1 CLASSIFICATION SPACE REDEFINED

INTRODUCTION

This chapter introduces the primary case of the Catalogue of Life, a relatively new kind of classification structure that pushes against the traditional assumptions and functions of biodiversity taxonomies. Historically and traditionally, one of the hallmark qualities of classifications that organize organisms is that they are internally consistent, meaning that those who build them employ the same methods and organizing principles throughout the classification system. In this way, taxonomies are epistemically charged—they convey particular models of the natural world that serve as scientific hypotheses. Each taxonomic opinion is based on observed evidence analyzed in light of a unified and consistent set of classificatory commitments. More often than not, because the nature of taxonomic work is such that scientists will specialize in particular organism groups, taxonomies are usually limited to a single genus or other small subset of the taxonomic tree. For example, a scientist may specialize in beetles, or annelids, or crustacea, and the taxonomies produced in each of these areas mirror the taxonomic commitments of the scientists who build them. And because of the niche-oriented production of these taxonomies, many classification systems circulate in parallel with one another within the scientific community. As we will find, that there are many approaches to building taxonomies creates an environment where many different structures are produced that can be incompatible with one another. Despite this, these different classifications can be equally valid because they were built by taxonomists

who use them to express a consistent, empirically produced hypothesis about how to order the world.

In direct tension with this traditional approach is the Catalogue of Life, which is focused on the efficient communication of species information with a global purview—that is to say, classification used primarily for data collection, communication, and retrieval. Specifically, the Catalogue is a particular kind of *composite classification*, which can be loosely defined as a structure that reconciles many local ("traditional") taxonomies within one space to serve as a unified taxonomic backbone structure for the organization of biological data. Backbone taxonomies, from this view, are user-oriented structures meant to serve primarily as ready-made hierarchical standards that can move biological data beyond what I call the "local" level of organization. In practice, composite classifications are also called consensus, common, or management classifications, terms I use interchangeably.

Not everyone is keen on these new composite structures, as many taxonomists feel the descriptive and hypothesizing paradigm of local taxonomies conflict with the epistemological commitments of consensus-based classifications. As taxonomies are brought together within the space of management taxonomies, various normalizing procedures are implemented to integrate these potentially conflicting subtaxonomic structures. Each taxonomy, ingested into the Catalogue as a database, must be negotiated with every other, meaning that conflicts naturally arise, and some taxonomies are altered to fit the management structure. To reconcile these conflicts, decisions must be made about which system, of two or more competing systems, is the best choice for inclusion into a common space. For example, if there is a conflict between two contributed taxonomies on the placement of the class Aves (the class level for birds), then the best taxonomy (as decided by the Catalogue editor) is based on prevailing scientific opinion. This exposes a significant difference between traditional and composite spaces: in the former, what is important is the correct placement of a taxon based on the internal requirements of the entire system to ensure that it is consistent; in the latter, the best taxonomic structure is that which represents community consensus within subfragments of the taxonomy, meaning, as a whole system, management taxonomies are

likely *inconsistent* throughout. To focus my analysis, the bulk of this chapter describes the Catalogue's consensus structure and assesses how a "global view" of taxonomic opinion is manufactured. As we will see, taxon concepts are layered historical entities, so organizing these concepts within composite systems of classification adds to this complexity enormously. The underlying assumption of these common systems is that efficient data retrieval is preferred over fragmented and distributed data practices, which hinder data interoperability. Such fragmentation, some say, limits scientific breakthroughs and more emergent types of knowledge that can occur when research groups engage in broadscale collaboration.

Taxonomic aggregation of this nature presents a number of problems, including that the systems produced through this activity are less outwardly legible to anyone but the classification builder. If, as Tony Rees indicates, a classification is intended "to provide a recognizable navigation structure so that persons entering the classification can (hopefully) find their way to their particular organisms of interest" (2009), then how is a user to understand the Catalogue's internal intellectual structure? As Patrick Wilson (1968, chap. 4) makes quite clear, a good bibliographical instrument is one that people can understand, in order to appropriately navigate. Users should be able to anticipate (or at least be able to discover) the epistemological commitments of any given system. Consensus taxonomies, however, consist of many taxonomic databases, and thus this epistemological context can be far more difficult to understand, especially if a user is not a specialist in biodiversity work. This problem is especially important for biodiversity students or scientists who might wish to use the Catalogue for research purposes. Complete and sound taxonomic metadata from contributing taxonomies is most essential in this aggregated environment, as it allows a user to track back to a taxonomic source for authority assessment and construction principles. But metadata at this level of perfection is rarely, if ever, present (a fact that applies to all data sets, not just those of the biodiversity persuasion).

