7.

How do business models shape online communities?

You can't understand technology in isolation--it is inseparable from the social systems that create it, and in which it is used. Researchers in human-computer interaction (HCI) study technology embedded in the world, thinking about "socio-technical systems." Dix et. al. write that "Socio-technical models are concerned with technical, social, organizational, and human aspects of design. They recognize the fact that technology is not developed in isolation but part of a wider organizational environment. It is therefore important to consider social and technical issues side by side" (Dix et al. 1998). Financial factors often don't receive enough attention as a component of the socio-technical system. In this chapter, I will argue that markets are a critical component of that broader social environment. *How things are financed changes what they become.*

The original computer network that later became the internet, the ARPANET, was created as a research project funded by the United States department of defense (Clark 2018). Today the internet (from the physical layer up to the application layer) is partly funded by governments and partly by companies. A number of smaller online platforms (like Mattermost and Mastodon) are non-profit, but most are run by corporations and supported either by advertising or subscription fees. How online sites are funded profoundly shapes them.

In this chapter, I will focus on three fundamental ways in which financial factors shape content moderation and online interaction. First, if there is a cost of joining a site or participating in a site, this shapes who can afford it, and who chooses to afford it. Second, customer service is expensive. How we manage bad behavior is fundamentally shaped by how much it costs. Third, how sites make money shape their priorities, design decisions, and policies.

Markets as a Regulator

As we saw in the last chapter, Larry Lessig describes four things that regulate behavior: laws, social norms, technology, and markets. To explain the impact of markets on behavior, Lessig uses the example of smoking. Raising cigarette taxes lowers the number of people who smoke. Markets are a classic way to regulate behavior, and it works for online behavior just like face-to-face behaviors.

Today people are accustomed to most online sites being free. On supposedly "free" sites, users are often paying by allowing their personal information to be mined, and by accepting targeted advertising. When there is a fee to participate, even a token fee can dramatically reshape user behavior. One of the first community networks, Berkeley Community Memory, leveraged this technique with great success. Faced with a deluge of low-quality posts, they began charging a small fee—just 25 cents—for each post. As a result, the quality of content on the site improved (Shuler 1994). In earlier chapters we discussed peer review as a powerful force for ensuring quality of content. This small fee encouraged self-review, which is also powerful.

Small fees can have big impacts. The irreverent online discussion site the Something Awful Forums (SAF, https://www.somethingawful.com/) prides itself on its high-quality discussion. New users are represented by an icon that says "stupid newbie" unless they pay a one-time small fee (formerly $5, now $9.95) for a paid account. This simple filter helps raise the level of conversation (Pater et al.).

One nice side effect of the small fee to join SAF is that it makes people think twice about bad behavior. On most online sites, if your account is banned, you can just create a new one. Some subreddits try to counter that by blocking submissions from accounts that were created too recently, or have low karma (points you get for upvotes on your posts and comments). On SAF, the fee to join is an elegant deterrent. If you get yourself banned, it's going to cost you $9.95 to come back!

A larger fee of course has a bigger impact. In the face-to-face world, many social clubs use high or extremely high fees to filter potential members. Clubs with high fees become places for elites to socialize. Such clubs play a role in reproducing inequality—when members of elites network, they inevitably help one another, which reinforces privilege. Not surprisingly, such elite clubs are being created online. For example, the site Rich Kids (https://richkids.life) charges €1000 per month for membership, and promises that tabloids will see your photos. The site Whispers is only for owners of new Rolls Royce cars (Mihalascu 2020). I'm sure there are many sites with expensive barriers to entry that are not visible to the public. An elite country club has a physical footprint that passersby can see, but digital ones can be invisible to outsiders.

Market forces like subscription fees may be key to managing bad content on the internet in the future. Over the last few years, there has been a growing recognition that much content on the internet is problematic. In his science fiction book *Fall; or, Dodge in Hell* (Stephenson 2019), author Neal Stephenson imagines a world in which individuals hire editors to manage content they see. Poor people use automated programs to filter their content, and thousands of people may use the same filter. Rich people, on the other hand, hire personal human editors who know their likes and interests and tailor an information feed for them personally. The result is that people who are economically advantaged get access to better information. In Stephenson's dystopian vision, reliable information is a privilege of the rich, and the less advantaged often sink into bizarre conspiracy and cult beliefs encouraged by a steady flow of unreliable information. With the current rise in belief in false conspiracies like QAnon, chemtrails, and a flat earth, Stephenson's prediction is increasingly plausible.

