CONTENT

KATE EICHHORN

The MIT Press | Cambridge, Massachusetts | London, England

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Chaparral Pro by New Best-set Typesetters Ltd.

Library of Congress Cataloging-in-Publication Data

Names: Eichhorn, Kate, 1971- author.
Title: Content / Kate Eichhorn.
Description: Cambridge : The MIT Press, 2022. | Series: The MIT press essential knowledge series | Includes bibliographical references and index.
Identifiers: LCCN 2021000773 | ISBN 9780262543286 (paperback)
Subjects: LCSH: Information services industry. | Online social networks. | Search engines. | Internet—Social aspects. | Internet—Psychological aspects. | Information technology—Psychological aspects.
Classification: LCC HD9999.I492 E33 2022 | DDC 302.2/402854678—dc23
LC record available at https://lccn.loc.gov/2021000773

10 9 8 7 6 5 4 3 2 1

USER-GENERATED CONTENT

In 2005, most writers were boycotting Amazon. At the time, the online retailer was viewed as an imminent threat to independent bookshops, though not yet as a threat to retailing as we knew it. But one writer, Kevin Killian, was secretly leaving reviews on the loathed platform. By the time of his death in 2019, Killian had published over 2,500 reviews of books, films, and miscellaneous household products on Amazon—enough reviews to earn him a spot in Amazon's "Reviewer Hall of Fame." In interviews, Killian explained that he started posting Amazon reviews after an illness left him unable to write anything of length or substance. Over time, his reviews, which were eventually compiled and published as an open-source book, would serve as more than writing therapy. In Killian's words, "I have often used Amazon reviews as a springboard to doing

other kinds of writing projects. So when you read them, yeah, they're reviews of a sort, but they also seem like novels. They're poems. They're essays about life. I adopt a different persona in them. . . . So yeah, I get a lot of my kinks out there, on Amazon."[1] But Killian's Amazon reviews—a conceptual writing project that spanned more than a decade—also served another purpose. Because his reviews, which represent a type of "user-generated content," always served his own needs and not simply Amazon's, they ultimately served as an ongoing reminder of user-generated content's original promise—to collapse the line between consumers and producers.

In the 2020s, readers may have difficultly appreciating that user-generated content once appeared ripe with creative, political, and even subversive possibilities. But a time existed when optimism about user-generated content abounded. It promised to revitalize the public sphere, lay the groundwork for the rise of citizen journalism, bring marginalized voices to the center, and enable writers, musicians, filmmakers, and artists without access to publishing houses, record labels, big-time production money, or gallery or museum space to reach increasingly broad audiences. Some of these things turned out to be true. Letters to the editor are no longer the only or primary way to publicly participate in political debates as regular folks. Citizen journalism has expanded around the globe. And,

for better or worse, platforms for user-generated content like Twitter, Instagram, and YouTube have created a massive audience for all sorts of rising stars, whether or not they exhibit any notable talents.

However, user-generated content—a concept that has at times been used contiguously with overlapping concepts ranging from "convergence culture" to "participatory culture" to "peer production"—wasn't just about regular people widely circulating their ideas, words, images, and performances. What most people didn't realize in the 1990s or even early 2000s is how, where, and to what ends the content they were now freely sharing online would help generate revenue for private companies and, in many cases, lay the foundation for the development of new technologies, including some that would eventually be turned against users. Concluding that the history of user-generated content is simply a history of deception would be unfair. The creators of some early user-generated platforms weren't fully aware of the potential profits they might reap from the data assets they were accumulating. Also, even well into the digital era, not all user-generated initiatives are driven by greed or big tech's desire to drive its own research and development agendas. For this reason, the history of user-generated content may be best understood as a history of conflicting desires and agendas.

The Long History of User-Generated Content

Although user-generated content may be a relatively new phenomenon, users did contribute to the creation or development of products before 2000. In fact, there is a long history of users making contributions to the production of texts and other products to which they never held copyright or ownership.