To fully understand the influence—and begin to explicate the conflicts—that these taxonomies present, in this chapter I also introduce how consensus structures act as social epistemological instruments, built as they are

through intellectual compromise. "Social epistemology" as articulated by Margaret Egan and Jesse Shera (1952), Jesse Shera (1965), and Steve Fuller (2007, 2009) can help us understand the cultural function of these aggregative mechanisms. In a general sense, social epistemology examines how knowledge is shaped by social relationships and institutions (Goldman and Blanchard 2016). Thus, I am not only interested in how common classifications influence how people think about biodiversity in general and the problems introduced by common classifications to consensus proper, but also in what we should expect of classifications *at all*, given that they operationally define the state of knowledge within a discipline. Should they strive to represent divergence or manufacture accordance? Inasmuch as backbone taxonomies represent global systems, they represent a *kind* of consensus between scientific professionals. But what kind of consensus do they represent, and how are such mechanisms situated to tell us something that local taxonomies cannot? Certainly, consensus structures do not represent complete agreement.

At the same time, consensus structures align with how those in the field of information studies typically approach the collection and distribution of situationally relevant information to many different communities (Wilson 1973). This is to say that information from many domains can and should be accessible in a unified space to users and their varieties of contexts. Another assumption is that these composite structures are meant to be *general* systems, and that, perhaps, the audience for them is not, necessarily or primarily, the "local" communities that produced them, but rather, other allied fields that also require access to biodiversity information collocated at global and aggregative levels. So much of the discourse in knowledge organization and classification studies at the present is focused on the pluralization of the classification record (Andersen and Skouvig 2017). Universal systems are critiqued for good reason since, as with any system, they predominantly represent some communities and simultaneously ignore or harm others. Given these expressive qualities, classifications are unavoidably also political entities: they include and exclude, express and repress, facilitate and restrict. Some classifications are visible and some

are layered behind a veneer of technics, algorithms, and system interfaces. The goal of IS, at least in part, is to make visible what is otherwise hidden from sight. Finally, universal classifications are fundamental to how we come to imagine, understand, and *position* ourselves embedded within a system of individuals, objects, laws, ideas, and ecologies in this world. Universal structures are about far more than mere schemas and access systems. If our goal is to collectively create classification and knowledge systems that are more flexible, diverse, and just—and I think these are and should be some of our main goals in practice—then pointed and fundamental discussions about what they do and do not offer are paramount.

However, to say that the Catalogue of Life, as a composite system, is a universal system is to stretch the term of universality beyond its reasonable limits, a consideration we will discuss in due course. First, an overview of how the Catalogue functions.

A COMPOSITE INSTRUMENT

> Our goal, therefore, is to provide a hierarchical classification for the CoL and its contributors that (a) is ranked to encompass ordinal-level taxa to facilitate a seamless import of contributing databases; (b) serves the needs of the diverse public-domain user community, most of whom are familiar with the Linnaean conceptual system of ordering taxon relationships; and (c) is likely to be more or less stable for the next five years.
>
> —MICHAEL RUGGIERO ET AL.
> "A Higher Level Classification of All Living Organisms"

One way to look at biological taxonomies is that they are meant, at least in part, to facilitate information retrieval. Even as far back as Linnaeus, systems of nomenclature and taxonomies were meant to improve information recall and facilitate memorization of species amid continuously growing stores of biological information (Müller-Wille and Charmantier 2012; Ereshefsky 2001, 366). As Patrick Wilson described in some detail (1968, chap. 7), these documentary systems are meant to be consulted—they

must be an avenue by which a user can select the objects that meet a current need. This selection is predicated on the fact that, first, we have a bibliographical universe that is available in full, and second, that the documents therein are described in such a way to make them exploitable. This is to say, good advice about which document meets a certain circumstance or situation is dependent on being able to assess the full spectrum of organizational possibilities. Traditional taxonomies such as those described above, even if they fully satisfy the second criterion, fall short with regard to the first in that they usually contain only those taxa of interest in the particular domain in which they were produced. A taxonomy for weevils will only reference concepts or documents related to weevils and thus is really only of use to those who want information on this particular group. This is not a shortcoming of these descriptive instruments, of course, merely a product of how the production of knowledge is necessarily compartmentalized within the discipline of biodiversity. Given the fragmentary existence of these many descriptive taxonomies, a broader framework within which to control the full documentary universe is necessary. It is this requirement that the *consensus classification* is meant to satisfy.