During the Middle Ages through much of the Renaissance, information was indeed the privilege of the rich. Wealthy merchants hired staff at great expense to go to court and report back on what was taking place. The advent of the first newspapers in the 17th century began a gradual process of making news available to a broader segment of society (Pettegree 2014). Wikipedia is the high point of this trend, making the world's knowledge more accessible than ever before.

Some of the first newspapers were sensationalist tabloids, created with the primary goal of making money and only a secondary goal of informing or enlightening. Information that is commercially available has long been of mixed

quality (Pettegree 2014). The problem of how to encourage the spread of good content and discourage the spread of bad is long-standing. At the time I am writing this, it feels like the problem is becoming more acute.

Market-based regulation mechanisms are likely to play a significant role in holding back the tide of bad content. How to achieve this without making better-quality information the privilege of the economically more advantaged is difficult and important.

The Cost of Customer Service

A second way that financial factors regulate online behavior is the cost of customer service. For many sites, customer service is their second-largest non-fixed cost (after paying for bandwidth). How much content regulation is affordable fundamentally shapes what kinds of sites are possible.

This was particularly evident in the United States after the passage of the Children's Online Privacy Protection Act (COPPA) in 1998 (Complying with COPPA: Frequently Asked Questions 2020). Before this legislation, companies were requiring children to give away personal information as a precondition of playing online games. COPPA ruled that you could only ask kids for personal information with parental consent. Getting parental consent costs money—employees need to be hired to provide customer support for the process. After the law took effect, many sites for kids closed, because their business models were no longer viable. New sites slowly opened over the next several years. In the end, sites for kids evolved into two groups: simple sites supported by advertising dollars that collect no personal information from kids, and more

complex sites supported by subscription fees. The subscription fees help defray the cost of verifying parental consent. Changes in required customer support and level of content regulation reshaped the landscape of what sites for kids exist in a fundamental way. All of this is also true in sites for adults.

Data on how much companies spend on customer service is typically proprietary, but the story of Yahoo Answers gives some insight. In the late 2000s, developers at the question and answer site Yahoo Answers had a problem with growing amounts of offensive content. Hiring customer service representatives to manage content was costing a million dollars a year, and it was taking on average eighteen hours to respond to complaints.

Randy Farmer and colleagues came up with a clever solution: they built a reputation system for site members who report bad content. If a member has a history of correctly reporting bad content, then Yahoo Answers trusts that person and takes down content they report immediately. The person who posted the content can file an appeal to dispute the removal if they wish. Those appeals are handled by a customer service representative. If an appeal is successful, it lowers the reliable reporting reputation of the person who reported it. Yahoo can tell who to trust to make accurate reports.

When the system was launched, Yahoo found that few post removals were repealed—the people posting offensive content knew an appeal wouldn't be successful. As a result of this system, the average time it takes to manage complaints went down from eighteen hours to *30 seconds*, and the cost to Yahoo went from a million dollars a year to less than ten thousand dollars a year, total. Farmer and Bryce Glass speculate that that part of their jaw-dropping results

was a sort of "broken window effect"—trolls quickly learned that Yahoo Answers was no longer a fun place to breach norms, because the content would be removed so quickly (Farmer and Glass 2010). This solution leverages human intelligence rather than artificial intelligence. It simultaneously creates a new way to volunteer to help the site and empowers users. It's a clever way to reduce the cost of customer service.

In 1994, I organized a panel discussion at the ACM Computer-Human Interaction (CHI) conference on "Approaches to Managing Deviant Behavior in Virtual Communities." In the panel proposal, I very sincerely wrote that I advocated a psychological approach to managing bad behavior. An administrator having a serious heart-to-heart chat with someone who is annoying others could solve the problem in a deeper sense than just sanctioning them (Bruckman et al. 1994). In one sense, I wasn't wrong—that approach really can help solve underlying problems in a more lasting way. However, with the scale of today's internet, the suggestion is absurd. Volunteer moderators don't have time to devote minutes much less hours to each incident, and staff moderators cost money. Software programs are less expensive than human staff.

