proposing an informal formalization of one or more natural language sentences in first-order predicate logic, and then examining given examples to see if the formalizations were able to accommodate all of the given examples of *bibliographic aggregates*. From this analysis, we have identified several broad implications for the general description of *bibliographic aggregates* through high-level metadata.

The ultimate goal of our analysis was to develop answers to the following research questions:

1. How comparable are the *mereological aggregate models* employed by conceptual standards like DC-CAP, FRBR, and FRBR_{oo}?
 - a. The texts of these standards often employ disparate part/whole terms like “container,” “component,” “member,” etc.; are we to take these terms at face value or are they being employed synonymously to indicate one specific part/whole relationship?
 - b. In the event that they are being used to indicate one specific part/whole relationship, which part/whole relationship is it?
2. How does the *content-artifact aggregate model* employed by LRM differ from the *mereological aggregate models* employed by our other three cases?

By answering these questions, we hope to better inform choices regarding how these standards should be employed to capture general metadata describing *bibliographic aggregates*. In LRM’s

case, we can show that the application of the *content-aggregate model* fails due to the model's highly context-sensitive nature. In this chapter, we summarize our findings, consider repairs for the LRM conceptual standard, and contemplate future avenues to expand upon the research that has been carried out here.

6.2. Summary of Findings

6.2.1. Finding 1 – Comparability of *Mereological Aggregate Models*

In Chapter 4, we compared five part/whole relationships (*is-gathered-into*, *[bibliographic]-part-of*, *[expression]-component-of*, *[work]-member-of*, and *incorporates*) from three different bibliographic standards (IFLA 1998, DCCDTG 2007, Bekiari et al. 2016) and found that they are all closely related. Based upon the sameness of the example aggregates that they are intended to be used with (e.g., anthologies, collections, journal issues, series, etc.), we determined that they all shared metaproperties and that their primary differences amounted to issues of scope.

More specifically, the part/whole relationships we examined in our cases evinced increasingly narrower domains and ranges, to the point that the domains and ranges employed can be described as sub-domains and sub-ranges of the part/whole relationships employed by other standards. For instance, the domain and range of FRBR's *[bibliographic]-part-of* was narrower than the DC-CAP's *is-gathered-into* relation. In turn, FRBRoo's *[expression]-component-of*, *[work]-member-of*, and *incorporates* were all narrower than FRBR's *[bibliographic]-part-of*. Based on the examples employed by the standards and other analyses of them, there seem to be no other differences among these part/whole relations.

Comparing these part/whole relationships to Winston, Chaffin, and Herrmann's collective-member relationship (1987), i.e., the *member-of* relationship, we also found that they shared a common group of metaproperties with that relationship. And so, we find DC-CAP's *is-gathered-into* relationship is exactly the same as Winston, Chaffin, and Herrmann's *member-of* relation, and we also find that the remaining part/whole relationships used by FRBR and FRBR₀₀ are all sub-properties of Winston, Chaffin, and Herrmann's *member-of* relation.

This has the very practical benefit of allowing practitioners to adopt any of the three standards (DC-CAP, FRBR, or FRBR₀₀) and expect to be able to accurately model *bibliographic aggregates* as general bibliographic entities within the established limitations of each of the standards. What this means is that in general, all three standards treat *bibliographic aggregates* exactly as general bibliographic entities. Therefore, each one of these is, in fact, missing some of the particular metadata facets that are often found to be desirable when one wants to differentiate between an anthology and an archival collection.

Fortunately, this is not an uncommon problem at implementation time for standards, and the existing best practices would dictate that the standards be extended with anything missing, so that finer-grained distinctions among different kinds of *bibliographic aggregates* can be made. Unfortunately, extensions to FRBR's WEMI model (also used in FRBR₀₀) have proven difficult to construct as we have shown with examples like McDonough et al. (2010) and Jett et al. (2015). Through the former's video game example, we can see that WEMI begins to break down when faced with the need to accommodate multiple language dimensions (e.g., computational language versus human language). In turn, Jett et al.'s solution for providing adequate metadata for video

games focuses on decomposing the fine-grained distinctions between what they call [video game] editions which are platform (i.e., computer-language) specific and local releases (or localizations) which are human-language specific entities. In more FRBR-specific language, Jett et al.'s solution calls for the decomposition of WEMI's *expression* entity into two entities—edition and localization—for video games.

6.2.2. Finding 2 – Gaps in FRBR₀₀'s Aggregate Model

With regard to the overall goals of FRBR's WEMI model, its goal has always been to illustrate the differences in the continuum from the content of a bibliographic entity (e.g., topicality, etc.) to its physical characteristics (e.g., how much space it takes up, etc.). The fact that FRBR₀₀ apparently lacks part/whole relationships having to do with two particular entities on the content-artifact continuum, namely WEMI's *manifestations* and *items*, is a limitation to what kinds of things can be said of the physical characteristics of *bibliographic aggregates*. We would expect that come implementation time, FRBR₀₀ will need supplementation with another existing model that describes bibliographic entities or extension with additional entities and relationships that ensure that metadata which would otherwise be missing can be recorded for *bibliographic aggregates*. Such supplementation is commonplace and to be expected at implementation time, because typical implementations possess more stringent constraints due to local contextual requirements than general models are required to address.

6.2.3. Finding 3 – Failings in the *Content-Artifact Aggregate Model*

Through counter-examples we were able to show that the *content-artifact aggregate model* set forth by the FRBR Working Group on Aggregates (2011) and codified into IFLA's LRM standard (Riva et al, 2017) creates several problems. Chief among these is that it forms barriers to users

trying to accomplish FRBR's user tasks (finding, identifying, selecting, and obtaining) when a particular item is both mass-produced and part of one or more *bibliographic aggregates*.

We showed that this occurs because WEMI2's *manifestation* entity is specifically defined as a kind of set. Since sets do not share characteristics with their elements or their sub-sets, it becomes difficult to reconcile the characteristics of an entity like a journal issue with the characteristics of the articles that are aggregated into it. From this analysis, it would seem that *bibliographic aggregates* are not successfully modeled with the LRM standard, and that the particulars of the *content-artifact aggregate model*, as defined through WEMI2, make it difficult to supplement with other vocabularies or extend with additional entities and relationships such that adequate metadata representing *bibliographic aggregates* can be recorded. However, we may be able to ignore WEMI2's *manifestation* entity altogether as a means to repair this. Further discussion of this solution is the focus of the next section.

6.2.4. Finding 4 – The WEMI2 *manifestation* Entity is Semantically Overloaded

As we showed in Chapter 4, *manifestations* play at least four distinct roles in the WEMI2 model.

This semantic overloading will complicate any implementation of the LRM standard. It may even make faithful implementations of the standard impossible when the implementation environments have strict requirements for clarity and precision as is the case in Semantic Web, Linked Data, and RDF-based computing environments.

As we have already mentioned in Finding 2 above, it is not unusual for standards to be supplemented or extended (i.e., customized) at implementation time. Unfortunately, the semantic overloading here is so severe that proper supplementation or extension of the standard is likely

difficult. However, as we mentioned in Finding 3, one potential alternative is to remove WEMI2's *manifestation* entity altogether. This is possible because several of the things it does, such as grouping WEMI2 *items* together according to special criteria, i.e., shared characteristics, are easily accomplished through traditional IR system infrastructures (e.g., relational database queries) without any need to model sets as full-fledged entities in our conceptual space. We provide more detail on this approach in the next section.

6.3. Excising WEMI2 *Manifestation* from the LRM Standard

In Chapter 5, we mentioned that portions of LRM's WEMI2 model had the potential to be a superior model for the deduplication of bibliographic entities. The first step for realizing this model's potential would be eliminating the *manifestation* entity and all relationships and attributes that rely on it from LRM and the WEMI2 model entirely.

We think this is an achievable action, as all of the roles that a WEMI2 *manifestation* plays are redundant with those played by other entities and relationships in LRM or WEMI2, or which are part and parcel of actual implementation infrastructure.

In particular, the ability to group entities into sets according to particular shared characteristics is better achieved by relying on the query languages that come with our data storage and IR systems. There is no need to model these sets as first-class members of our domain space. The only time we might want to model them is when we want to model the behaviors of those queries. Modeling such behaviors is well outside the scope (and norm) of bibliographic standards like LRM and conceptual models like WEMI2.

We also do not need WEMI2 *manifestations* to arrive at a perfectly adequate model for *bibliographic aggregates*. WEMI2’s existing *[work]-part-of* and *[expression]-part-of* relationships already accomplish two-thirds of the task. All that is lacking is a comparable *[item]-part-of* relationship. One easily accomplished approach to extending LRM with an *[item]-part-of* relationship is to simply rehabilitate LRM’s *[manifestation]-part-of* relationship by replacing instances of the word “manifestation” with the word “item,” and then eliminating all language referring to publication plans. After all, what we really want to account for here are differences among carriers that carry the same signs (conveying the same content). We do not need to refer to publication plans as some kind of intermediary between *expressions* and *items*. Like *bibliographic aggregates*, plans are themselves first-class bibliographic entities.




