Choose a size for the type: Large Medium Small
Go to Prologue
Go to Everything is Miscellaneous home page

C h a p t e r   O n e

The New Order of Order

Before the Web, the word browsing was usually a polite way of telling a salesperson to buzz off. "May I help you?" a salesperson asks. "I'm just browsing," you reply with a little smile. With that word, a customer declares a lack of commitment. With that smile, she asserts that she's within her rights: "Just try and stop me, salesboy!"

Browsing is more than window-shopping, fantasizing about what it would be like to own something or resenting those who do. You browse when you intentionally ignore the organizational structure the store has carefully imposed on its stock. You have a hankering to read something about the Civil War, but the bookstore has Civil War books strewn about the fiction, nonfiction, biography, and travel sections, all neatly arranged in individual aisles and on individual shelves. Or you're in the mood for some light reading—nothing more specific than that—and appealing books pop up on just about every shelf. The store helps you violate its order by providing tables at the front with staff picks, new books, and bargains, but it can't anticipate all the ways every customer who walks in the door is going to want to browse. So it has to depend on its sales staff to help its customers find the exact book they want when they say, "I need something as a sweet sixteen present." Sometimes they get it right; sometimes—say, if the available staff person thinks all the kids still have a favorite Beatle—they lose the sale as the customer decides it would be safer just to give her niece a check.

The normal organization of a store works well enough if you come in knowing what you want: Go to the fiction shelf, find the "A" section conveniently located at the beginning of the alphabetized authors, and locate that copy of Pride and Prejudice for your niece. But discovering what you want is at least as important as finding what you know you want. Our bookstores look like they prefer seekers over browsers because the usual layout works well for people trying to find what they came in for, whereas there are almost as many ways to organize for browsers as there are browsers. An order that works for one interest may not for others—clustering all the books about the Civil War would help the Civil War buffs but would pull Gone with the Wind off the shelf where the historical fiction buffs browse. On the other hand, dumping the shelves into a bookstore-sized pile of books would turn browsing into pawing through printed rubble.

If only there were a way to arrange the stuff in stores so that every possible interest could be captured. When we know what we want, we'd find it immediately. When we want to browse, the store would rearrange itself based on our needs and interests, even when we aren't quite sure what those are.

At Apple Computer's iTunes music store, it's already happened. For decades we've been buying albums. We thought it was for artistic reasons, but it was really because the economics of the physical world required it: Bundling songs into long-playing albums lowered the production, marketing, and distribution costs because there were fewer records to make, ship, shelve, categorize, alphabetize, and inventory. As soon as music went digital, we learned that the natural unit of music is the track. Thus

was iTunes born, a miscellaneous pile of 3.5 million songs from over a thousand record labels. Anyone can offer music there without first having to get the permission of a record executive. Apple lets customers organize the pile any way they want and markets through their customers' choice of tracks and playlists rather than to the mass market. By making music miscellaneous, Apple has captured more than 70 percent of the market.

And the iTunes store isn't even all that miscellaneous. It's a spreadsheet that can be sorted by the criteria iTunes provides: the track's name, length, artist, album, genre, and price. If you want to browse, first you pick the genre, artist, and album, in that order. If you want to browse by the artist and then by genre, you can't. If you want to browse by mood, language, or date, you can't. Even though iTunes is all digital, there are still more ways you can organize your physical collection of CDs. The problem isn't that iTunes has chosen an inappropriate set of criteria for sorting, although that certainly could be argued. The real problem is that iTunes accepts the premise we've had to operate under in the physical world: that there is a set of appropriate criteria.

While iTunes is parsimonious in its built-in ways of sorting, it generously enables customers to create their own playlists, pulling together songs from across the entire ocean of tunes. By allowing customers to then publish their playlists—and rate and comment on other people's—iTunes provides as many ways to navigate its inventory as there are customer moods and interests. That lets in an important breath of the miscellaneous and shows that iTunes has learned the right lesson: To get as good at browsing as we are at finding—and to take full advantage of the digital opportunity—we have to get rid of the idea that there's a best way of organizing the world.