In the early years of movable type, for example, books were frequently printed dozens of times, with subsequent editions incorporating changes and corrections. In some cases, the changes and corrections were driven by the author, but other early printed books were revised and republished due to additions and revisions proposed by readers. Sebastian Munster's *Cosmography*, for example, first appeared in 1544 and went through eight editions in Munster's lifetime and another thirty-five editions by 1628. The revisions were largely the result of additions and corrections offered by his readers. In her seminal work on the history of the printing press, Elizabeth Eisenstein describes the culture of the early print world as one where editors and publishers "did not merely store data passively in compendia" but also "created vast networks of correspondents" by soliciting criticism of each edition and "sometimes publicly promising to mention the names of readers who sent in new information or who spotted the errors which would be weeded out."[2] In this respect, early

printers and readers, according to Eisenstein, were engaged in endeavors that resonate with the crowd-sourced and user-generated projects that have come to define contemporary communications.

As the popularity of encyclopedias and compendiums declined, readers continued to contribute to the research and development of at least some types of texts. In 1857, the Unregistered Words Committee of the Philological Society of London issued a circular calling for volunteers to read specific books and copy out quotations that offered examples of "unregistered" words and meanings (i.e., items not yet recorded in other dictionaries). The original proposal wasn't intended to develop a complete dictionary, but as volunteers continued to respond to the Philological Society's call for assistance, a decision was made to create a new English dictionary. This user-generated experiment eventually led to the publication of the first *Oxford English Dictionary*.[3] Beyond dictionaries, user-generated content has also long been a mainstay in the travel guide industry. In the twenty-first century, TripAdvisor relies entirely on user-generated content. The platform, which was built with just $3 million in investments but was worth an estimated $7 billion by 2016, has transformed how people make decisions as they travel around the world.[4] Yet, as early as the late nineteenth century, travel guides were already being developed with the help of travelers. As John Muirhead, the English-language editor of the *Baedeker*,

observed in 1889, "A guidebook is not made, it grows." Muirhead further explained, "When a new edition is being prepared, the first thing we do is to go carefully through the mass of correspondence, generally very voluminous, which has come to hand. This consists of hotel bills, notes, complaints, and suggestions."[5]

From encyclopedias to dictionaries to travel guides, there is a long history of readers contributing to the research and development of texts. But since the 1990s, three things have radically transformed how regular folks contribute to the production of texts and images of all kinds: an expanded capacity to engage in the production of audio and visual content; an expanded capacity to broadcast these creations; and most importantly, an expanded capacity for private companies to turn such creations into assets.

The Ability to Generate Content across Media

In a print culture, user-generated content was generally restricted to one type of content—text. In a digital culture, users can now share more than words. Since the launch of digital photography and mobile devices, user-generated content is now just as likely to take the form of photographs, videos, and sound recordings as it is to take the form of text. With user-generated content now being produced across media, there is more user-generated content and, more importantly, such content is no longer

restricted to wordsmiths. Today, anyone, even very young children or people with limited literacy skills, can easily contribute to the content pool.

The Ability of Users to Publish/Broadcast Content

In the past, users could submit content—whether suggesting words to the editors of the *Oxford English Dictionary* or travel advice to the editors of the *Baedeker* guides—but in a print culture, such users were still entirely dependent on the editors of these volumes to put their ideas into print. In a digital era, sites may be moderated, but users can generally make a small edit to a Wikipedia page, submit a review on Yelp, or post a dance performance on YouTube without waiting for anyone's approval.

Capacity to Capture, Manage, and Profit from User-Generated Content

In the twenty-first century, user-generated content can be easily captured, managed, and transformed into an asset. To begin, user-generated content can now be captured across media twenty-four hours a day, seven days a week; and in some cases, users are not even aware that their content is being captured. Second, the ability to manage large data sets has expanded. No longer dependent on editors manually sorting through the mail, data can now be automatically collected and mined. As such, volume is no longer a concern, and this has led to another important shift.

Because we now have the capacity to capture, collect, and mine increasingly large sets of data, we can deploy user-generated content to achieve entirely different ends. Facebook's users may upload photographs in order to share memories with friends and family, but these photographs are valuable to Facebook for an entirely different reason. With millions of tagged photographs, Facebook can support the development of facial recognition technologies, among other things. Simply put, brought to scale, user-generated content exceeds its original purpose and, in the process, becomes increasingly valuable as an asset.