A consensus classification is a comprehensive and pragmatic hierarchy intended to serve as a "framework for data integration" (Ruggiero et al. 2015a). These types of classifications—of which the Catalogue of Life is a prominent example but certainly not the only one—are intended to commingle the organizing structure of many traditional taxonomies in a singular space, with the intention of organizing data about, and potentially applying it to, numerous taxon groups. The Global Biodiversity Information Facility (GBIF) Backbone Taxonomy, for example—a management taxonomy of its own that builds on the Catalogue's structure—ostensibly brings data from anywhere together in one space for a multitude of purposes. To archive these distributed data, and to use them effectively as an integrated set of information, one needs a consensus classification as a practical tool to bring it all together. The taxonomic instrument as a whole is not intended to describe one unifying point of view but rather must reconcile many (perhaps, and often, contradictory) opinions in one space.

In a sense, consensus structures are a very familiar concept in IS—one can think of the Dewey Decimal System (DDC), for example, as a series of subclassifications brought together in one universal structure through which all produced documents can be organized and classed. The DDC is maintained with a coherent set of rules and guidelines that dictate how it can be revised over time. Each DDC class level—the main classes and subdivisions—is maintained as a separate domain-specific entity, each being updated and revised according to the documentation produced within the disciplinary area it represents. And because so-called general systems such as the DDC use academic disciplines as the primary method of classing knowledge (Langridge 1992, 21), we look to those disciplines (or perhaps more accurately, a discipline's produced documents) as the authority for that domain of knowledge: this is the notion of bibliographic warrant (Beghtol 1986). The recent revision of the Angiosperms in the DDC, for example, looked to phylogenetic and botanical literature for warrant to change their class structure (Green and Martin 2013). Biological consensus taxonomies function similarly: the authority for each taxon rests with the experts that research it, typically assessed by examining the relevant literature. In the sciences, however, a consensus classificatory approach is a much more problematic path, since its normative internal taxonomic structure can potentially conflict with the many other taxonomies that circulate in biodiversity spaces. Library classifications do not have the same epistemic and hermeneutic qualities as those that define biodiversity classifications.

A Constellatory Taxonomy

The Catalogue's management hierarchy uses the standard higher classifications of the traditional Linnaean system, including, but not exclusively, the ranks of kingdom, phylum, class, and order. Such an approach is suited to accommodating not only the largest amount of taxonomic information derived from contributing databases, but also to catering to the general public, who is most acquainted with this system (Ruggiero et al. 2015a). A consensus classification creates a backbone schematic that most easily facilitates communication between people and information systems. For example, if a person were looking for birds, they would look under Aves,

and for beetles under Coleoptera. Having a unified standard for data communication makes good sense—it helps people navigate an otherwise complex system, because they are generally acquainted with its principles. For example, a few years back I moved to the Midwest from Southern California. In the Midwest, a popular supermarket chain is owned by the same corporation as the supermarket I frequented in Los Angeles. Luckily, both stores were organized more or less the same way, so I was able to locate fairly easily items I wished to purchase. The same logic goes for biological consensus systems: since most people know the Linnaean binomial nomenclature system, for example, it makes sense to keep that system in place to facilitate access. Versus, for example, creating a classification or nomenclature system that completely eschews Linnaean order.

As of the publication of Ruggiero et al. (2015a), 1.3 million of the then 1.8 million species contributed by subsidiary taxonomies were at the class level or below. Given this, the higher-level approach for the Catalogue's taxonomy was implemented to fill in blanks, so to speak, for taxonomies that didn't provide categories above the class rank, as well as to account for the intense disagreements regarding how these higher levels should be oriented if they were, in fact, provided. The end result is an incredibly detailed taxonomic hierarchy that provides a working and iterative structure on which more detailed databases can be appended (Ruggiero et al. 2015a, 10–55).