Figure 6.1: Further Simplified WEMI2 (or WEI) Model

At this point the question becomes, do we need *manifestations* to model anything at all? If we have no *manifestations*, then we will have arrived at much leaner WEMI2 model (illustrated in Figure 6.1 above). We can see that it appears very similar to a content-artifact model proposed by Richard Smiraglia some 18 years ago (Figure 6.2 below). While there is no comparable entity to “bibliographic entity” in Figure 6.1, we take it to be the case that WEMI2 is intended by Riva, Le Bœf, and Žumer to illustrate a particular view of bibliographic entities in general.⁷⁰

⁷⁰ Indeed, the entire thesis presented here hinges on the interpretation that WEMI and WEMI2 are describing aspects of bibliographic entities.

We note that Smiraglia’s “text” entity is playing the same intermediary role as WEMI2’s *expression* entity.⁷¹ Indeed, the definition for *expression*—“signs conveying content”—suggests that “text” is within its purview. However, we must contend with an objection one can have with these accounts: How does one differentiate between different *items* carrying the same *expression* and different *items*? Or more simply, how do we account for differences in format?

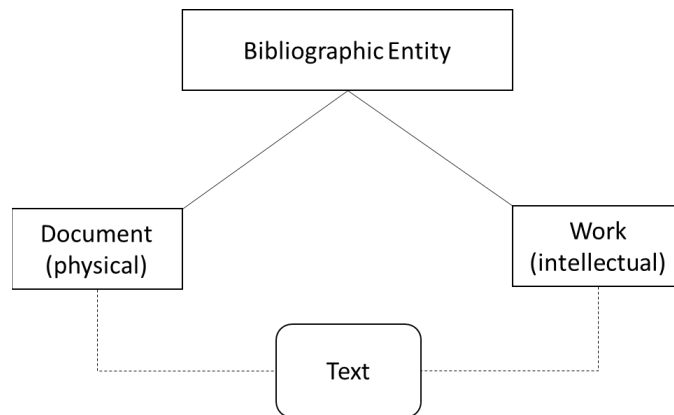


Figure 6.2: Smiraglia’s Bibliographic Entity Model⁷²

Originally, the concept of *manifestation* was likely intended to help with this issue. But as we have seen from LRM, defining it so that it can accomplish this task is a tricky business. And where the item of a particular bibliographic entity is a singleton (e.g., a manuscript, a painting, a sculpture, etc.), we must confess the whole business is quite unnecessary. An intervening entity here only gets in the way in these cases.

How then should we proceed?

⁷¹ It may be important to note that we are uncertain to what extent that Smiraglia is committed to the model he proposes. As early as 2002 he moves from defining *works* as “intellectual content” to sets of metadata records.

⁷² Reproduced from Figure 1.1 in Smiraglia 2001, p 4.

One way is to make a conceptual interpretation of what we mean by “format.” One interpretation that might be justified is considering “format” to be a kind of “sign.” We know (from cryptography, among other sources) that it is possible for signs to have the role of carrying other signs (indeed, blockchain is a highly pertinent example). We also know that WEMI2 defines *items* as “[physical] objects carrying signs conveying content.” However, a further interpretation, that “objects” might in some circumstances mean “signs,” does not necessarily invalidate the given definition if we relax the constraints given by LRM’s scope notes concerning WEMI2 *items*. Finally, if it is the case that the *carries* relationship is a *transitive* one, then it seems possible that we might have a situation where an artifact is carrying signs which are carrying signs which are carrying signs, and so on, and so forth, until we reach some signs conveying content. We illustrate what this might look like below in Figure 6.3 (below).

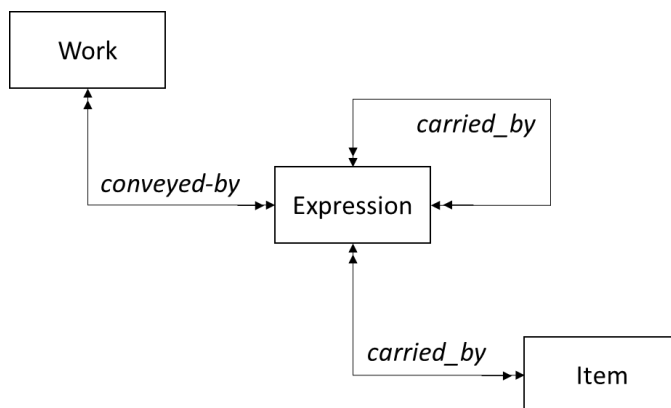


Figure 6.3: Modified WEMI2 (or WEI) Model

The primary objection to the model in Figure 6.3 is that format-level differences do not seem like an *expression*-level differences. Formats do not intuitively seem to be signs.

However, it seems to us that when an author, editor, curator, or other content creator determines to *express* their content in a particular format, then they are also ontologically committed to the

limitations dictated by that format. Many of these formats evince features, often reflecting best or common (and sometimes normative) practices of the time and place in which an expression is first created and which place certain limitations on the kinds and amounts of content which can be conveyed. Thus, format acts as a constraint upon the entity of *work* in our model, by limiting how much of a *work* can be *conveyed* by any given granule of an *expression*. Format is also suspiciously grammar-like in some instances. It may be a non-normative grammar, but it is an accepted grammar which conveys content nonetheless. It seems counter-intuitive but, considering the constraints that format places upon the conveying of content, it also seems quite rational to accept that format is, in fact, a kind of sign.

This is not an unusual position. The original WEMI model used format as an attribute of *work* (defined there as a “creation” rather than content). If we accept that format is a kind of sign, then differences in formats are *expression*-level differences where the relationship between the expressions is that of one carrying the other exactly as Figure 6.3. illustrates. We can then leverage this modified WEMI2 model (really *work-expression-item* [WEI] at this point) to help us disambiguate all of the different copies of *works* that have resulted from mass-production and digitization.

6.4. The Paratext Issue Revisited

Whether we excise the *manifestation* entity from the WEMI2 model or not, there still remains one unresolved issue for *bibliographic aggregates*—paratext. Paratext is an extremely difficult problem for cataloging. That publishers include additional content in all manner of diverse forms as Coyle points out (2016b) is an incontrovertible fact. The question, again, is how to proceed?

Should we record metadata about paratextual features? If so, how does this impact the general conceptual model for *bibliographic aggregates*?

To some extent, the answers to these questions have serious implications for the entire cataloging endeavor. If we answer, “yes, paratextual features need to be recorded, especially at the *manifestation* level,” then the result for *bibliographic aggregates* will be a combinatorial explosion in metadata for and about them.

Let us take the *Shannara* series (Figure 6.4 below) as an example once again. For the *works* that compose the *work* that is the series, there may be several *expressions* of each *member-work* and so there a number of distinct *expressions* of the series as a whole. We will use a more specific example to illustrate this, by considering the *Shannara* series as it was in the early 80’s when only the first three novels in the series had been published. We will say there are three editions of the first book, *The Sword of Shannara*—one with illustrations by the Hildebrandt brothers, one with alternate illustrations of the same scenes but which are also by the Hildebrandt brothers, and one where the text comprises a third of an omnibus edition called *The Sword of Shannara Trilogy*. We will say that there are only two editions of the other two books, *The Elfstones of Shannara* and *The Wishsong of Shannara*—in one edition they are distinct novels packaged with cover art by the Hildebrandt brothers, but sans any internal illustrations, and another edition where each comprises a third of an omnibus called *The Sword of Shannara Trilogy*.