With unlimited funds, truly outstanding content moderation is possible. If I can hire enough paid staff and train them well, I can handle difficult content with speed and finesse. With a realistic budget, the situation is different. I organized another panel discussion at CHI on the same topic twelve years later in 2006 (and again in 2018). Reviewing what we knew in 2006 that we didn't know in 1994, a key observation was the fundamentally financial nature of the challenge (Bruckman et al. 2006).

What sort of behavior is allowed fundamentally shapes what a site becomes. Behavior management is shaped by the software systems and human workers (paid and volunteer) who manage it. Both software and human labor cost money. Even using volunteer labor, you need to invest money to pay administrators to respond to volunteer concerns, and build tools for volunteers to use. To understand the management of online behavior, you need to appreciate the connections between the social, technical, and financial aspects.

The Evolution of Business Models for Online Communities

The third way that markets regulate online content is through the business models that support platforms. All the syllabi for my Design of Online Communities class are online, from 1998 to the present. I enjoy looking at past years, and thinking about how my knowledge and the field's knowledge have evolved. One of the things that has changed the most over the years are the business models. The things I taught my class about topics like community and identity in 1998 still broadly hold true today. But I couldn't have told you much about future business models to support the internet. Those business models drive site design, so they're of central importance.

The evolution of the business model of Cartoon Network is a good example. Turner Broadcasting is near Georgia Tech, and Chris Waldron from Turner and Cartoon Network came to give guest lectures in online communities class for many years. In the late 1990s when he first visited, online ads generated negligible revenue for Cartoon Network. Their revenue came from television ads

and merchandising. Ads on the website were initially free if you purchased a television ad. The idea of targeted advertising wasn't yet generally understood. Today, online ads are a significant source of revenue for Cartoon Network and much of the technology sector of the economy.

A fundamental assumption behind our market-based economy is that consumers will make rational decisions in selecting products and services. Companies that provide quality products at a fair price will thrive, and others will fail. In this model, everyone simply needs to look out for themselves, and the magic power of the market will ensure that good things will happen. Except that our actual markets are fairly far from idealized markets, and things don't really turn out that way. The current state of the internet is an example.

Business Models as Regulation

Commercial internet sites are driven primarily by the profit motive. Since the 1970s, it's been fashionable in financial circles to argue that publicly traded companies have an obligation to their stock holders to "maximize shareholder value" (Zuboff 2019). This is neoliberalism, "an ideology of unswerving loyalty to the logic of the market" (Field 2019, 5). Stock in those corporations is likely in your retirement fund, so the idea that they might try to thrive financially has merit. However, if that is their *only* guiding light, it leads to problems.

Underlying this ideology is the assumption that making financial gain the top priority is a kind of objective function that removes bias from the system. In other words, what makes the most money is the right choice for the economy and society. There is no evidence that this is true.

Consider the case of the YouTube recommendation algorithm. To make the most money, YouTube designed their algorithm to try to maximize each person's time on the site—so they would see the most advertisements. They discovered that leading people to more and more niche content and more controversial content keeps them on the site longer. This increases revenue from advertisements. That sounds like a reasonable design choice until you learn that much of that more-specific content advocates conspiracy theories and ideas not supported by mainstream science. Zeynep Tufecki notes that a simple search for mainstream political figures led her to conspiracy videos claiming that the attacks on the US in September 2001 were perpetrated by the government. The algorithm quickly takes people to low-quality content, like videos promoting false conspiracy theories and hate groups. You can maximize shareholder value better if you encourage people to watch extremist content—it can be quite engrossing. Tufekci concludes that "YouTube leads viewers down a rabbit hole of extremism, while Google racks up the ad sales" (Tufekci 2018). By maximizing profits, YouTube became an engine for convincing people of crazy nonsense.

The driving goal of YouTube and similar platforms is to show people more ads. The more you know about individuals, the more you can both keep them on your site and also better target ads to them. Shoshana Zuboff calls the business model of gathering trace data of people's online behavior for the purpose of targeted marketing "surveillance capitalism." In her book *The Age of Surveillance Capitalism*, she points out that many sites deliberately manipulate people to reveal more and more about themselves, for the purpose of marketing to them. As a result, she argues that companies like YouTube are treating people as means

to an end—not as ends in themselves. This is *the definition of unethical* (the second formulation of the categorical imperative) as articulated by philosopher Immanuel Kant (Kant and Patton 1964). In the YouTube example, they are trying to maximize ad views without any thought to the impact of false content on individuals or society (Zuboff 2019).