Like adding ornaments one by one to a tree, the Catalogue of Life attaches contributed taxonomic data to the consensus classification at the highest rank represented in the database until the tree is filled out and complete (figure 1.1). If a contributed Global Species Database (GSD) taxonomy (a taxonomy that includes all extant species for a group throughout the world, versus one that focuses on only a limited geographic region), for example, covers the Droseraceae family (Culham and Yessen 2018)—a group consisting of carnivorous sundew species including the Venus flytrap—it will be appended to the Catalogue at that family level, while the management hierarchy would complete the higher levels not included in that contributed database. This allows the Droseraceae data to be interconnected with the broader taxonomic tree. From this connecting node

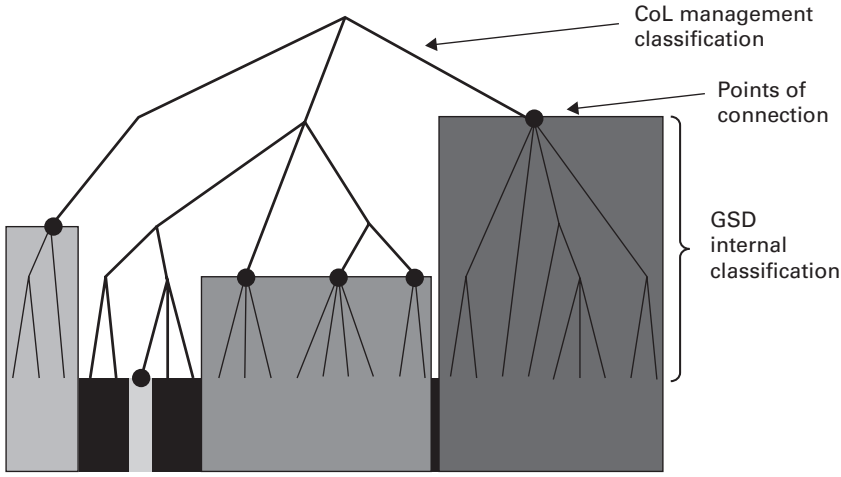

Figure 1.1
Schematic of the Catalogue of Life management hierarchy interacting with the GSD internal classifications. Original image label: "The Catalogue of Life retains the GSD's own classification below points of connection and uses the management classification above" (Species 2000 2016a). CC-BY 4.0, Catalog of Life, used by permission.

downward, the general composition of the Droseraceae of the contributed GSD is essentially maintained. Experts within the domain represented by a contributed taxonomy serve as the authority for its composition. It is sometimes the case that a database includes higher ranks that are not included in the overall Catalogue system, especially if the higher-level categories do not fall in line with the Catalogue's established hierarchy. In such cases, those higher levels are redacted before entry in the management classification.

The node of attachment for any taxonomic contribution is carefully and individually assessed through editorial review (Ruggiero et al. 2015a; Gordon 2009), and each species database is linked at only one node in the classification. Given the unique challenges posed by each taxonomy, no uniform standard exists by which this integration can take place. To reconcile the complexities and conflicts that arise from assembling these taxonomies, the Catalogue editorial team has articulated standards for certain required metadata elements to facilitate this reconciliation (Species 2000 2014). These standards provide a stable set of elements on which editors

can make an initial assessment of taxonomies in relation to the larger taxonomic whole. For example, as part of the Catalogue's standard data set, contributing databases are required to provide metadata identifying the highest taxon covered within the database, as well as all the classification ranks below that point down to the species rank.

The system is managed by a team of editors, led by the executive editor. The executive editor handles all the decisions regarding data transformation and maintains contact with the various data providers that contribute taxonomic systems to the management classification. Additionally, the executive editor has between seven and ten individuals serving on a board of editors that helps negotiate with regional hubs and umbrella providers such as those located in Europe, China, and South America (Roskov 2016b).

It is important to note the fundamental tension between subsidiary taxonomic authority and editorial review. Not all taxonomists agree on any given taxonomic alignment. But assuaging these disagreements is less a concern for the Catalogue than is representing, as accurately as possible given the constraints, the diversity of taxonomic opinion within one consensus space.