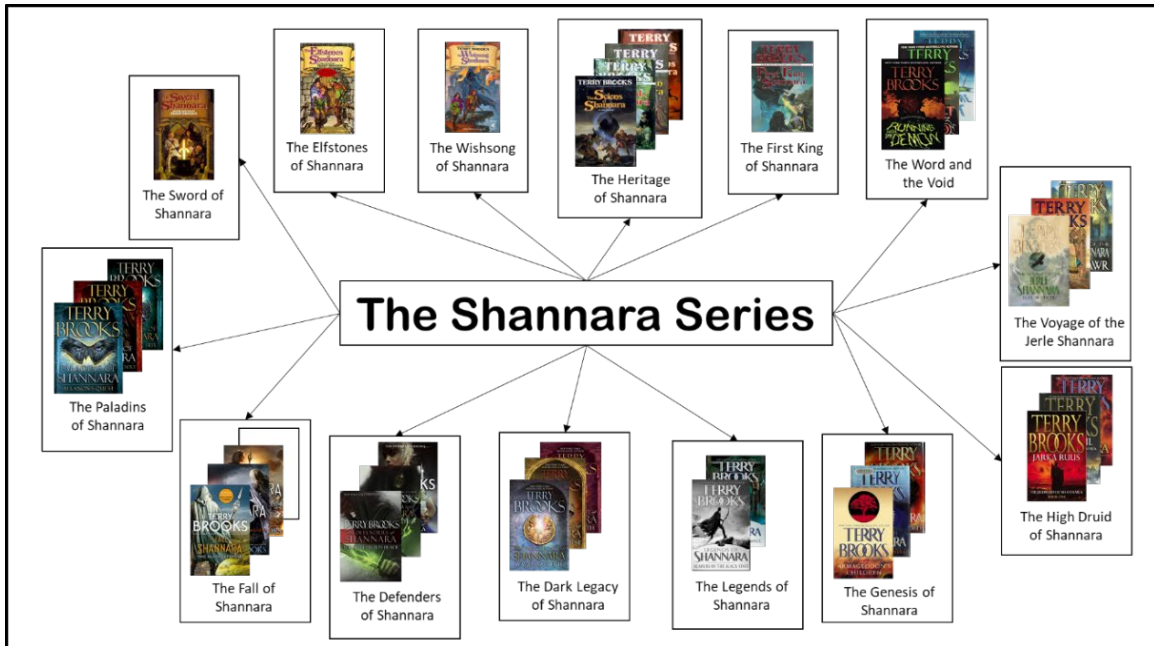


Figure 6.4: *The Shannara Series* by Terry Brooks

It is clear to us at this point in time that the *expression* of *The Sword of Shannara Trilogy* [an omnibus] is one *expression* of the work, *The Shannara Series* (since the other works have yet to be published). However, it is also the case that any particular combination of the two *expressions* of *The Sword of Shannara*, in combination with the expressions of *The Elfstone of Shannara* and *The Wishsong of Shannara*, also compose *expressions* of *The Shannara Series*. At this point we have three distinct *expressions* for one *work*.⁷³

This combinatorial expands when we begin to consider the combinatorials of *manifestations* for our *expressions* of *The Shannara Series*. Say for instance that there is a hardcover edition, a special collector’s hardcover edition, a mass-market “book club” hardcover edition, a trade paperback edition, and two mass-market paperback editions (featuring different cover art) for each of the

⁷³ It is possible that we might have two additional expressions if we allow parts of certain expressions, e.g., *The Sword of Shannara Trilogy*, to fulfill the role of whole monographs in the expression of *The Shannara Series* (circa the early 1980s).

individual *expressions*, including *The Sword of Shannara Trilogy*. Now we have a situation where our one *work*—*The Shannara Series*—is linked to 55 different *manifestations*. But this is to be expected, and this tree-like grouping structure of *works*, *expressions*, *manifestations*, and *items* is the whole point of the original WEMI model.

However, things quickly go downhill if we apply the *content-artifact aggregate model* used in LRM, as the *manifestations* are the series, and so, the only metadata we have is of these 55 different versions of what appears to be the same series (at least from the perspective of *work*, i.e., the content). Now imagine this is true for the entirety of the series illustrated in Figure 6.4. One imagines that where *bibliographic aggregates* like series are defined as *manifestations* then there will be hundreds upon hundreds of *aggregate manifestations* (due to the combinatorials of *aggregate manifestations* comprising individual *manifestations*) and that for each, metadata distinct to its particular combination of individual *manifestations*, i.e., distinct to that *aggregate manifestation* will be recorded as though each was in fact a distinct series, a distinct *bibliographic aggregate*, even though the content remains the same across all of them. This seems to be contrary to the kind of disambiguating agenda set out by bibliographic theorists like Verona, Wilson, Lubetzky, and the like. What we really want is an approach like those set out by those scholars that lets us say something like, all of the different combinations, all the different bibliographic units, are in fact all the same literary unit.

Unfortunately, things can once again be made complex, if we consider whether or not the content communicated through paratextual features is content that affects the context of a *work's* “primary” content. If this is the case, then like the case we make for *bibliographic aggregates*

themselves, any combination of particular paratext with a primary text is actually a distinct *work* in its own right.

Now a new question presents itself—how should we represent this kind of information? There would seem to be two approaches, we might model them as *bibliographic aggregates* or we might model them as *derivative works*. This is not a question the research carried out in this dissertation was designed to answer, or even to contribute an answer towards. However, some pragmatic issues present themselves. If we model the contextual relationship between paratextual content and primary textual content as a *bibliographic aggregate*, then it seems as though all or almost all bibliographic entities are *bibliographic aggregates*. Examples of dependent works in our standards in particular suggest that this might be the intention of standards writers. Such a solution is going to cause the same kinds of combinatorial problems that defining *bibliographic aggregates* as *manifestations* does; specifically, we can expect an explosion in the numbers of *works* which on the surface all seem alike. Such an explosion directly detracts from our ability to identify duplicate content.

An alternative solution might be called for. One possible alternative is that when paratextual content alters the primary text content to such an extent that we feel that recording metadata about the paratextual content (and its paratext) is vitally important to the end user (as will likely be the case for scholarly editions and similar specialized products). Then, we employ an entirely different part of our bibliographic conceptual model—the linking device of the *derivative* relationship.

We can use this as a simplifying tactic for our metadata and the links among particular assertions that it makes. In the case of a scholarly edition, say for example of Brooks' *The Sword of Shannara*, modeling the scholarly edition as a *derivative work* has the benefit of subtracting it from the context of the overarching *work*—Brooks' *The Shannara Series*—that it is a part of. This is because the scholar may remark on the *work*, *The Sword of Shannara*, directly, but can only remark on the series indirectly. One imagines if a scholar were to desire to remark on the series directly itself, then they would need to produce a scholarly edition of the series in its entirety.

In the end, both of these approaches are only tactics that we can propose here, but which we cannot definitively state as being better or worse with respect to one another, or with any other approach to the handling of paratextual features, without additional research. Ultimately, we leave this task to future meditations regarding paratextual features and their impact on metadata practices.

6.5. Avenues for Future Research

As we have seen, many of LRM's and WEMI2's shortcomings can be overcome when we remove the *manifestation* entity altogether. However, more work needs to be done to ascertain the full implications of such a radical change to both LRM and the WEMI2 model. A potential avenue for further research in this respect would be the comparison of the modified WEMI2 model proposed in Figure 6.3 with other similar conceptual models.

One such model (the Basic Representation Model) was proposed in Wickett et al. (2012) in their analysis of data. The figure below, reproduced from Wickett et al. (p 4)⁷⁴ showcases how similar the modified model we proposed with Figure 6.3 is to their model.

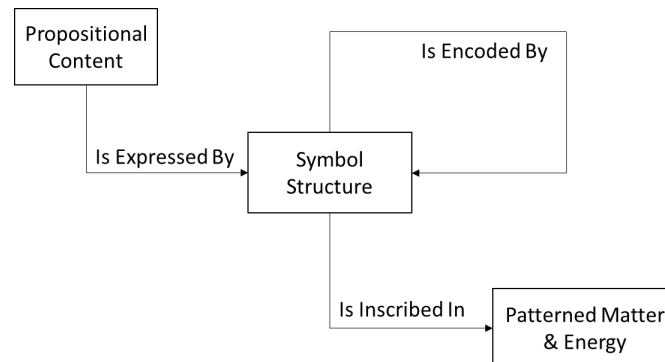


Figure 6.5: Basic Representation Model (BRM)

Here we have three entities that are content, signs, and objects. There are several subtle differences. The content-level is limited to just propositional content, but this is appropriate for the data objects that Wickett et al. are examining. Like WEMI2 *items*, the object-level is clearly intended to be limited to physical artifacts. And finally, the relationships seem somewhat clearer; there is, at least, a distinction made between when the carriers of symbol structures (i.e., signs) are other symbol structures, and when the carriers are physical objects onto which the symbol structures are inscribed.

BRM is only a part of the story that Wickett et al. tell us about data. Equally important is an interpretive framework called the Systematic Assertion Model (SAM [Dubin, Wickett, and Sacchi 2011; Wickett et al. 2012]). SAM is the part of their approach to describing data that does the heavy lifting of capturing information, such as metadata, provenance, etc., that are important for

⁷⁴ This is Wickett et al.'s Figure 1.

end users to successfully understand what the data is, why it was made, and how they might go about assessing it for their own uses.

This picture of data that Dubin, Wickett, and Sacchi give us is not so different to how the rest of Riva, Le Boëf, and Žumer's LRM standard relates to WEMI2. One future avenue of research then, is to more closely examine the parallels between the BRM and SAM models and the WEMI2 and LRM models.