So, what's changed since the days of crowd-sourced encyclopedia, dictionary, and travel guide publishing? Back in the sixteenth century, if you contributed an update to Munster's *Cosmography*, your update likely remained your update—a correction or new factoid that would simply help make *Cosmography* a more accurate and relevant resource. Today, some user-generated content—for example, a correction or addition to a Wikipedia page—maintains its original purpose, but a lot of other user-generated data is collected under one pretense and then used for different purposes. User-generated content is now more likely to also circulate as, to use Lyotard's term, "investment knowledge"—knowledge exchanged to optimize the performance of an entirely separate endeavor. This may sound problematic, and in many respects it is, but this

doesn't mean user-generated content has always been or currently is entirely at odds with user needs.

The Promise of User-Generated Content

To appreciate the original promise of user-generated content, consider some of the other ways in which it has been described over time. While *user-generated content* has always been the favored term in a corporate context, alternative terms such as *convergence culture*, *participatory culture*, and *peer production* have often been favored by scholars, cultural workers, and digital activists. Whether or not you opt to embrace these terms, they matter. Above all else, they remind us that before user-generated content was viewed as merely a part of the endless data stream required to fuel tech companies, it was envisioned as something that held the potential to transform culture, media, and production endeavors of all kinds.

Consider the terms *convergence culture* and *participatory culture*. Media theorist Henry Jenkins is generally credited with coining the former term and popularizing the latter. Early on, Jenkins recognized that we live in a world where content can and will be increasingly monetized by private corporations. As he observed in his 2006 book *Convergence Culture*, "corporations—and even individuals within corporate media—still exert greater

power than any individual consumer or even the aggregate of consumers." Still, in retrospect, Jenkins's use of the concepts of *convergence culture* and *participatory culture* reflected a certain optimism about the world users might build online. "In a world of media convergence," Jenkins observed, "every important story gets told, every band gets sold" (cynically, he also noted that "every consumer gets courted across multiple media platforms"). He further insisted that in this new online world, "Rather than talking about media producers and consumers as occupying separate roles, we might now see them as participants who interact with each other according to a new set of rules that none of us fully understands."[6] Jenkins's position on convergence culture and participatory culture was never naïve, but it also wasn't dripping with cynicism. "Convergence culture is where old and new media collide," he concluded. It is "where grassroots and corporate media intersect, where the power of the media producer and the media consumer interact in unpredictable ways."[7] Ultimately, Jenkins imagined a world where consumers might be increasingly repositioned as producers and once-marginalized voices might finally have at least a chance of taking center stage.

Around the time that Jenkins was writing about convergence culture and participatory culture, legal scholar Yochai Benkler was sharing another vision for this new world order. Using the already-established practice of open-source

software as a central example, Benkler envisioned a world where problems might be tackled with increased efficiency. In his 2002 *Yale Law Journal* article "Coase's Penguin, or, Linux and *The Nature of the Firm*," he accurately predicted that we were about to enter a "third mode of production in the digitally networked environment." Before YouTube, Facebook, or Flickr had even launched, Benkler foresaw the arrival of an era of "commons-based peer production." He used the term *peer production* to distinguish this new mode of production from earlier modes of property- and contract-based production. With peer production, Benkler observed, "groups of individuals successfully collaborate on large-scale projects following a diverse cluster of motivational drives and social signals, rather than either market prices or managerial commands."[8]

While Benkler, like Jenkins, never assumed that this emerging mode of peer production would be immune to cooptation, in 2002, his optimism seeps through. Among other things, Benkler believed that commons-based peer production would take the guesswork out of sourcing talent. After all, as he argued, "peer production has a systematic advantage over markets and firms in matching the best available human capital to the best available information inputs to create information products." That is, if in the past one had to rely on serendipity to bring together a dream team of designers or engineers to solve a problem, in the emerging commons-based peer production

world, serendipity would no longer be necessary for two key reasons. First, peer production increases the chances that the best person for any given job will actually end up doing the job, because jobs that were previously closed (e.g., limited to employees with specific credentials) will be opened up to the broader public. Second, peer production removes traditional obstacles posed by property and contracts by decreasing the cost of "allowing larger clusters of potential contributors to interact with large clusters of information resources in search of new projects and opportunities for collaboration."[9] Put simply, in the world of commons-based peer production, problems are presented and solutions are crowd-sourced by the best minds around the world without all of the hassle of old-world models with their credential vetting, contracts, and expectation that collaborators will necessarily assemble in the same room and work face to face for extended periods of time.