Everything Has Its Places

It won't be easy. The world started out miscellaneous but it didn't stay that way, because we work so damn hard at straightening it up. Take eating, the most basic bodily activity we do on purpose. Preparing to serve a meal is a complex dance of order. We have separate storage areas for each of the implements of eating: silverware, plates, glasses, napkins. Each of these areas has its own principle of organization: plates stacked by type, clustered by pattern and how "good" they are, juice glasses separated from wineglasses and from tumblers, silverware tucked cozily into special compartments. We then break these items apart to form patterns (salad fork to the left of dinner fork) determined by tradition, one set per participant in the meal. When we clear the table, we recluster the items by type because we generally don't wash the implements in the same pattern we used when we were eating with them: We do all the plates at once, put the silverware into a pot to soak, and line up the glasses next to the sink. If we stack the dishwasher, we cluster the items yet another way. When it's all clean and dry, we again store everything in its initial order, completing our nightly choreography.

We juggle multiple principles of organization without even thinking about it. You know what goes in your spice rack and what doesn't, even though the principle of order is hard to find: What makes dried leaves (oregano), dried seeds (nutmeg), and dried bark (cinnamon) all spices? All add a little more flavor to a dish? But so do chocolate sprinkles, and they don't count as a spice, despite coming in a shaker the size and shape of an oregano bottle. And even if you count salt and black pepper as spices, you probably don't keep them in the spice rack, because you use them too often. Without pausing for thought, you have coordinated four intersecting sets of criteria: how big the bottles are, what the contents are used for, which part of a meal they're applied to, and how frequently you need them.

The same is true for every room, every closet, and every tabletop in our houses. Even the oddest, most random items have their place. The gin-flavored floss someone gave you as a joke goes on the top shelf of the bathroom cabinet, while the Miss Piggy night-light goes in with the rarely used electrical equipment in the box in the basement. If you're genuinely stumped, you'll probably throw it in a box of things to give away.

The two processes by which new things are introduced into our homes are typical of how we handle information: We go through new arrivals and then we put them away. We go through the mail and file it in the special places we have for bills (the desk), cards from relatives (the refrigerator door), and junk mail (the trash). We go through bags of groceries and put the food away within minutes of bringing it into our house. We address these elements of disorder—unsorted mail in the mailbox, groceries sorted by relative weight into bags by a clerk in the store—with remarkable alacrity.

There isn't a part of our homes that is truly unordered, except perhaps under our beds, and for many

of us even that is the site of the spontaneous ordering of dust into bunnies.

We invest so much time in making sure our world isn't miscellaneous in part because disorder is inefficient—"Anybody see the gas bill?"—but also because it feels bad. Knowing where things are and where things go is essential to feeling at home. If cleanliness is next to godliness, then slovenliness is next to The Odd Couple's Oscar Madison. And who wants to be next to Oscar Madison?

We've been raised as experts at keeping our physical environment well ordered, but our homespun ways of maintaining order are going to break—they're already breaking—in the digital world. The most visible breach so far: the folder on the family computer that stores the digital photos.

If you're managing to keep your digital photos well organized, congratulations. But you're probably going to lose the battle sooner rather than later. It's a simple matter of numbers. A typical album you buy at your local camera store holds between fifty and two hundred paper-based photos. You likely have a thousand photos in all your albums put together. A thousand photos, each with its own story. "Oh, remember how Mimi was always wearing that silly cowboy hat? And there's Aunt Sally on the beach. She was so sunburned we had to take her to the hospital, and that funny doctor said we should just baste her in barbecue sauce."

Now check your computer. If you have a digital camera, you may well have saved over a thousand photos in just the past few years. It's only going to get worse. Digital cameras started outselling film cameras in the United States in 2003 and worldwide in 2004. And, in 2004, 150 million cell phones with cameras were sold, almost four times the number of digital cameras. Because digital photos are virtually free, we're tempted to take more and more pictures, sometimes just in the hope that one will come out well. We're also keeping more of the photos, and not always because we want them: Since our cameras apply names like "DSC00165.jpg" to our photos, it's easier to keep bad photos than to throw them out. To keep them, we just press a button to move them from our camera. To get rid of them, we have to look at each one, compare it with the others in the series, select the bad ones, press the Delete button, and then confirm our choice.