Additionally, while the thesis here closely examines *bibliographic aggregates* in the context of four bibliographic standards, those four standards are hardly all of the standards that exist. In particular, the Library of Congress's new Bibframe 2.0 standard seems poised to rival IFLA's LRM as the quintessential high-level library metadata standard. And so, a natural expansion of the research carried out here is an examination of Bibframe 2.0's treatment of *bibliographic aggregates* as full-fledged bibliographic entities. Similarly, the new cataloging bible, RDA,⁷⁵ is being affected by our evolving understanding of standards like FRBR, LRM, and Bibframe 2.0. So, it too, is another natural expansion avenue for the research that has been carried out here. Unfortunately, RDA is a subscription service, which complicates its study by the need for greater financing.

A final avenue for expanded research lies in more closely examining the various kinds of *bibliographic aggregates* themselves. The standards here all take high-level, generalized outlooks on *bibliographic aggregates* in order to provide overarching advice for the creation of minimal

⁷⁵ Poised to be the successor to the ALA's AACR2.

metadata for them. While some of the standards succeed at this more than others, the fact remains that the examples of *bibliographic aggregates* are a disparate grouping of bibliographic entities themselves. An archival collection is a very different kind of bibliographic entity than an anthology or a monograph series, even though all three of them are *bibliographic aggregates*.

There is a need to more closely examine each kind of *bibliographic aggregate* and its user-base to better understand where gaps in metadata representing qualities unique to that particular kind of *bibliographic aggregate* (e.g., anthology, collection, series, etc.) exist. In turn, where such gaps are found to exist, there is often a clear need to propose extensions to existing standards to fill those gaps and thereby better help users accomplish the essential tasks of finding, identifying, selecting, and obtaining.

6.6. Conclusion

Through our analysis here we have shown that:

- All of the bibliographic standards all using *mereological aggregate models* all use the same aggregate model or portions of the same model.
- In the case of the FRBR₀₀ standard, gaps in *mereological aggregate model* indicate that the standard requires supplementation or extension in order for *bibliographic aggregates* to receive equal treatment as bibliographic entities.
- The *content-artifact aggregate model* used in LRM lacks the ability to capture *work-level* and *item-level* qualities of *bibliographic aggregates* and is thereby unable to model *bibliographic aggregates* as full-fledged bibliographic entities in their own right.
- WEMI2's *manifestation* entity is semantically overloaded and performs a variety of disparate functions in the LRM standard. This implies that it will be difficult to implement

in RDF-based and Linked-Data computational environments which have much stricter requirements with regard to semantic overlap of entities and relationships than simpler computational environments.

Through these findings, we have better characterized how the *mereological aggregate models* used in the high-level conceptual standards DC-CAP, FRBR, and FRBR_{oo} are all, in essence, the same model of parts and wholes. The primary difference among the three standards' approaches to modeling *bibliographic aggregates* is one of ever-narrowing scope. In the case of DC-CAP the *is-gathered-into* relationship is general enough to be considered synonymous Winston, Chaffin, and Herrmann's *member-of* relationship. One implication of this is that the DC-CAP aggregate model is actually general enough to model general groupings such as groups of people, forests, fleets, and other arbitrary groupings beyond *bibliographic aggregates*. While this may seem outside of the intentions of the DC-CAP designers, it is in keeping with the DCMI vocabulary's extremely generalized semantics. In turn, with their narrower domains and ranges, it becomes more obvious that FRBR and FRBR_{oo} are both better fits for modeling *bibliographic aggregates*, specifically.

Conversely, we have seen that because the *content-artifact aggregate model* is closely coupled to the WEMI model, it is extremely sensitive to variations in the definitions for WEMI's entities and relationships. In the LRM case, we see that the employment of a set-theoretic definition for the *manifestation* entity has serious implications for how well the *content-artifact aggregate model* can articulate information regarding *bibliographic entities*. In this case, it fails on account of irreconcilable differences among the entities being modeled and the characteristics that describe them. In part, this is because the role that *manifestations* play at any given time is not clear in LRM

contexts, due to the significant semantic overloading of the *manifestation* concept. However, this problem has revealed opportunities to simplify both the WEMI conceptual model and how employment of *mereological* models for aggregates is beneficial by decoupling accounts of parts and wholes from accounts of content and artifacts.

References

- American Library Association (ALA). (2002). *Anglo-American Cataloging Rules, 2nd Edition, 2002 Revision*. Chicago: ALA.
- Antelman, K. (2004). Identifying the serial work as a bibliographic entity. *Library Resources & Technical Services* 48(4), pp 238-55.
- Avram, H. (1968). *The MARC Pilot Project: Final Report on a Project Sponsored by the Council on Library Resources, Inc.* Washington, D.C.: Library of Congress.
- Baca, M. & Clarke, S. (2007). FRBR and works of art, architecture, and material culture. In A. Taylor (Ed.), *Understanding FRBR* (pp. 103-110). Westport, CT: Libraries Unlimited.
- Barbakoff, A. (2017). Balancing connections and collections: Making space for users to interact with materials—and one another. *Library Journal* [online]. Retrieved from: <https://www.libraryjournal.com/?detailStory=balancing-connections-collections-library-design>
- Beaney, M. (2016). Analysis. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy (Spring 2016 Edition)*. Retrieved from: <http://plato.stanford.edu/entries/analysis/>
- Bekiari, C., Doerr, M., Le Bœuf, P. & Riva, P. (eds.) (2015). *Definition of FRBROO: A conceptual model for bibliographic information in object-oriented formalism (version 2.4)* [revised 2016]. The Hauge: IFLA.
- Berners-Lee, T., Hendler, J. and Lassila, O. (2001). The semantic web. *Scientific American* 284(5), pp 29-37.
- Bittner, T. & Donnelly, M. (2005). Computational ontologies of parthood, componenthood, and containment. In L. Kaelbling (ed.) *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence 2005 (IJCAI05)*, pp 382-7.

- Brack, E., Palmer, D. & Robinson, B. (2000). Collection level description – The RIDING and AGORA experience. *D-Lib Magazine* 6(9). Retrieved from:
<http://mirror.dlib.org/dlib/september00/brack/09brack.html>
- Buckland, M. (1991). Information as thing. *Journal of the American Society for Information Science* 42(5), pp 351-60.
- Buckland, M. (1997). What is a ‘document’? *Journal of the American Society for Information Science* 48(9), pp 804-9.
- Charles, V. ([written 2017]). The digital art exhibition example. In Jett, J., Charles, V., Humpal, N., Dubin, D., & Lee, J.-H. (unpublished manuscript) Bibliographic aggregates: A comparison of the FRBR and IFLA-LRM aggregate models.
- Chen, P. (1976). The entity-relationship model: Toward a unified view of data. *ACM Transactions on Database Systems* 1(1), pp 9-36.
- Cotnoir A. J. & Bacon, A. (2012). Non-wellfounded mereology. *The Review of Symbolic Logic* 5(2), pp 187-204.
- Cotnoir, A. J. (2018). Is weak supplementation analytic? *Synthese* (2018), pp 1-17.
- Coyle, K. (2016a). The evolving catalog: Cataloging tech from scrolls to computers. *American Libraries* [online]. Retrieved from:
<https://americanlibrariesmagazine.org/2016/01/04/cataloging-evolves/>
- Coyle, K. (2016b). *FRBR, Before and After: A Look at Our Bibliographic Models*. Chicago: American Library Association.
- Currall, J., Moss, M. & Stuart, S. (2004). What is a collection? *Archivaria* 58, pp 131-46.
- Cutter, C. A. (1891). *Expansive Classification: Part I: The First Six Classifications*. Boston: C. A. Cutter.