On the Taxonomic Commitment to Access

The Catalogue's founder, Frank Bisby, believed that in order to truly understand the extent of biodiversity knowledge, a mechanism was needed to aggregate data in a unified and coherent manner. Then, and only then, would scientists be able to address gaps in species knowledge on a global level. The methodology of the Catalogue was motivated by the concept of access, not description, as the predominant taxonomic function. As the Catalogue's executive editor conveyed, "We were firmly dedicated to the position that we should provide users [with] a single view of taxonomy. A single, simplified, view unified through different codes: zoological, botanical, bacterial codes of nomenclature" (Roskov 2016a). The Catalogue was intended to "disjoin" taxonomic practices "from the desk" of individual scientists, facilitating a cross-taxonomic discussion in the unified space of a hierarchy (Schalk 2016a). Many users, after all, are less interested in taxonomic order specifically, and more interested in finding information about a certain species—particularly users outside the domain of biodiversity (Bisby 2000).

The process of creating a consensus hierarchy was difficult and complicated for the Catalogue; it lasted approximately two years, full of a great deal of debate and deliberation (Orrell 2016; Schalk 2016b). And this debate continues in many sectors of taxonomic practice. The editorial group knew that the purposes of the Catalogue would eventually be challenged as being antithetical to the function and commitments of description-based, traditional systems. Two primary concerns arose as central to conceptualizing the hierarchy: (1) how to best establish taxonomic authorities for each taxon group and (2) how to articulate the function of the classification in contradistinction to local, description-based structures. To address the first concern, each portion of the management hierarchy was assessed in consult with experts in the field, as well as through examination of more than two hundred sources of literature (Ruggiero et al. 2015b). Both the consulted experts and the literature acted as the warrant for a final taxonomic view that would ultimately be adopted over any other. Clearly identifying and citing taxonomic authorities was central to this process, primarily because, as Peter Schalk reminded me, no final product would match the views of all scientists, and thus an established hierarchy needed to be justified at every level—a process eventually documented in the article, "A Higher Level Classification of All Living Organisms" (Ruggiero et al. 2015a).

Regarding the second concern, the Catalogue is about *managing data* and maintaining this sharing functionality, not about managing internally consistent taxonomic opinion. It is a pragmatic structure. Although the Catalogue is generally "reflective of phylogeny" within its hierarchy (Ruggiero et al. 2015a, 2), it does not purport to be a cohesive phylogenic tree built to answer phylogenetic questions. For the Catalogue, constructing a taxonomy has nothing to do with a correspondence to some a priori notion of how evolutionary relationships should be presented—at least insofar as its classification is distinct from, or similar to, prevalent scientific opinion. Warrant is used to assess how each subtaxa should be presented, meaning that the Catalogue can potentially have many distinct opinions commingled in its classificatory space. As Catalogue executive editor Yury Roskov indicated (2017), it doesn't matter how contributed individual classifications are built but rather that they include the full, global reach of any

given taxon and that it be up to date with regard to taxonomic method and opinion, citing appropriate literature for its construction. The assessment of any given contributed taxonomy often comes down, at least in large part, to the reputation of its contributor—both their individual reputation and that of their affiliated institution. In an environment where taxonomy cannot be pinned down and stabilized, all the Catalogue has to go on is whether or not a scientist has done due theoretical and methodological diligence in producing their classifications.

The underlying access functionalities of the Catalogue were articulated through careful examination of its potential use as a tool within the biodiversity world as well as in ancillary domains that equally needed access to the information. The Catalogue strives to manage a balance among four basic user groups: research scientists, policy makers, citizen-scientists, and educational users (Species 2000 2015a). Each of these user groups would come to understand the Catalogue in a vastly different way. Research scientists, for example, have used the database to understand the extent and distribution of species, as well as the valid names and synonyms for species concepts (Parr et al. 2012). The original impetus for the Catalogue is rooted in a policy-oriented mind-set, especially given that the instrument precipitated from the Convention on Biological Diversity (2003). In specific, the Global Taxonomy Initiative (Convention on Biological Diversity 2017a) identified how international scientific activities depend on taxonomic structures, including research on food security, conservation, and vector control. Concerning education, we see the Catalogue's backbone as integral to the structure of popular sites such as the Encyclopedia of Life (2018), which focuses on providing educational resources for a variety of curriculum levels.