As a result, we are loading onto our computers thousands of photos with automatically generated names that mean nothing to us. When you have ten, twenty, or thirty thousand photos on your computer, storing a photo of Aunt Sally labeled "DSC00165.jpg" is functionally the same as throwing it out, because you'll never find it again.

We're simply not going to be able to keep up. Even obsessive-compulsives have only twenty-four hours in a day. Perhaps technology will get better at automatically figuring out what and who is shown in a photo. Or perhaps labeling photos will become a social process, with others pitching in to help us organize them. The user-based organizing of photos is already happening on a massive scale at Internet sites like Flickr.com, where people can post their photos and easily label them, allowing others to search for them. Moreover, anyone can apply descriptive labels to photos and create virtual albums made up of photos taken by themselves and strangers. What's clear is that however we solve the photo crisis, it will be by adding more information to images, because the solution to the overabundance of information is more information.

We add information in the real world by putting a descriptive sign on the shelf beneath a product, sticking a label on a folder, or using a highlighting pen to mark the passages that we think will be on the test. The real world, though, limits the amount of additional data we can supply: Staples has to keep the product information labels on the shelves small enough so they won't obscure the product; a manila folder's label can't have more than a few dozen characters on it without becoming illegible; and if previous students have already highlighted every other sentence in your textbook, the marks you make won't add much information at all. In the digital world, these restrictions don't hold. The product listing on the Staples Web site can link to entire volumes of information, our computers can store more information about a desktop folder than is actually in the folder, and if the digital textbook has had every word highlighted by previous readers, a computer could show us which sections have been highlighted by the majority of A students.

Such features are not just cool tricks. They change the basic rules of order. When we come across the paper photo from 2005 of Aunt Sally on a beach in Mexico at sunset celebrating cousin Jamie's birthday, with the twins in the background playing badminton, we have to decide which one spot in one album we're going to stick it into. If it were a digital album, we wouldn't have to make that choice. We could label it in as many ways as we could think of: Aunt Sally, Mexico, 2005, beach, birthday,

twins, badminton, sunset, trips, foreign countries, fun times, relatives, places we want to go back to, days we got sunburned. That way we could have the computer assemble albums based on our interests at the moment: all the photos of the cousins, all the trip photos for the past five years, all the photos of Aunt Sally having fun. The digital world thereby allows us to transcend the most fundamental rule of ordering the real world: Instead of everything having its place, it's better if things can get assigned multiple places simultaneously.

Yours, Mine, Ours

These types of changes create effects that are rippling through our social world. Recently, my sister-in-law organized some of our parents-in-law's physical photos into a traditional album as an anniversary gift from all of us. She balanced the pictures so that her kids didn't dominate, remembered to include photos of dear friends even if the only available snaps were rather unflattering, and carefully placed them in an album in chronological order. Had she gone off and done this without asking, it would have been highly presumptive because building an album is often a ritual a family shares. We hand the pictures around, clucking at silly expressions and worse haircuts, laughing at the escapades the photos capture. Together we construct our past for the future, making the decisions about which photographs to put next to which, "chunking" the smoothness of experience into lumps of meaningful memories.

We do something quite different with our digital photos. A digital album is the same as an iPod music playlist: a way to remember a particular arrangement of photos. A single photo can go on dozens of playlists at virtually no cost. So if my sister-in-law takes five hundred photos during her trip to Paris, she can make one digital album that focuses on her kids' reactions, another of interesting faces of Parisians, and another of every plate of food she ate. If she chooses to share her digital photos with her extended family, perhaps one of us will cluster the photos of public art or of her kids making faces. There's no limit to how many albums we can assemble. So, we're no longer forced to carefully construct a single shared path through memory. Rather, the more ways our digital photos can be sorted, ordered, clustered, and made sense of—the more miscellaneous they are—the better. We lose the requirement that a family get on the same page (literally) about its memories. And if albums are the archetypes of memory, memory becomes less what we have assembled and locked away and more what we can assemble and share.