- Dewey, M. (1876). *Classification and Subject Index for Cataloging and Arranging the Books and Pamphlets of a Library*. Hartford, CN: Case, Lockwood, & Brainard Company.
- Dewey, M. (1885). *Decimal Classification and Relative Index for Arranging, Cataloging, and Indexing Public and Private Libraries for Pamphlets, Clippings, Notes, Scrap Books, Index Rerums, etc.: Second Edition, Revised and Greatly Enlarged*. Boston: Library Bureau.
- Donnelly, M. (2010). Parthood and multi-location. *Oxford Studies in Metaphysics* 5, pp 203-43.
- Dubin, D., Wickett, K. M. & Sacchi, S. (2011). Content, format, and interpretation. In *Proceedings of Balisage: The Markup Conference 2011. Balisage Series on Markup Technologies 7* (Montréal, Canada, 2-5 August 2011). DOI: 10.4242/BalisageVol7.Dubin01.
- Dubin, D., Senseney, M. & Jett, J. (2013). What it is vs. how we shall: Complementary agendas for data models and architectures. In *Proceedings of Balisage: The Markup Conference 2013. Balisage Series on Markup Technologies 10(2013)*. DOI: 10.4242/BalisageVol10.Dubin01
- Dublin Core Collection Description Task Group (DCCDTG). (2007). *Dublin Core Collections Application Profile*. Seoul: DCMI. Retrieved from: <http://dublincore.org/groups/collections/collection-application-profile/2007-03-09/>
- Duranti, L., Eastwood, T. & MacNeil, H. (2002). *Preservation of the Integrity of Electronic Records*. New York: Springer Publishing.
- Europeana (2016). *Definition of the Europeana Data Model v5.2.7*. The Hague: Europeana Foundation. Retrieved from: https://pro.europeana.eu/files/Europeana_Professional/Share_your_data/Technical_requirements/EDM_Documentation/EDM_Definition_v5.2.7_042016.pdf

- Fenlon, K., Senseney, M., Green, H., Bhattacharyya, S., Willis, C. & Downie, J. S. (2014). Scholar-built collections: A study of user requirements for research in large-scale digital libraries. Paper presented at *The 77th ASIS&T Annual Meeting*. (Seattle, WA, 31 October – 5 November, 2014).
- Fine, K. (1994). Compounds and Aggregates. *Noûs* 28, pp 137-58.
- Foulonneau, M., Cole, T. W., Habing, T. G. & Shreeves, S. L. (2005). Using collection descriptions to enhance an aggregation of harvested item-level metadata. *Proceedings of the 5th ACM/IEEE-CS Joint Conference on Digital Libraries* (Denver, CO, 7-11 June 2005).
- Frænkel, A. & Bar-Hillel, Y. (1958). *Foundations of set theory*. Amsterdam: North-Holland Publishing Co.
- FRBR Working Group on Aggregates (FRBR-WGA). (2011). *Final report of the Working Group on Aggregates*. The Hague: IFLA.
- Furner, J. (2004). Information studies without information. *Library Trends* 52, pp 427-447.
- Furner, J. (2010). Philosophy and information studies. *ARIST* 44, pp 161-200.
- Furner, J. (2016). “Data”: The data. In Kelly, M. & Bielby, J. (eds.) *Information Cultures in the Digital Age*. New York: Springer Publishing, pp 287-306.
- Galton, A. (2010). How is a collection related to its members? *Interdisciplinary Ontology: Proceedings of the Third Interdisciplinary Ontology Meetings* 3, pp 9-17.
- Gartner, R. (2016). *Metadata: Shaping Knowledge from Antiquity to the Semantic Web*. New York: Springer Publishing.
- Genette, G. (1997). *Paratexts: Thresholds of Interpretation*. Cambridge: Cambridge University Press.

- Glock, H.J. (2008). *What is Analytic Philosophy?* Cambridge: Cambridge University Press.
- Gonçalves, M. A., Fox, E. A., Watson, L. T. & Kipp, N. A. (2004). Streams, Structures, Spaces, Scenarios, Societies (5S): A formal model for digital libraries. *ACM Transactions on Information Systems* 22, pp 270-310.
- Gorman, M. (1998). AACR3? Not! *ALCTS Papers on Library Technical Services and Collections* 1998, pp 19-29.
- Guarino, N. & Welty, C. (2004). An overview of OntoClean. In Staab, S. and Studer, R. (eds.) *The Handbook on Ontologies*. New York: Springer Publishing.
- Guizzardi, G. (2005). *Ontological foundations for structural conceptual models*. PhD Thesis, Telematica Institute, Twente University, Enschede, the Netherlands.
- Hadro, J. (2015). NYPL digital collections platform: An introduction. *NYPL Labs* [blog]. New York: New York Public Library. Retrieved from:
<https://www.nypl.org/blog/2015/01/21/digital-collections-platform-intro>
- Heaney, M. (2000). An analytic model of collections and their catalogues. Bath, UK: UKOLN.
- Hughes, R. I. G. (1997). Models and representation. *Philosophy of Science* 64(supplement) [Proceedings of the 1996 Biennial Meetings of the Philosophy of Science Association. Part II: Symposia Papers (Dec. 1997)], pp S325-S336.
- Hunter, N. C., Legg, K. & Oehlerts, B. (2010). Two librarians, an archivist, and 13,000 images: Collaborating to build a digital collection. *The Library Quarterly: Information Community, Policy* 80(1), pp 81-103.
- Husserl, E. (1970). *Logical Investigations*. J.N. Findlay (trans.). London: Routledge & Kegan Paul. (Original publication 1900/1901.)

- International Federation of Library Associations (IFLA). (1961). *Statement of Principles Adopted by the International Conference on Cataloging Principles*. Paris: IFLA.
- IFLA Study Group on FRBR (IFLA). (1998). *Functional requirements for bibliographic records: Final report* [revised 2009]. München: K.G. Saur Verlag.
- Jett, J., Sacchi, S., Lee, J.-H. & Clarke, R. I. (2015). A conceptual model for video games and interactive media. *JASIS&T* 67(3), pp 505-17.
- Jett, J. (2015). Modeling worksets in the HathiTrust Research Center. *CIRSS Technical Report WCSA0715*. Champaign, IL: Center for Informatics Research in Science and Scholarship, School of Information Sciences, University of Illinois at Urbana-Champaign. Available via: <http://hdl.handle.net/2142/78149>
- Jett, J., Cole, T. W., Maden, C. & Downie, J. S. (2016a). The HathiTrust Research Center workset ontology: A descriptive framework for non-consumptive research collections. *Journal of Open Humanities Data* 2, p. e1. DOI: <http://doi.org/10.5334/johd.3>
- Jett, J., Han, M.-J., Cole, T. W. & Sarol, M. J. (2016b). Disambiguating descriptions: Mapping digital special collections metadata into Linked Open Data (LOD) formats. Poster presented at the *ASIS&T 79th Annual Meeting* (Copenhagen, Denmark, 14-18, October 2016).
- Jett, J., Humpal, N., Charles, V. & Lee, J.-H. (2017). What is a series, really? *Knowledge Organization* 44(1), pp 24-36.
- Jett, J. & Dubin, D. (2018). How are dependent works realized? *Proceedings of Balisage: The Markup Conference 2018* (Rockville, MD, 31 July – 3 August 2018).

- Jett, J., Fenlon, K., & Downie, J. S. (2018). Towards unifying our collection descriptions: To LRMize or not? *Proceedings of the 8th Conference of the Japanese Association for Digital Humanities* (Tōkyō, Japan, 9-11 September 2018).
- Joint Steering Committee for Development of RDA (JSCD-RDA). (2014). *RDA: Resource Description and Access*.
- Jonsson, G. (2005). Cataloguing of hand press materials and the concept of expression in FRBR. *Cataloging & Classification Quarterly*, 39(3/4), 77-86.
- Keet, C. M. (2006a). Introduction to part-whole relations: mereology, conceptual modelling and mathematical aspects. *KRDB Research Centre Technical Report KRDB06-3 (Tutorial reader)*, Faculty of Computer Science, Free University of Bozen-Bolzano, Italy.
- Keet, C. M. (2006b). Part-whole relations in Object-Role Models. *Lecture Notes in Computer Science* 4278, pp 1116-27.
- Kemp, R. (2008). Catalog/cataloging changes and Web 2.0 functionality: New directions for serials. *The Serials Librarian* 53(4), pp 91-112. DOI: 10.1300/J123v53n04_07.
- Kiryakos, S. & Sugimoto, S. (2018). Building a bibliographic hierarchy for manga through the aggregation of institutional and hobbyist descriptions. *Journal of Documentation* 75(2), pp 287-313. DOI: 10.1108/JD-06-2018-0089.
- Krier, L. (2012). Serials, FRBR, and library linked data: A way forward. *Journal of Library Metadata* 12(2/3), pp 177-87. DOI: 10.1080/19386389.2012.699834.
- Lagoze, C. & Fielding D. (1998). Defining collections in distributed digital libraries. *D-Lib Magazine* 4(11). Retrieved from:
<http://mirror.dlib.org/dlib/november98/lagoze/11lagoze.html>

- Le Bœuf, P., Doerr, M., Ore, C. E., & Stead, S. (eds.). (2018). *Definition of the CIDOC Conceptual Reference Model*. Version 6.2.3 (May 2018). Paris, France: ICOM/CIDOC.
- Lee, H.-L. (2000). What is a collection? *Journal of the American Society for Information Science* 51(12), pp 1106-13.
- Lee, H.-L. (2005). The concept of collection from the user's perspective. *The Library Quarterly* 75(1), pp 67-85.
- Lee, J.-H., Renear, A. H. & Smith, L. C. (2006). Known-item search: Variations on a concept. *Proceedings of the 69th ASIS&T Annual Meeting* (Austin, TX, 3-8 November 2006).
- Lee, J.-H., Tennis, J. T., & Clarke, R. I. (2012). Domain analysis for a video game metadata schema: Issues and challenges. *Lecture Notes in Computer Science, Theory and Practice of Digital Libraries 7489*, pp 280-5.
- Lee, J.-H., Jett, J. & Perti, A. (2015). The problem of "additional content" in video games. *JCDL'15 Proceedings of the 15th ACM/IEEE-CS Joint Conference on Digital Libraries* (Knoxville, TN, 21-25 June 2015). DOI: 10.1145/2756406.2756949.
- Lee, J.-H., Jett, J., Cho, H., Windleharth, T., Disher, T., Kiryakos, S. & Sugimoto, S. (2018). Reconceptualizing superwork for improved access to popular cultural objects. *Proceedings of the 81st ASIS&T Annual Meeting* (Vancouver, British Columbia, Canada, 10-15 November 2018). DOI: 10.1002/pr2.2018.14505501030.
- Lewis, D. W. (2013). From stacks to the web: The transformation of academic library collecting. *College & Research Libraries* 14(2), pp 159-76.
- Lubetzky, S. (1969). *Principles of Cataloging. Final Report. Phase I: Descriptive Cataloging*. Los Angeles: UCLA.