One common thread among all these user groups is that their work, to greater or lesser extent, depends on the collocation of *all* species in one coherent space. With such pooling of intellectual knowledge, the Catalogue can facilitate global browsing and discovery mechanisms that any one local taxonomy would be unable to accommodate. Some of the possible access-oriented functionalities include providing (Species 2000 2015a):

- a single entry point into all distributed GSDs that make up its data space;
- a sense of global species coverage;
- universal searching tools for all extant species documentation and concepts, both inter- and intra-taxon;
- a hierarchy that can be fully browsed in tree form;
- a multilingual browsing mechanism accessible in eleven languages other than English.

Though these functionalities are productive, they point to the utility of the Catalogue as, primarily, a closed and self-referential entity—that is, these are all functions performed *within* the space of the Catalogue. Bringing the Catalogue into conversation with other, traditional taxonomies problematizes the veracity of its internal space. A host of other access functions require taking the Catalogue's taxonomic structure en bloc and using its backbone as the information architecture for other database and interface systems, such as the Global Biodiversity Information Facility (GBIF) and the Encyclopedia of Life (EoL). These systems incorporate the Catalogue's management structure into their systems, rather than creating a management taxonomy wholesale on their own.

Quite apart from the question of whether systems like the Catalogue or GBIF can be thought of as reflecting or conflicting with the aims of local (or description-based) taxonomies, one cannot ignore, as Slota and Bowker (2015) have conveyed, the broadscale impact these backbone structures can have on, for example, the articulation of conservation policy and new directions for scientific endeavors. In this way, systems like these directly shape how we come to understand the purchase and influence of this biodiversity knowledge on society. For better or for worse, global systems like the Catalogue provide us a mechanism for understanding the global state of affairs of biodiversity knowledge.

As a product of scientific collocation, the Catalogue is an instrument of *social epistemology*, as both a system that expresses a certain plurally driven, epistemic point of view, and one that modified how we understand the import of these data outside the biodiversity domain.

CONSENSUS TAXONOMIES AND SOCIAL EPISTEMOLOGY

The use of the term *social epistemology* deserves a bit of attention here, given its specific context within the history and literature of information studies. Social epistemology was introduced as a term within IS (Furner 2004; Fuller 2007, 69) through the publication of Margaret Egan and Jesse H. Shera's (1952) influential article, "Foundations of a Theory of Bibliography," which was later expanded on in Shera's monograph, *Sociological Foundation of Librarianship* (1970). In the former, Egan and Shera relate how the practice of bibliography had become fragmented as a discipline and thus needed more integrative approaches to communicate ideas both within and without the subdomain. The then-current field of bibliography (as it existed in 1952) was functioning as a series of separate groups, each "running around and around its own little circuit" (1952, 125). To illustrate a possible remedy, they introduce the metaphor of a railroad system to represent the pinnacle of information communication: an interconnected system primed to take advantage of the emergent knowledge that comes from cross-group information exchange. To Egan and Shera, complex forms of knowledge emerge from circumstances that comprise multilayered communication channels that exist *within* each specialist group, *among* these groups, and *between* specialists and the broader global information community (1952, 126).

In such a model, each specialist group could have a specific, and perhaps even conflicting, purpose from any other—and this diversity produces unexpected and sophisticated discoveries about the state of bibliographical knowledge. Part of this new approach was to take a more nuanced view of the *process* of information communication and to think about how the act of acquiring information had particular social impacts related to how one person (or group) could situate themselves "in relation to the total [intellectual] environment" (Egan and Shera 1952, 132). That the Catalogue embeds conflicting taxonomic opinion within its structure, for example, may not be a problem if we also understand that the biodiversity field of knowledge is equally defined by conflict of opinion. This metaphorical railroad system represented the total information environment—the entire documentary universe—and without the proper exchange of these "intellectual products" (Furner 2004, 793) within a bibliographic system, there

could be no way to situate individual needs within a broader social environment. A critical impetus for the advancement of a discipline comes from the incisive insights that come from extra-disciplinary spaces. On top of this, we shape our research trajectories based on the findings of others and often act in response to their ideas. Thus, the more we share, the more we expand our disciplinary boundaries. Egan and Shera called this new proposed disciplinary emphasis *social epistemology*, which focused on the recursive "impact of knowledge on society" (Shera 1965, 107) and scientists.