Jenkins's and Benkler's early writings on user-generated content are by no means driven by an identical agenda, but they had at least one thing in common. Both theorists agreed that with the shift to participatory culture or peer production, more people—regardless of their identity, credentials, or locations—will be able to come together to engage in cultural production, debate issues, and solve problems. Of course, by the time Jenkins and Benkler were weighing in on this subject in scholarly publications, digital activists had been making similar

arguments for well over a decade, albeit not necessarily in the form of peer-reviewed articles or university press books.

One of the earliest and most well-known groups to express the arguments later popularized by theorists such as Jenkins and Benkler was the WELL (Whole Earth 'Lectronic Link)—a freewheeling virtual community started by the same commune-supplying hippies responsible for the *Whole Earth Catalog* in the early 1970s. The WELL has always recognized that "users" (what the WELL would describe as "members") aren't just people who generate stuff to occupy virtual spaces. What members produce isn't content per se but rather a virtual form of bricks and mortar. Kevin Kelly, a cofounder of the WELL who later served as an executive editor at *Wired*, has said that when the WELL started, it was driven by seven design goals, including a commitment to making the WELL a "self-designing experiment"—that is, a place where "early users were to design the system for later users."[10] In this respect, from its inception, the WELL understood user-generated content as raw material laying the foundation for new forms of community.

Precisely because content generated by users was once viewed as both the stuff required to build new containers and the stuff that might occupy these new containers, many cyber optimists believed early on that users might eventually rebuild the world online and, in the process,

overcome the entrenched inequities and access barriers of real life. On February 8, 1996, when many people were still figuring out how to connect their home modem to their desktop computer, John Perry Barlow—a cattle rancher, early member of the WELL, and founding member of the Electronic Frontier Foundation (EFF)—published his now-infamous cyber manifesto, "A Declaration of the Independence of Cyberspace." The manifesto wasn't just a call for governments around the world to back off and leave cyberspace alone. It reflected the utopian ethos prevalent at the time: the belief that users were building a new and better world online. From its opening statement—"Governments of the Industrial World . . . On behalf of the future, I ask you of the past to leave us alone. You are not welcome among us. You have no sovereignty where we gather"—to its conclusion—"We will create a civilization of the Mind in Cyberspace. May it be more humane and fair than the world your governments have made before"—Barlow meant business.[11] He wasn't alone. The manifesto and the work he would ultimately carry out through the EFF remain deeply rooted in a belief that online, users can change the world by reimagining its parameters, altering who is invited to engage in its design, and most important, by wrestling its control out of the hands of governments and handing it back to the people.

When looking back on the early days of cyberspace and writings on convergence culture, participatory culture,

peer production, and the electronic frontier, one finds a lot of hope about the potential ways in which users might build new cultures, systems of exchange, and political structures. Also painfully apparent is the fact that not everyone was participating in these utopian conversations. Certainly, user-generated or participatory media may have once appeared ripe with possibilities, but most of the people writing about these possibilities (and promoting them) were still white men living in highly industrialized nations—a point that visionaries like Barlow never fully appreciated. By 2000, media theorist Lisa Nakamura could say with considerable certainty that "the Internet 'revolution' is over," and worse yet—as had already happened during most other so-called revolutions—people of color were left on the margins. As Nakamura observed at the time, even as terms such as *digital divide* started trending, people of color were functionally absent from the internet.[12] Yet, as Nakamura also contends, even with people of color occupying a marginal role online, the internet was already a place where racial identities, divides, and inequalities were being produced and reproduced.