- Madison, O. M. A. (2009). The origins of the IFLA study on functional requirements for bibliographic records. *Cataloging & Classification Quarterly* 39(3-4), pp 15-37. DOI: 10.1300/J104v39n03_02.
- Masolo, C., Borgo, S., Gangemi, A., Guarino, N., & Oltramari, A. (2003). *The WonderWeb Library of Foundational Ontologies and the DOLCE ontology* (WonderWeb (EU IST Project 2001-33052) Deliverable D18).
- Masolo, C., Vieu, L., Bottazzi, E., Catenacci, C., Farrario, R., Gangemi, A., & Guarino, N. (2004). Social roles and their descriptions. In D. Dubois & C. Welty (eds.) *Knowledge Representation and Reasoning, KR2004*. Menlo Park, CA: AAAI Press, pp 267-77.
- McDonough, J., Kirschenbaum, M., Reside, D., Fraistat, N., & Jerz, D. (2010). Twisty little passages almost all alike: Applying the FRBR model to a classic computer game. *Digital Humanities Quarterly*, 4(2).
- Meghini, C., Spyrtos, N. and Yang, J. (2010). A data model for digital libraries. *International Journal on Digital Libraries* 11(1), pp 41-56.
- Miller, D. & Le Bœuf, P. (2005). Such stuff as dreams are made on: How does FRBR fit performing arts? *Cataloging & Classification Quarterly*, 39(3/4), 151-178.
- National Information Standard Organization (NISO). (2005). *Bibliographic references. ANSI/NISO Z39.29* [revised 2010]. Baltimore, MD: NISO.
- Nicolas, Y. (2005). Folklore requirements for bibliographic records: oral traditions and FRBR. *Cataloging & Classification Quarterly*, 39(3/4), 179-195.
- Nimer, C. L. & Daines, J. G. (2013). The development and application of U.S. descriptive standards for archives, historical manuscripts, and rare books. *Cataloging & Classification Quarterly* 51(5), pp 532-49. DOI: 10.1080/01639374.2013.764373

- Nurmikko-Fuller, T., Jett, J., Page, K. R., Maden, C., Willcox, P., Cole, T. W., Fallaw, C., Senseney, M. & Downie, J. S. (2015). Building complex research collections in digital libraries: A survey of ontology implications. *Proceedings of the 12th ACM/IEEE-CS Joint Conference on Digital Libraries* (Knoxville, TN, 21-24 June 2015).
- Palmer, C. L. (2004). Thematic research collections. In Schreibman, S., Siemens, R. & Unsworth, J. (eds.) *A Companion to Digital Humanities*. Hoboken, NJ: Blackwell Publishing.
- Palmer, C. L., Knutson, E., Twidale, M. & Zavalina, O. (2006). Collection definition in federated resource development. *Proceedings of the 69th ASIS&T Annual Meeting* (Austin, TX, 3-8 November 2006).
- Palmer, C. L., Zavalina, O. L. & Fenlon, K. (2010). Beyond size and search: Building contextual mass in digital aggregations for scholarly use. *Proceedings of the 73rd ASIS&T Annual Meeting* (Pittsburgh, PA, 22-27 October 2010). DOI: 10.1002/meet.14504701213.
- Panizzi, A. (1841). Rules for the compilation of the catalogue. *Catalogue of Printed Books in the British Museum 1*, p [v]-ix.
- Pomerantz, J. (2015). *Metadata*. Cambridge, MA: The MIT Press.
- Powell, A., Heaney, M. & Dempsey, L. (2000). RSLP collection description. *D-Lib Magazine* 6(9). Retrieved from: <http://www.dlib.org/dlib/september00/powell/09powell.html>
- Ranganathan, S. R. (1962). International conference on cataloging principles, Paris 9-18 October 1961 and its findings (Paper B). *ALIS* 9(1), pp 15-38.
- Renear, A. H. & Choi, Y. (2006). Modeling our understanding, understanding our models: The case of inheritance in FRBR. *Proceedings of the 69th ASIS&T Annual Meeting* (Austin, TX, 3-8 November 2006).

- Renear, A. H., Wickett, K. M., Urban, R. J., Dubin, D. (2008a). The return of the trivial: Problems formalizing collection/item metadata relationships. *Proceedings of the 5th ACM/IEEE-CS Joint Conference on Digital Libraries* (Pittsburgh, PA, 16-20 June 2008).
- Renear, A. H., Wickett, K. M., Urban, R. J., Dubin, D. & Shreeves, S. L. (2008b). Collection/item metadata relationships. *DCMI'08 Proceedings of the 2008 International Conference on Dublin Core and Metadata Applications* (Berlin, Germany, 22-26 September 2008).
- Riley, J. (2017). *Understanding Metadata: What is Metadata, and What is it For?* [Niso Primer Series.] Baltimore, MD: NISO. Retrieved from:
https://groups.niso.org/apps/group_public/download.php/17446/Understanding%20Metadata.pdf
- Riva, P., Le Bœuf, P., & Žumer, M. (2017). *IFLA library reference model: A conceptual model for bibliographic information*. The Hague, Netherlands: IFLA.
- Salaba, A. and Zhang, Y. (2007). Functional requirements for bibliographic records: From a conceptual model to application and system development. *Bulletin of the ASIS&T* 33(6), pp 17-23.
- Simons, P. (1987). *Parts: A Study in Ontology*. Oxford: Oxford University Press.
- Shadbolt, N., Hall, W. and Berners-Lee, T. (2006). The semantic web revisited. *IEEE Intelligent Systems Journal* 21(3), pp 96-101.
- Shera, J. (1966). *Documentation and the Organization of Knowledge*. Hamden, CT: Archon Books.
- Shreeves, S. & Cole, T. W. (2003). Developing a collection registry for IMLS NLG digital collections. *DC-2003 Proceedings of the 2003 International Conference on Dublin Core and Metadata Applications* (Seattle, Washington, 28 September – 2 October 2003).

- Smiraglia, R. P. (2001). *The Nature of "A Work": Implications for the Organization of Knowledge*. Lanham, MD: The Scarecrow Press, Inc.
- Smiraglia, R. P. (2002). Further reflections on the *Nature of 'A Work'*: An introduction. *Cataloging & Classification Quarterly* 33(3-4), pp 1-11.
- Smiraglia, R. P. (2005). Instantiation: Toward a theory. Presented at *Data, Information and Knowledge in a Networked World; Annual Conference of the Canadian Association for Information Science* (London, Ontario, 2-4 June 2005).
- Smiraglia, R. P. (2008). A meta-analysis of instantiation as a phenomenon of information objects. *Culture Del Testo e Del Documento* 9(25), pp 5-25.
- Smith, B. (2001). Fiat objects. *Topoi* 20(2), pp 131-48.
- Sugimoto, S., Kiryakos, S., Wijesundara, C., Monika, W., Mihara, T. & Nagamori, M. (2018). Metadata models for organizing digital archives on the web: Metadata-centric projects at Tsukuba and lessons learned. *DCMI'18 Proceedings of the 2018 International Conference on Dublin Core and Metadata Applications* (Porto, Portugal, 10-13 September 2018).
- Sullivan, D. (2014). The Yahoo directory — once the internet's most important search engine — is to close. *Search Engine Land*. Retrieved from: <https://searchengineland.com/yahoo-directory-close-204370>
- Svenonius, E. (2000). Bibliographic entities. In Svenonius, E. *The Intellectual Foundation of Information Organization*. Boston, MA: MIT Press, pp 31-51.
- Sweet, M. & Thomas, D. (2000). Archives described at collection level. *D-Lib Magazine* 6(9). Retrieved from: <http://www.dlib.org/dlib/september00/sweet/09sweet.html>
- Taniguchi, S. (2002). A conceptual model giving primacy to expression-level bibliographic entity in cataloging. *Journal of Documentation* 58(4), pp 363-82.