The changes we're facing are not just personal. We have major institutions dedicated to keeping the world from slipping into the miscellaneous. The Library of Congress owns 130 million items, including 29 million books on 530 miles of shelves. Every day, more books come into the library than the 6,487 volumes Thomas Jefferson donated in 1815 to kick-start the collection after the British burned the place down. The incoming books are quickly sorted into cardboard boxes by topic. Then the boxes are delivered to three to four hundred catalogers, who between them represent eighty different subject specializations. They examine each book to see which of the library's 285,000 subject headings is most appropriate. Books can be assigned up to ten different subject headings. Keeping America's books nonmiscellaneous is a big job.

So is building and maintaining the subject headings themselves. "We create eight thousand new subject headings per year, with about as many adjustments to existing subject headings," says Barbara Tillett, chief of cataloging policy. The structure of the library's catalog is open-ended enough to allow a subject heading to be created to accommodate books on new topics, rather than insisting on cramming books into the existing headings. The library is also willing to modify the system's larger-scale arrangement, as when the new category of Environmental Sciences was created because, says Tillett, catalogers were bouncing books back and forth: "This one is yours." "No, it's yours."

Catalogers, using their years of experience, are free to propose new subject headings for an up-or-down vote by a committee composed of senior librarians. At one typical weekly meeting, just one out of eighty proposed classifications was turned down—a subclassification of the Photographic Portraits subheading to cover Sri Lankans. The committee decided that that heading's subclassifications were intended to classify people by type—children, physicians, and so on—rather than by country of origin. Such issues may not be glamorous, but they are the type faced daily by those patrolling the conceptual borders of book categorization.

It takes hundreds of professionals with centuries of cumulative expertise to keep the Library of Congress well ordered. But even though the Library of Congress has itself become a standard unit of measurement for large objects—for example, NASA says it maintains information about the

environment that would "fill the Library of Congress 300 times"—it's only dealing with seven thousand new books a day. The Washington Post estimates that seven million pages are added to the Web every day. Search on Google for "American history," which is just one Library of Congress subheading, and you'll get 750 million Web pages—about twenty-six times the number of books in the Library's entire book collection. The Library of Congress's carefully engineered, highly evolved processes for ordering information simply won't work in the new world of digital information. Not only is there too much information moving too rapidly, there are no centralized classification experts in charge of the new digital world we're rapidly creating for ourselves, starting with the World Wide Web but including every connected corporate library, data repository, and media player.

If the Library of Congress's well-proven approach won't work as we digitize our information, ideas, and knowledge, what will?

The Three Orders of Order

Bill Gates bought the Bettmann Archive, the most prestigious collection of historic photos in the United States, so he could bury it. In 2001 he hired nineteen trucks to move it from the melting summers of Manhattan to a cool limestone cave 220 feet underground in the middle of Pennsylvania. There, dehumidifiers the size of closets hold down the moisture level, and security guards patrol brightly lit streets carved out of stone. The site, run by the records management company Iron Mountain, is modern in every way, but if you walk far enough you'll come to a dead end where a hole in the wall reveals an underground lake illuminated only by the light from the opening through which you're looking.

The photographs in the Bettmann Archive are stored in a long narrow cavern whose arched walls of rough rock have been painted white but otherwise left unfinished. In rows of filing cabinets that stretch to the vanishing point are 11 million priceless photographs and negatives. They are arranged by the originating collections the Bettmann purchased over time. Within the collections the photos and negatives are generally ordered chronologically. The room is being slowly lowered to -4° Fahrenheit because Henry Wilhelm, a leading authority on film preservation, believes that at that temperature it will take five hundred years for the collection to deteriorate as much as it did in a single year when it was kept in Manhattan. Wilhelm was inspired to make film preservation his life's work when, as a member of the Peace Corps deep in the Bolivian rain forest, he saw treasured family photographs in thatched houses. "They were deteriorating badly, and there was nothing I could do about it," he says, still sounding frustrated decades later. The Bettmann facility he designed is the polar opposite of a Bolivian rain forest.