Since Nakamura's *Cybertypes: Race, Ethnicity, and Identity on the Internet* first appeared, similar arguments have been made by other theorists about gender, ability, and geographic location. At least early on—back when a much higher percentage of online users were still building virtual infrastructure and doing so at arm's length from big

tech, in alternative communities like the WELL—users may have been dreaming big, but those users certainly didn't represent everyone who could have been participating in this world-building endeavor. To be clear, virtual communities would eventually become far more diverse. The Pew Research Center found that in 2000, 38 percent of Black Americans reported using the internet compared to 53 percent of white Americans. By 2018, the gap had closed.[13] What changed during that time is important. In 2000, more spaces online were still transaction-free, and a higher percentage of users were still actively engaged in building online communities from the ground up rather than simply adding content to sites and platforms owned by outside entities. As access expanded, virtual communities became increasingly diverse, but unfortunately this expansion coincided with the decline of most early utopian virtual communities and the rise of transactional communities (e.g., those supported by established social media platforms).

When User-Generated Content Became an Asset

In the mid-1990s, many people were optimistic about the potential ways that users might rebuild the world online, but they weren't all cyber enthusiasts. By that time, user-generated content was also attracting attention from

another demographic—big business. With digital business still in a nascent stage, business analysts were even actively looking at underground communities, including the WELL, to figure out how to start profiting from the world's growing digital outputs.

In "Real Profits from Virtual Communities," an article published in a 1995 issue of the *McKinsey Quarterly*, a team of McKinsey & Company analysts observed, "Electronic communities have actually existed for many years. Groups like The Well have spawned strong relationships and developed their own norms and sense of history." The authors of the article recognized that these communities are "noncommercial"—that is, "Communication, entertainment, and information are their reasons for being, and contracts tend to be in barter form. Financial transactions are rare, and often resisted." Still, they were also confident that "new kinds of community that are more commercially focused will emerge soon."[14]

The consultants at McKinsey were right. Shortly after the publication of "Real Profits from Virtual Communities," online communities started to become increasingly transactional. As the McKinsey analysts predicted, while some online community members did remain resistant to turning their communities into transactional environments (to this day, the WELL is still a nonprofit endeavor), as more people went online, consumers—rather than visionaries—increasingly started to shape the online

world. But this isn't the only thing McKinsey's analysts accurately predicted in 1995. They also foresaw the potential for virtual communities to become valuable assets in a unique new way. Without entirely discounting the possibility that profits might still be derived from eventually making virtual communities subscription-based, they predicted that the real profits would ultimately come from the "unique content" virtual community members would generate. The shift from old-world approaches to profit (i.e., subscriptions) to new-world approaches (i.e., turning user-generated content into an asset that can be collected, mined, and sold) wouldn't be simple. Among other challenges, McKinsey's analysts cautioned, "Aspiring community organizers will have to decide which communities they will try to own. Their decision will be based partly on the assets that they already possess, such as brands, content, special skills, and relationships with other communities."[15] Many other questions remained unanswered in 1995: How will we structure ownership of user-generated content? Must these virtual communities be heavily moderated to ensure only the right type of content is generated by users? Who will be charged with the moderation of these communities? What types of skills will these moderators need and where will they receive their training? Yet what was clear as early as 1995 was the potential to turn virtual communities—their participants, infrastructure, and especially their unique content—into assets.

In retrospect, what is most interesting about this period of internet history is the notable disconnect between what business analysts and businesses were plotting and what most online users understood about these future plans. After all, in the 1990s and even well into the early 2000s, while business analysts and businesses were actively exploring how to turn users' comments and, eventually, users' digital photographs, videos, and sound files into assets, most online users remained largely in the dark about the ways in which their digital output might be monetized. This is not to suggest that early online users were completely oblivious. Most people anticipated that the internet would become more commercialized. In the 1990s, many people were already worrying that free sites might eventually be accessible only with a paid subscription and that websites would eventually be cluttered with advertisements and, as a result, start to look a lot more like the average newspaper or magazine. What most people wandering through the "electronic frontier" for the first time didn't realize in 1995 or even 2000 was that the comments they were freely sharing online were already being viewed as a potential asset by a new class of entrepreneurs. This lack of awareness is presumably why few users were concerned about freely sharing information, even private information, on social media platforms such as Facebook and photo- and video-sharing sites such as YouTube and Flickr when they began to appear in the early

2000s. While a real concern existed that these newfangled forms of entertainment and communication might eventually require a high subscription fee or be colonized by advertisements, few people were worrying that these sites might eventually transform everything—even conversations with friends or exchanges of personal photographs of birthdays and bat mitzvahs—into assets. But this is precisely what happened: private companies learned to exploit users' online interactions and creativity (i.e., their content).