A second significant proponent of an IS-centric notion of social epistemology is Steve Fuller, a professor of social epistemology at the University of Warwick, whose work emphasizes the specific material and documentary concerns of the discipline. Fuller defines social epistemology as the normative processes by which we "seek to provide guidance on how and what we should know on the basis of how and what we actually know" (2007, 177). He claimed that the institutional locus of the practice of social epistemology lies within the field of information studies, given that it is engaged in activities that validate and recommend access to, and use of, reliable and relevant information to society at large (2007, 6). Librarians—and, by extension, modern information professionals broadly speaking—"regard documents as sites for studying the multiply embedded social relationships among producers, consumers and objects of knowledge" (2007, 69). This curation of documents—assessment of a document's aboutness, value, quality, and relationship to a specific user—in turn affects how society interprets, and puts into use, the products of cultural and intellectual thought. The focus on the "document" as a core aspect of what we might call a socially geared epistemic agenda in IS makes good sense, especially since it is these collectively produced and organized knowledge systems (classifications, in this context) that shape our aesthetic and material relationships to the external world. Classifications help us see and outline what we know, and by what terms. In this way, the information scientist is not only a curator of the past, but also must be able to anticipate a need and direction for future intellectual work (Shera 1965, 109).

Of course, Steve Fuller (2007, 69–73) has certainly pointed out the dangers of a classification system oriented to its "dominant users," rather

than, say, focused on the internal "organizational mapping functions" of its hierarchy. Moving away from the latter detrimentally alters a classification's epistemic orientations to one in which "success" proper is not defined as truth, but rather by a systemic alignment with expected public concerns and needs. Such an approach produces a reactive and service-oriented normative ethic, rather than one that focuses on curation and validity.

This existent dichotomy can be seen regularly in the discourse of biodiversity science. The divide usually manifests in conflicts between those who see taxonomies as hypothesis-driven arguments, and thus disagree with the aims of consensus structures, and those that see taxonomic backbones as user-oriented systems that facilitate the work of biodiversity informaticians. The former believe that sound inference comes from taxonomies that are internally consistent and, thus, facilitate more effectual work. The latter believe that such consensus structures actually unearth disciplinary inconsistencies and that they further a more global understanding of biodiversity issues. Of course, the reality of this situation is that the supposed "ends" of both of these approaches is to facilitate sound biodiversity work—the conflict lies in how they believe scientific consensus and agreement can come to be realized and represented in information spaces. Are the systems that we design to organize knowledge meant to impose order artifically onto the process of disciplinary knowledge production that is, by definition, defined by disagreement? What these critiques tell us is that, in addition to their use outside the domain of biodiversity work, consensus systems also are used to examine the shape of the knowledge *within* the domain itself. In the next chapters we closely examine these possibilities and critiques, so as to better understand these respective positions and the potentialities and limitations of consensus systems.

INSTRUMENTS OF DIVERSE OPINION

One function of this chapter has been to show how consensus taxonomies serve a particular purpose in biodiversity work that differs markedly from that of their "descriptive," or epistemically unified counterparts. These purposes align with access-oriented commitments that promise broad

taxonomic representations across the full suite of taxa. Whereas description-oriented taxonomies are situated for local use, often and primarily within the domain of biodiversity taxonomic work, consensus structures allow that local knowledge to gain traction in domains that extend beyond their disciplinary applications. Consensus structures allow the discipline of biodiversity taxonomy to assess its internal contours—identify data gaps, for example—as well as to influence the circulation and constitution of knowledge beyond these borders, such as in policy work, conservation, and education. The Catalogue acknowledges this broader influence of its work as a tool for societal communication (Ruggiero et al. 2015a) and biodiversity knowledge implementation.

Returning to Egan and Sher's railroad metaphor, we can with very little effort relate and extend this notion to consensus structures within the biodiversity world. As mechanisms designed primarily for the retrieval of information—insofar as they are formed to understand the global extent of knowledge-via-data—they relate individual, description-based taxonomies within the broader biodiversity knowledge landscape. Backbones facilitate communication, however imperfect those structures may be. By constructing a centralized repository, we can better understand, for example, how biological knowledge favors certain charismatic species; what parts of the globe require more attention and investment; and finally, where our knowledge structures have erred with regard to syntactic and semantic qualities.