- Taniguchi, S. (2003). *A Conceptual Model Giving Primacy to Text-Level Bibliographic Entity in Cataloging: A Detailed Discussion*. [technical report]. Tsukuba, Japan: Institute of Library and Information Science, University of Tsukuba. Retrieved from: <http://user.keio.ac.jp/~taniguchi/report200301.pdf>
- Tillett, B. B. (1987). *Bibliographic relationships: toward a conceptual structure of bibliographic information used in cataloging*. PhD diss., Los Angeles, CA: University of California.
- Tillett, B. B. (1991a). A taxonomy of bibliographic relationships. *Library Resources & Technical Services* 35(2), pp 150-58.
- Tillett, B. B. (1991b). A summary of the treatment of bibliographic relationships in cataloging rules. *Library Resources & Technical Services* 35(4), pp 393-405.
- Tillett, B. B. (1992a). The history of linking devices. *Library & Technical Services* 36(1), pp 23-36.
- Tillett, B. B. (1992b). Bibliographic relationships: An empirical study of the LC machine-readable records. *Library Resources & Technical Services* 36(2), pp 162-88.
- Tillett, B. B. (2005). FRBR and cataloging for the future. *Cataloging & Classification Quarterly* 39(3-4), pp 197-205. DOI: 10.1300/J104v39n03_12.
- Tillett, B. B., Kuhagen, J. A., Cato, A. & Murotmaa, E. (2014). Letter to the editor. *Cataloging & Classification Quarterly* 52(3), pp 359-61. DOI: 10.1080/01639374.2014.897900
- Tourney, M. (2003). Caging virtual antelopes: Suzanne Briet's definition of documents in the context of the digital age. *Archival science* 3, pp 291-311.
- Unsworth, J. (2000). Scholarly primitives: What methods do humanities researchers have in common and how might our tools reflect this? Presentation given at *Humanities*

- Computing, Formal Methods, Experimental Practice* [symposium] (Kings College, London, 13 May 2000).
- Varzi, A. C. (1996). Parts, wholes, and part-whole relations: The prospects of mereotopology. *Data and Knowledge Engineering* 20, pp 259-86.
- Varzi, A. C. (2016). "Mereology." In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition). Retrieved from:
<http://plato.stanford.edu/archives/spr2016/entries/mereology/>
- Verona, E. (1959). Literary unit versus bibliographic unit. *Libri* 9, pp 79- 104.
- Vieu, L. & Aurnague, M. (2007). Part-of relations, functionality and dependence. *The Categorization of Spatial Entities in Language and Cognition* 20, pp 307-36. doi: 10.1075/hcp.20.18vie.
- Wetzel, L. (2009). *Types & Tokens: On Abstract Objects*. Cambridge, MA: The MIT Press.
- Wickett, K. M. & Renear, A. H. (2009). A first order theory of bibliographic objects. *Proceedings of the 72nd ASIS&T Annual Meeting* (Vancouver, BC, Canada: 6-11 November 2009)
- Wickett, K. M., Renear, A. H. & Urban, R. J. (2010). Rule categories for collection/item metadata relationships. *Proceedings of the 73rd ASIS&T Annual Meeting* (Pittsburgh, PA, 22-27 October 2010).
- Wickett, K. M., Renear, A. H., & Furner, J. (2011). Are collections sets? *Proceedings of the 74th ASIS&T Annual Meeting* (New Orleans, LA, 9-13 October 2011).
- Wickett, K. M. & Renear, A. H. (2012). The logical form of a metadata record. *Proceedings of the 9th ACM/IEEE-CS Joint Conference on Digital Libraries* (Washington, DC, 10-12 June 2012).

- Wickett, K. M., Sacchi, S., Dubin, D. & Renear, A. H. (2012). Identifying content and levels of representation in scientific data. *Proceedings of the 75th ASIS&T Annual Meeting* (Baltimore, MD: 26-31 October 2012).
- Wickett, K. M., Isaac, A., Fenlon, K. S., Doerr, M., Meghini, C., Palmer, C. L. & Jett, J. (2013). Modeling cultural collections for digital aggregation and exchange environments. *CIRSS Technical Report 201310-1*. Champaign, IL: Center for Informatics Research in Science and Scholarship, Graduate School of Library and Information Sciences, University of Illinois at Urbana-Champaign. Available via: <http://hdl.handle.net/2142/45860>
- Wickett, K. M., Isaac, A., Doerr, M., Fenlon, K., Meghini, C. & Palmer, C. [L.] (2014). Representing cultural collections in digital aggregation and exchange environments. *D-Lib Magazine* 20(5/6). Retrieved from:
<http://www.dlib.org/dlib/may14/wickett/05wickett.html>
- Wickett, K. M. (2018). A logic-based framework for collection/item metadata relationships. *Journal of Documentation* 74(6) pp 1175-89.
- Wilson, P. (1968). *Two Kinds of Power: An Essay on Bibliographic Control*. Berkeley, CA: University of California Press.
- Winston, M., Chaffin, R. & Herrmann, D. (1987). A taxonomy of part-whole relations. *Cognitive Science* 11, pp 417-44.
- Yeo, G. (2012). The conceptual fonds and the physical collection. *Archivaria* 73, pp 43-80.
- Žumer, M. & O'Neill, E. T. (2012). Modeling aggregates in FRBR. *Cataloging & Classification Quarterly* 50(5-7), pp 456-72. DOI: 10.1080/01639374.2012.679547.

Appendix A: Discussion of Functionality

In many cases, it appears that, especially for *components*, a part is *essential* to the whole. Some common examples used in the literature (Bittner & Donnelly 2005; Guizzardi 2005) are hearts for human bodies, and engines for automobiles. For our more library-centric examples, we might consider things like particular characters (i.e., people, fictional or otherwise) essential to particular works, but particular articles might not be essential to particular journal issues. Developing a good sense of *functionality* will give us a method of marking such distinctions.

Regarding *functionality*, we might begin by noting that if x is *functionally-part-of* y , then it is almost certainly the case that y is *dependent-on* x . While Husserl (1970) discusses *dependence*, both existentially and generically, the following definitions are adapted from Guizzardi (2005), who provides a treatment of parthood expressly designed for formal ontological work.

Guizzardi's first relevant definition here tells us that something is *existentially-dependent* on something else if it is necessarily the case that when one exists, then so does the other.

D.8 (Existential Dependence): $\forall x\forall y(\text{existentiallyDependentOn}(x, y) =_{df} \square (\text{exists}(x) \rightarrow \text{exists}(y)))$

Definition D.8 states that for all entities x and y , entity x is *existentially-dependent* on entity y , if and only if, it is necessarily the case that the existence of entity x implies the existence of entity y .

This next definition tells us that if something is an *essential-part-of* of something else, then it must be the case that something else is *existentially-dependent* on that that thing and the thing is necessarily *part-of* it.

D.9 (Essential Parthood):

$$\forall x \forall y (\text{essentialPartOf}(x, y) =_{df} (\text{existentiallyDependentOn}(y, x) \wedge \Box \text{partOf}(x, y)))$$

Definition D.9 states that for all entities x and y , entity x is an *essential-part-of* entity y , if and only if, entity y is *existentially-dependent-on* entity x and it is necessarily the case that entity x is *part-of* entity y .

In some cases, parts are existentially co-dependent with their wholes. A common example in the literature is that of a brain and the person it is *part-of*. In these cases, we say that the thing is an *inseparable-part-of* the whole.

D.10 (Inseparable Parthood):

$$\forall x \forall y (\text{inseperablePartOf}(x, y) =_{df} (\text{existentiallyDependentOn}(x, y) \wedge \Box \text{partOf}(x, y)))$$

Definition D.10 states that for all entities x and y , entity x is an *inseparable-part-of* entity y , if and only if, entity x is *existentially-dependent-on* entity y and it is necessarily the case that entity x is *part-of* entity y .

From these definitions, we can develop a series of axioms that allow us to formalize which metaproperties (*functionality*, *homeomerosity*, and *separability*) Winston, Chaffin, and Herrmann's mereonymous relationship pairings possess.

We first consider what it means to be a functional part of something, by rooting the concept of *functionality* with the concepts of *essentialness* and *inseparability*. So, we can say that when

something is *functionally-part-of* a whole, then it seems to either be an *essential-part-of* that whole, or an *inseparable-part-of* that whole.

A.53 (Functional Parthood): $\forall x\forall y (functionallyPartOf(x, y) \rightarrow (essentialPartOf(x, y) \vee inseperablePartOf(x, y)))$

Axiom A.53 states that for all entities x and y , if entity x is *functionally-part-of* entity y , then it is the case that either entity x is an *essential-part-of* entity y or entity x is an *inseparable-part-of* entity y .