As Christian Fuchs argued in his 2013 essay "Class and Exploitation on the Internet," user-generated data is best understood as a commodity that is partially produced by users and partially produced by the corporations that build and maintain the platforms adopted by users. However, as Fuchs further commented, one notable difference exists between these two players—"users are unpaid and therefore infinitely exploited."[16] Being unpaid doesn't mean that these users aren't, as scholars such as Jenkins have long emphasized, also active participants. As Fuchs suggests, "On Facebook, Twitter, and blogs, users are fairly active and creative . . . but this active character is the very source of exploitation."[17] Ultimately, the ability of private companies to exploit users' online interactions and creativity (i.e., their content) would drive the success of the digital economy.

The Classification of User-Generated Content

Three decades after the arrival of the web, user-generated content has become ubiquitous. In the process, it has transformed nearly every sector imaginable from entertainment and education to health and finance. User-generated content has also restructured our everyday practices—it has altered how and what we cook, how we learn, where we travel, and even who we date. Still, disagreements about user-generated content abound. Optimists argue that such content has brought new efficiencies into our lives (e.g., user-generated content can be collected and mined to strengthen the supply chain and ensure the right volume of products are available in the right place at the right time). Pessimists say that user-generated content has eroded our privacy and, worse yet, turned us all into digital laborers who are expected to toil away—without compensation—twenty-four hours a day. Despite the fact that user-generated content—love it or loathe it—has transformed and continues to transform how we work, play, and live, the classification of user-generated content remains a challenge.

On Wikipedia, which is just one example of user-generated content, the entry on "user-generated content" (at least the entry that existed when I wrote this chapter) lists eight "types" of user-generated content: blogs, websites, video games, advertising, retailers, educational, photo

sharing, and short video sharing. What's striking about the entry is that the types of user-generated content listed don't appear to reflect any classificatory logic at all. A blog is arguably a specific genre of online writing; a website is a collection of web pages; a video game is a type of game; and advertising, photo sharing, and video sharing are all practices that can take place either online or off. Simply put, the Wiki entry attempts but fails to classify different types of user-generated content because it endeavors to compare genres to collections to practices. Listing blogs, websites, and video games under the same broad category is a bit like listing novels, archives, and games in the same category. As a classification, it doesn't make much sense. Yet, once one starts to contemplate how to update the Wiki entry on user-generated content, classifying such content becomes far more challenging.

To begin, one might approach the update by classifying user-generated content in relation to established media categories (e.g., text, image, sound). At the very least, this approach would reveal how user-generated content operates across media. But since all digital user-generated content is streamable, old-world divisions between text, image, and sound ultimately reveal little about the differences between different types of user-generated content. Said another way, while the differences between a book, a printed photograph, and a vinyl record may be significant, the differences between a blog post, a digital photograph,

and an MP3 are less clear. So, what if one focused instead on the function of different types of user-generated content?

Classifying user-generated content by function might draw attention to content whose primary function is to review products or services (comments left on a platforms like Yelp or Amazon); share opinions (reader comments on the *New York Times* or Fox News site); entertain (videos on YouTube or Vimeo); offer feedback on the user's health, habits, or lifestyle (Fitbit); connect with family and friends (Facebook or Instagram); and so on. This classification might be helpful, but it would remain entirely user-centric, and as previously noted, user-generated content is often employed in ways entirely unknown to the people who generated the content. A more difficult but by no means less important endeavor, then, would be to classify user-generated content in relation to how it is used by the owners of digital platforms that rely on user-generated content. Such a classification system might separate user-generated content produced for the pharmaceutical industry (the user-generated data about monthly periods left on Clue) from user-generated content produced to drive facial recognition technologies (photographs posted on Facebook and Instagram). But again, this classification would still fail to capture the fact that the content in question isn't just generated or used for one purpose or by one entity. The user-generated content shared by users on the

family-history site Ancestry.com, which includes genetic data, not only supports the platform's own research and development (helping the company to improve its platform) but also is used to support research in other fields, including the pharmaceutical industry.