As you stand in the long cavern, you are in the midst of a huge first-order organization. In the first order of order, we organize things themselves—we put silverware into drawers, books on shelves, photos into albums. But when you go through the air lock that Wilhelm designed to connect the back chamber and the front one, you confront a prototypical example of the second order of order: a card catalog containing information about each of the eleven million objects in the back cavern. The catalog separates information about the first-order objects from the objects themselves, listing entries alphabetically by subject so that you can find, say, all the photos of soldiers across all of the archive's collections. A code on this second-order object, the catalog card, points to the physical place where the first-order photo is stored in the back room. But quite a few of the Bettmann's photographs are not listed in the card catalog. Some of the older collections arrived with catalogs entered in hand-written ledger books, one line per photo, listed in the order in which the photo was received. Finding a photo in one of those collections requires looking through the ledgers' yellowing pages line by line, hoping to come across a description of the image you're seeking. The ledgers are also a form of second-order classifications, just a much less efficient method than the card catalog.

The Bettmann's second-order organization works, but it's expensive to maintain, and retrieval times are sometimes measured in days. And there are limits inherent in the second order. Not all the information about the objects is recorded; a photograph of a Massachusetts soldier in the Civil War eating in a field, his rifle by his side, might be listed under "Civil War" and "soldier," but probably not also under "Massachusetts," "rifle," "weapons," "uniforms," "dinner," and "outdoors." That means if you were to ask the Bettmann's curators if they had a photo of a Civil War soldier eating outdoors, they would have to send someone into the stacks and stacks of filing cabinets to do a search through the photos themselves. Even if all that data were recorded, it would swell the size of the card catalog to the point of unusability: Searching through eleven million cards at one per second would take over four months of round-the-clock riffling.

All that work—a long line of trucks to move the archive, a hole dug deep into the earth, an ambient temperature growing so cold that you have to don arctic gear to enter the vault—and we still get so little use of those valuable—and expensive—assets. Indeed, a first- and second-order archive the size of the Bettmann literally cannot know everything it has.

The problems with the first two orders of order go back to the fact that they arrange atoms. There are laws about how atoms work. Things made of atoms tend to be unstable over time—paper yellows and disintegrates, negatives turn to soup—so we have to take measures to sway nature from its course. Atoms take up room, so collections of photos can get so large that we have to build card catalogs to remind us of where each photo is. And things made of atoms can be in only one spot at a time, so we have to decide whether a photo of a soldier eating should go into the Civil War folder or the Outdoor Meals folder.

But now we have bits. Content is digitized into bits, and the information about that content consists of bits as well. This is the third order of order and it's hitting us—to use a completely inappropriate metaphor—like a ton of bricks. The third order removes the limitations we've assumed were inevitable in how we organize information.

For example, the digital order ignores the paper order's requirement that labels be smaller than the things they're labeling. An online "catalog card" listing a book for sale can contain—or link to—as much information as the seller wants, including user ratings, the author's biography, and the full text of reviews. You can even let users search for a book by typing in any phrase they remember from it—"What's the title of that detective novel where someone was described as having a face like a fist?"—which is like using the entire contents of the book as a label. That makes no sense when all that information has to be stored as atoms in the physical world but perfect sense when it's available as bits and bytes in the digital realm.

You can see the third order in action by flying across the country from the Bettmann Archive to Seattle, where Corbis, the Bettmann's parent company, has its headquarters. Corbis has charmingly renovated an old bank, knocking down walls to let in light and air, and even retaining the old circular vault door, symbolically open and inviting. Corbis holds over four million digital images, a collection smaller than the Bettmann's but subject to the same issues of organization and control. Because the images as well as the information about the images are all fully electronic, Corbis organizes its photos without regard to the physical constraints that limit the curators back at the Bettmann. At Corbis, you can find a digital image of a Civil War soldier eating dinner by typing "soldier," "Civil War," and "meal" into the search engine, or by browsing a list of categories and subcategories. You can find what you need in seconds. If you don't, you can be pretty sure it just isn't in the Corbis collection.

Of course, it took Corbis many hours of preparation to reduce these search times to seconds. A team of nine full-time catalogers categorize each image Corbis acquires, anticipating how users are going to search. When a new image comes into the collection, one of the catalogers uses special software to browse the 61,000 "preferred terms" in the Corbis thesaurus for those that best describe the content of the image, typically attaching 10 to 30 terms to each one. The system incorporates about 33,000 synonyms (searches on "beach" turn up images labeled as "seaside"), as well as more than 500,000 permutations of names of people, movies, artworks, places, and more. That broadens the side of the barn so wide that if you misspell Katharine Hepburn's name as Katherine or Catherine, you'll still find all the images of the high-cheekboned screen legend in the Corbis collection. And if you're looking for Muammar Gadhafi, at least seventeen different ways of spelling his name will get you what you want.