But the collocation of taxonomies into a unified space is not without its controversy. Any taxonomic instrument functioning on consensus paradigms—biodiversity, documentary, or otherwise—should be sensitive to the fact that these structures have drastic ramifications in a vast array of social spaces, within and without its domain of production, and with expected and unintended consequences. Our organizational tools mobilize bodies in the natural and social spaces that we inhabit. They also dictate which objects and organisms are prioritized and visible over any other. As Matthew Hull (2012, 260) makes clear, documents "generate larger-scale forms of sociality . . . not only directly as instruments of control but also as vehicles of imagination."

In "The Modernity of Classification" (2011), Mai recommends that we should begin to understand and interpret classifications, not as accurate

models of the world per se, but as models of the way the world appears to those that construct them. There are many ways to view the organization of the world, and so our classifications should embrace this diversity. Extending this notion, some posit that there is a need for consensus structures in biodiversity taxonomy precisely because such diversity of opinion exists within the discipline and, therefore, there is a need to represent this discord in functional and operational ways. So long as taxonomists maintain divergent opinions, consensus structures can help operationalize this collective activity in structures that bring attention to these differences.

If the Catalogue exposes anything, it is that integrating one (unique, consistent, and local) taxonomy into a space of many (diverse and global) is a delicate process that questions the very function that taxonomies serve as hermeneutic instruments within biodiversity work. This becomes especially important because no single person can control what and how these structures are implemented or repurposed once they are made available to the wider public.

One of the biggest impediments to the ready communication of taxonomic knowledge is, as articulated by Thomas Orrell, the adherence to one and only one way to understand taxonomic practice. "Why does dogma overtake certain areas of the discipline?" Orrell asked, "Why is it that traditions can't be seen as separate from functional classifications of information?" (2016). Far too often, knowledge producers cannot see past their own methods and implementations to understand the utility of common systems that introduce a host of benefits that traditional, description-oriented systems simply cannot offer. The presence of multiple rules of nomenclature is an example of this dogmatic division Orrell mentions. Having only one rule of nomenclature would be much more effective, yet we have many. Our traditions maintain divisions and inadvertently also make it difficult to communicate the fruits of our research. What is it that stops us from agreeing on one classification, Orrell inquired, even as that one classification is meant to function differently from what we have come to historically expect from these structures?

Unearthing what common consensus structures can and cannot do is essential. First, a primary goal of the Catalogue is to provide a unified

classification, not a universal one. The Catalogue's intent is to bring knowledge together for access and rearticulation, not to create a structure that must maintain authority in all domains of taxonomy. Second, the Catalogue is intended to provide a structure that can communicate diverse opinion, not (necessarily) one that seeks to provide and argue a singular taxonomic perspective. These distinctions are important in a domain that depends on taxonomies as argument-making instruments at the core of its work. Taxonomic work can and should function on two very distinct professional levels: one that requires a taxonomy to facilitate data transfer *and* one that can (and must) be used to serve as the structure by which taxonomic arguments are made. How can we divorce the expectation that a taxonomy must be an opinion from the idea that a taxonomy can serve primarily as a facilitator of information? The answer, in part, lies in better understanding the possibilities and limitations of the latter type of structures.

My hope is that unpacking these distinctions brings a broader awareness to their utility in various social spaces, even in spaces outside biodiversity work, such as in information studies. Systems like the Catalogue relinquish a certain level of conceptual and informational control in deference to necessarily embracing the fact that knowledge is diverse. Let it be known that, below, I argue that the space of the Catalogue is not as diverse as it might otherwise seem—nor is it truly plural in any secure or structural sense. This is especially true when we think about the kinds of data the Catalogue is unable to absorb due to its epistemic assumptions. However much one might value composite systems, it is my suspicion that the library world should look to biodiversity work for models on how to best represent knowledge, especially given the pressure placed on these institutions to embrace more pluralistic and diverse approaches to the organization of knowledge. We can learn from other disciplines as much through their trials and errors as through their successes.

Next we examine some higher-level properties of taxonomies and describe how taxonomies are instruments of power that exert in both material and epistemic ways. We then return to the Catalogue and examine its benefits and drawbacks in greater detail.