As data is posted, collected, combined, mined, and traded, both its original medium and message cease to matter. As such, any attempt to classify user-generated content based on medium or message is bound to run into problems. Rather than attempting to classify types of user-generated content, it may be more valuable to focus on tracing its life cycle. Ultimately, what distinguishes digital user-generated content from early forms of user-generated content and other types of digital content (e.g., content produced by the owners of platforms) is its capacity to morph over time—to transform from a communicative act originating from a single user to one small bit of data in a larger database, and then to a form of investment knowledge that exists to optimize services, products, and schemes that the original user may have never imagined possible.

NOTES

Preface

1. Jeff Desjardins, "How Much Data Is Created Each Day?," World Economic Forum, April 17, 2019, https://www.weforum.org/agenda/2019/04/how-much-data-is-generated-each-day-cf4bddf29f/.
2. Ella Koeze and Nathaniel Popper, "The Virus Changed the Way We Internet," *New York Times*, April 7, 2020, https://www.nytimes.com/interactive/2020/04/07/technology/coronavirus-internet-use.html.
3. Quoted in Jim Turner, "'Starved for Content': Ron DeSantis Explains Why Sporting Events Are Needed Now," *Tampa Bay Times*, April 14, 2020, https://www.tampabay.com/news/health/2020/04/14/in-defense-of-wwe-decision-ron-desantis-said-he-wants-more-sporting-events/.
4. Among other recent studies, see Jean M. Twenge and W. Keith Campbell, "Associations between Screen Time and Lower Psychological Well-Being among Children and Adolescents: Evidence from a Population-Based Study," *Preventive Medicine Reports* 12 (October 18, 2018): 271–283, doi:10.1016/j.pmedr.2018.10.003.

Chapter 1

1. Jodi Dean, *Democracy and Other Neoliberal Fantasies: Communicative Capitalism and Left Politics* (Durham: Duke University Press, 2009), 26.
2. Bharat Anand, *The Content Trap* (New York: Random House, 2016).
3. Theodor W. Adorno and Max Horkheimer, *Dialectic of Enlightenment*, trans. John Cumming (New York: Verso, 1997), 136.
4. William Safire, "The Summer of This Content," *New York Times*, August 9, 1998, section 6, p. 18.
5. John Unsworth, editor's introduction to *Postmodern Culture* 4 (2) (1994), http://pmc.iath.virginia.edu/text-only/issue.194/intro.194.
6. Joanna Brenner, "3% of Americans Use Dial-Up at Home," PEW Research Center, August 21, 2013, https://www.pewresearch.org/fact-tank/2013/08/21/3-of-americans-use-dial-up-at-home/.
7. Adorno and Horkheimer, *Dialectic of Enlightenment*, 136.
8. Jean-François Lyotard, *The Postmodern Condition: A Report on Knowledge* (Minneapolis: University of Minnesota Press, 1979), 4–5.
9. Ibid., 5–6.

10. McKenzie Wark, *A Hacker Manifesto* (Cambridge, MA: Harvard University Press, 2006), 57–58, 26.
11. Michael Bhaskar, *The Content Machine: Towards a Theory of Publishing from the Printing Press to the Internet* (London: Anthem Press, 2013), 6.