Now, we can say that when something is *separably-part-of* a whole (e.g., like a book in a library collection or a tool in a toolbox), then it cannot be *inseparably-part-of* the whole.

A.17 (Separable Parthood): $\forall x\forall y (seperablyPartOf(x, y) \rightarrow \sim inseperablyPartOf(x, y))$

Axiom A.17 states that for entities x and y , if entity x is *separably-part-of* entity y , then it is not the case that entity x is *inseparably-part-of* entity y . Rather importantly, this axiom relies on definition D.10 which tells us that *inseparable* parts are both *existentially-dependent* on the whole and necessarily *part-of* it.

Together, these additional axioms allow us to construct a table that accommodates Winston, Chaffin, and Herrmann's meronymous relationships (Table A.1 above).

At first glance, the contents of Table A.1 seem plausible, but there are counter-examples for several of the underlying axioms. In particular, axiom A.53 (functional parts are either essential or

inseparable) seems vulnerable to counter-examples regarding existential necessity. There also seems to be some problems with axiom A.8 (homeomerous parts share properties).

Fortunately, Vieu and Aurnague (2007) leverage previous work by Husserl (1970), Simons (1987), Fine (1995), and Masolo et al. (2003, 2004) to provide a definition for *General Functional Dependence (GFD)* using a temporalized existence primitive—*exists-at-[time]* and lexical types.

While temporalization is very helpful for the variety of cases Vieu and Aurnague are examining, it is of somewhat limited benefit in the context this dissertation analyzes. Thereby, simplified versions of the supporting definitions and axioms for a simplified formalization of *GFD* are given below.

As a first step, Vieu and Aurnague define a *dependent-on* relationship, which says that something is *dependent-on* something else if it is necessarily the case that when one exists the other exists and they are not the same thing.

$$\mathbf{D.21 (Dependence):} \forall x \forall y \left(\text{dependentOn}(x, y) =_{df} \square \left((\text{exists}(x) \rightarrow \text{exists}(y)) \wedge \sim(x = y) \right) \right)$$

Definition D.21 states that for all entities x and y , entity x is *dependent-on* entity y if and only if it is necessarily the case that when entity x exists, then entity y exists and, it is not the case that entity x and entity y are not the same entity.

Table A.1: Meronymous Relationship Metaproperties

Metaproperty	part-of	proper-part-of	component-of	member-of	sub-quantity-of	constituted-of	feature-of	located-in
<i>Reflexive</i>	+	-	-	-	-	-	-	-
<i>Irreflexive</i>	-	+	+	+	+	+	+	+
<i>Symmetrical</i>	-	-	-	-	-	-	-	-
<i>Asymmetrical</i>	-	+	+	+	+	+	+	+
<i>Antisymmetrical</i>	+	+	+	+	+	+	+	+
<i>Transitive</i>	+	+	+	+	+	+	+	+
<i>Functional</i>	-	-	+	-	-	-	+	-
<i>Homeomeric</i>	-	-	-	-	+	-	-	+
<i>Seperable</i>	-	-	+	+	-	-	-	-

Next they reuse Masolo et al.'s (2004) *classified-as[-an]-X-at[-time]* axiom (simplified here). In simple terms this axiom says that if something is *classified-as* a member of a class then it exists.

A.54 (Classification): $\forall x(\text{classifiedAs}(x, X) \rightarrow \text{exists}(x))$

Axiom A.54 states that for all entities x , if entity x is *classified-as* kind X , then entity x exists.

Classification is an indispensable part of Vieu and Aurnague's definition for General Dependence.

D.22 (General Dependence):

$\forall X \forall Y \left(\text{isGenerallyDependentOn}(X, Y) =_{df} \Box \forall x \left(\text{classifiedAs}(x, X) \rightarrow \right. \right.$
 $\left. \left. \exists y (\sim(y = x) \wedge \text{classifiedAs}(y, Y)) \right) \right)$

Definition D.22 states that for all kinds X and Y , kind X is *generally-dependent-on* kind Y if and only if it is necessarily the case for all entities x that when entity x is *classified-as* kind X , then there exists an entity y such that y and x are not the same entities and entity y is *classified-as* kind Y .

Finally, we arrive at the simplified version of Vieu and Aurnague's definition for *GFD*.

D.23 (General Functional Dependence [GFD]):

$\forall X \forall Y \left(\text{isGenerallyFunctionallyDependentOn}(X, Y) =_{df} \Box \forall x \left(\left(\text{classifiedAs}(x, X) \wedge \right. \right. \right.$
 $\left. \left. \text{functioningAs}(x, X) \right) \rightarrow \exists y \left(\left(\begin{array}{l} (\sim y = x) \wedge \\ \text{classifiedAs}(y, Y) \wedge \\ \text{functioningAs}(y, Y) \end{array} \right) \right) \right) \right)$

Definition D.23 states that for all kinds X and Y , kind X is *generally-functionally-dependent-on* kind Y if and only if it is necessarily the case for all entities x that when entity x is *classified-as* and *functioning-as* kind X , then there exists some entity y such that y and x are not the same entity and entity y is both, *classified-as* and *functioning-as* kind Y .

Vieu and Aurnague go on to provide several definitions and axioms that support the development of theorems showing the transitivity of GFD and its propagation from types to supertypes. Simplified versions of these axioms, definitions, and theorems are provided below.

D.24 (Subtype): $\forall X \forall Y \left(isSubtypeOf(X, Y) =_{df} \forall x (classifiedAs(x, X) \rightarrow classifiedAs(x, Y)) \right)$

Definition D.24 states that for all kinds X and Y , kind X is a *subtype-of* kind Y if and only if it is the case that for all entities x , when entity x is *classified-as* kind X , then entity x is *classified-as* kind Y .

A.55 (Sub-functioning): $\forall x \forall X \forall Y \left((functioningAs(x, X) \wedge isSubtypeOf(X, Y)) \rightarrow functioningAs(x, Y) \right)$

Axiom A.55 states that for entities x and kinds X and Y , if entity x is *functioning-as* kind X and it is the case that kind X is a *subtype-of* kind Y , then it is the case that entity x is *functioning-as* kind Y .

D.25 (Disjoint Classes):
 $\forall X \forall Y \left(isDisjointFrom(X, Y) =_{df} \sim \exists x (classifiedAs(x, X) \wedge classifiedAs(x, Y)) \right)$

Definition D.25 states that for all kinds X and Y , kind X is *disjoint-from* kind Y if and only if there exists no entities x such that entity x is both, *classified-as* kind X and *classified-as* kind Y .

$$\mathbf{A.56 \text{ (Transitive GFD):}} \forall X \forall Y \forall Z \left((isGFD(X, Y) \wedge isGFD(Y, Z) \wedge isDisjointFrom(X, Z)) \rightarrow isGFD(X, Z) \right)$$

Axiom A.56 states that for all kinds X , Y , and Z , if kind X is GFD on kind Y and kind Y is GFD on kind Z and kind X is *disjoint-from* kind Z , then kind X is GFD on kind Z .

$$\mathbf{A.57 \text{ (Propagation of GFD):}} \forall X \forall Y \forall Z \left((isGFD(X, Y) \wedge isSubtypeOf(Y, Z)) \rightarrow isGFD(X, Z) \right)$$

Axiom A.57 states that for all kinds X , Y , and Z , if kind X is GFD on kind Y and kind Y is a *subtype-of* kind Z , then kind X is GFD on kind Z .

From here, it is necessary to step beyond the linkages between lexical types, to the linkages between entities. Fortunately, Vieu and Aurnague provide additional definitions, which are provided here in simplified form.

$$\mathbf{D.26 \text{ (Individual Functional Dependency [IFD]):}} \\ \forall x \forall y \left(\begin{array}{c} isIndividuallyFunctionallyDependentOn(x, y) =_{df} \\ \forall X \forall Y \left((isGFD(X, Y) \wedge isClassifiedAs(x, X) \wedge isClassifiedAs(y, Y) \wedge \right. \\ \left. isFunctioningAs(x, X) \rightarrow isFunctioningAs(y, Y)) \right) \end{array} \right)$$

Definition D.26 states that for all entities x and y , entity x is IFD on entity y if and only if it is the case that for all kinds X and Y : kind X is GFD on kind Y , entity x is *classified-as* kind X , entity y is *classified-as* kind Y , and it is the case that when entity x is *functioning-as* kind X , then entity y is *functioning-as* kind Y .

At this point Vieu and Aurnague's account takes a turn from complex to extremely complex, as they begin attempt to account to structural issues and the subsumption of one lexical type with another (via class hierarchy relations). However, for our purposes, we have accounted for where both axiom A.17 (separable parts) and the IFD metaproperty come from, and so we do not labor to explicate these matters any further.