The differences between the Bettmann's second-order organization and Corbis's third-order method affect every aspect of their businesses.

The Bettmann is an attic that's never been fully explored. It doesn't know all the photos it owns; a ledger entry may be buried among thousands of others, and it may not describe the photo in a way that enables people who want it to find it. At Corbis, every image has been carefully cataloged and can be found by using the company's search engine.

The Bettmann has to be parsimonious with its information: Create too many catalog cards for each photo and your card catalog bloats to unriffability. Like Staples, it bumps against the limits of the physical world. Corbis's approach to information is sprawling and extravagant: The immateriality of bits encourages Corbis to put its images in every place where people might look for them.

Because the Bettmann collection's second-order information—known as metadata because it's information about information—is incomplete and spread across catalogs and ledgers, it's accessible to only a handful of trained experts. Corbis, because its collection is third-order and thus fully digital and cataloged, is designed to be searched by any customer.

Finding photos among the Bettmann's assets can be a slow, manual, expensive process. Corbis, on the other hand, not only can derive benefit from every image but knows so much about them that it can offload much of the job of searching for photos to its users.

Corbis's digital images do not deteriorate with age and require far less physical maintenance. Overall, Corbis spends less per image, can make more per image, and is able to turn more of its images into productive assets.

The differences in the order of order even drive differences in the lighting. The Bettmann's archives are brightly lit because their images are made of atoms that are visible only when light reflects off them. Corbis's catalogers work in semidarkness because digital images on monitors are their own source of light.

But here's the kicker: Like the iTunes Store, Corbis isn't even a particularly good example of third-order organization. It's doing what's right for its business at the moment, but it's still doing the basic second-order task of having professionals stick things into folders. Granted, the things and the folders are electronic, so Corbis can get more value from its assets at a lower cost. But there are other organizations that are able to move further down the third-order path. Corbis gives us only a taste of the revolution that's under way. Just take a look at Flickr to see one way this is unrolling. With over 225 million photos already uploaded by users and almost a million added every day, Flickr's collection dwarfs that of Corbis and Bettmann. Flickr has no professional catalogers. It relies solely on the labels users make up for themselves, without control or guidance. Yet it is remarkably easy to find photos at Flickr on almost any topic and to pull together collections of photos on themes that mix and match those topics at will. Want to find photos of dogs wearing red clown noses? A search at Flickr finds nineteen of them. Researching car-crash art? Flickr finds thirty-two photos that may help your studies.

The digital revolution in organization sweeps beyond how we find odd photos and beyond how we organize our businesses' information assets. In fact, the third-order practices that make a company's existing assets more profitable, increase customer loyalty, and seriously reduce costs are the Trojan horse of the information age. As we all get used to them, third-order practices undermine some of our most deeply ingrained ways of thinking about the world and our knowledge of it.

For example, medical information that used to come only through the careful filters of medical experts and medical publications is now available to everyone prior to the basic housekeeping processes of being gone through and put away. The miscellanizing of this information not only breaks it out of its traditional organizational categories but also removes the implicit authority granted by being published in the paper world. Second-order organization, it turns out, is often as much about authority as about making things easier to find.

We have entire industries and institutions built on the fact that the paper order severely limits how things can be organized. Museums, educational curricula, newspapers, the travel industry, and television schedules are all based on the assumption that in the second-order world, we need experts to go through information, ideas, and knowledge and put them neatly away.

But now we—the customers, the employees, anyone—can route around the second order. We can confront the miscellaneous directly in all its unfulfilled glory. We can do it ourselves and, more significantly, we can do it together, figuring out the arrangements that make sense for us now and the new arrangements that make sense a minute later. Not only can we find what we need faster, but traditional authorities cannot maintain themselves by insisting that we have to go to them. The miscellaneous order is not transforming only business. It is changing how we think the world itself is organized and—perhaps more important—who we think has the authority to tell us so.

Copyright © 2007 by David Weinberger. All rights reserved.