Chapter 2

1. Kevin Killian, interview with Ruby Brunton, "Developing a Creative Process," *The Creative Independent*, March 4, 2019, https://thecreativeindependent.com/people/writer-kevin-killian-on-being-unlikeable-in-your-work/.
2. Elizabeth L. Eisenstein, *The Printing Press as an Agent of Change* (Cambridge: University of Cambridge Press, 1979), 109.
3. Oxford English Dictionary, "'Your Dictionary Needs You': A Brief History of the OED Appeals to the Public," *OED* blog, October 4, 2012, https://public.oed.com/blog/dictionary-needs-brief-history-oeds-appeals-public/.
4. Linda Kinstler, "How TripAdvisor Changed Travel," *Guardian*, August 17, 2018, https://www.theguardian.com/news/2018/aug/17/how-tripadvisor-changed-travel.
5. As cited in Ana Alacovska, "The History of Participatory Practices: Rethinking Media Genres in the History of User-Generated Content in Nineteenth-Century Travel Guidebooks," *Media, Culture and Society* 39, no. 5 (2017): 662–663.
6. Henry Jenkins, *Convergence Culture* (New York: New York University Press, 2006), 3.
7. Ibid., 259–260.
8. Yochai Benkler, "Coase's Penguin, or, Linux and *The Nature of the Firm*," *Yale Law Journal* 112, no. 3 (December 2002): 369.
9. Ibid.
10. As cited in Howard Rheingold, *The Virtual Community: Homesteading on the Electronic Frontier*, 2nd ed. (Cambridge, MA: MIT Press, 2000), 31.
11. John Perry Barlow, "A Declaration of the Independence of Cyberspace," Electronic Frontier Foundation, February 8, 1996, https://www.eff.org/cyberspace-independence.
12. Lisa Nakamura, *Cybertypes: Race, Ethnicity, and Identity on the Internet* (New York: Routledge, 2002), xii.
13. Pew Research Center, "Internet Use by Race and Ethnicity," January 11, 2017, https://www.pewresearch.org/internet/chart/internet-use-by-race.
14. Arthur Armstrong and John Hagel III, "Real Profits from Virtual Communities," *McKinsey Quarterly* 3 (1995): 2.
15. Ibid., 7.

16. Christian Fuchs, "Class and Exploitation on the Internet," in *Digital Labor*, ed. Trebor Scholz (New York: Routledge, 2013), 218.
17. Ibid., 217.

Chapter 3

1. Peter H. Lewis, "Personal Computers; Cruising the Web with a Browser," *New York Times*, February 7, 1995, C0008.
2. David Rothenberg, "How the Web Destroys the Quality of Students' Research Papers," *Chronicle of Higher Education*, August 15, 1997, https://www.chronicle.com/article/How-the-Web-Destroys-the/75214.
3. Susan Colaric and Alison A. Carr-Chellman, "Speeding on the Information Superhighway," *TechTrends* 44, no. 6 (November 2000): 28.
4. Howard Rheingold, *The Virtual Community: Homesteading on the Electronic Frontier*, 2nd ed. (Cambridge, MA: MIT Press, 2000), xxiii.
5. Robert H. Reid, *Architects of the Web: 1,000 Days That Built the Future of Business* (New York: John Wiley & Sons, 1997), 300–308.
6. In 2001, Chip Bayers reported in *Wired*, "With AdWords, an advertiser can purchase one of the small, rectangular boxes that appear to the right of the search results on selected pages for an initial fee that ranges from a $15 CPM for the top spot to $8 for the bottom four in the list of eight." Chip Bayers, "I'm Feeling Lucky," *Wired*, October 1, 2001, https://www.wired.com/2001/10/google-6/.
7. On June 18, 2003, *Wired* reported that Google was rolling out a new "self-serve program called AdSense, aimed to help operators of less-trafficked sites quickly format their Web pages to receive advertisements that match key words in the content of those individual pages." "Wired News Report," *Wired*, June 18, 2003, https://www.wired.com/2003/06/oracle-responds-with-mo-money/.
8. Daniel Roth, "The Answer Factory: Demand Media and the Fast, Disposable, and Profitable as Hell Media Model," *Wired*, October 19, 2009, https://www.wired.com/2009/10/ff-demandmedia/.
9. All the aforementioned job postings were found on Upwork.com on December 28, 2018.
10. Richard Florida, *The Rise of the Creative Class, Revisited* (New York: Basic Books, 2012).
11. Angela McRobbie, *Be Creative: Making a Living in the New Culture Industries* (Cambridge, UK: Polity Press, 2016).
12. While working on this book, I was unable to find any quantitative studies to support something I have found to be consistently true anecdotally—that