In surveillance capitalism, companies gain more and more knowledge about individuals. Zuboff comments that "unequal knowledge about us produces unequal power over us, and so epistemic inequality widens to include the distance between what we can do and what can be done to us" (Zuboff 2020).

Some targeted advertising may be useful, but letting the desire to do better targeted advertising *drive the design of the entire system* is leading to bad outcomes for us as consumers and citizens. One possible solution to this dilemma is for companies to *better articulate and prioritize corporate values*. YouTube's mission statement says their goal is "to give everyone a voice and show them the world"("YouTube About" n.d.). It would be better if they aimed to show people the world *as it really is*—to spread knowledge. As we saw in Chapter Three, "knowledge" is justified, true belief. YouTube has begun to make constructive changes to the platform by no longer recommending content about conspiracies (Rosenblatt 2019) and banning misinformation related to public health (*BBC* 2020).

If platforms better articulate their values, then users can ideally choose among alternatives and select sites that appeal to them. If YouTube decides that they are going to prioritize "good" content, then they are in the uncomfortable position of having to decide what good content is. If there is only one platform,

then it makes sense for it to allow all sorts of content. But you can imagine a world in which there is a marketplace of venues with different content policies, and then each person can choose a site that suits them. This already exists to some extent but there aren't yet enough alternatives. When a few big sites dominate, there is no marketplace.

Another critical problem is that the values of existing platforms are often not clearly expressed to members. Some politically radical sites have been founded with clearly stated, explicit values. This has the benefit that users know what to expect. If you go on the alt-right site Gab and post in favor of liberal ideas (like advocating free college tuition or government-funded health insurance), you would expect to be downvoted or maybe even banned. Some of my students have tried it, and that's what happened! Where things get fraught is when users don't have clear expectations for what is allowed on a site, and then are surprised by a content removal. ("What do you mean I can't say that here??")

Even if a platform tries not to have specific values, they end up being forced to make hard decisions. For example, in 2016 Facebook removed a photo of a naked 9-year-old girl and banned the person who posted it. Without more detail, that seems not only reasonable but mandatory. However, the specific photo was the 1972 Pulitzer-Prize-winning photo of a girl in Vietnam burned by napalm. The photo is historically significant and catalyzed a shift in American opinion about the war. Protests arose in response to the removal and the photo was restored (Shahani 2016). If it's hard to decide whether to allow a photo that is already widely recognized as historically important, imagine how hard the decision will be when the next shocking but important photo emerges. As

Tarleton Gillespie eloquently documents in his book *Custodians of the Internet*, there is no way today for platforms to avoid hard decisions (Gillespie 2018). These decisions might be easier if corporations articulated a vision for their impact on the world more nuanced than maximizing shareholder value.

In 2020, Facebook launched an "Oversight Board" (Oversight Board n.d.) to help it make such hard decisions in the future, and is aiming for that board to function independently (Complying with COPPA: Frequently Asked Questions 2020). This approach is promising and it will be interesting to see if it functions as hoped, and if other companies follow suit.

An advertising business model pressures platforms to maximize time on the platform and hence ad views. On the other hand, a flat fee, subscription-based business model means it benefits the platform financially if people participate just enough to love it and feel they get value but no more (since usage incurs bandwidth and support costs). It's intriguing to consider what other business models are viable, and how those would shape incentives for platforms and users.

As we have seen, the business model of a site shapes its content policies. It also shapes where the site invests its precious, scarce development time for software engineers. What a site becomes depends on what its leaders value, and the bottom line can dominate executive's thinking. A bit more attention to what kind of world a site is creating would help. My optimistic hope is that sites that pay more attention to values will actually make more money in the long run. People switch to the email provider Protonmail because it preserves privacy and security. We need lots more ProtonMail's—alternative sites guided by values.

What New Business Models are Possible?

We have not yet invented all the forms of online community. Some of them we may stumble into. In fact, the collaboration platform Slack began as an internal development tool for a development team who were building a game—and then the engineers realized that their collaboration tool was a better product than the game itself (Baer 2016). Other innovations are likely to be carefully planned. Understanding the nature of community, the value it brings to people, and the ways to support its growth may lead us to design the next form and the next business model.

We can choose business models that better incentivize the development of platforms that add to our individual and collective well-being (eudaimonia). If someone creates a new site that is like YouTube but everything is consistent with mainstream science, I'm going to switch to it. The choices you make as a consumer change what platforms become. I believe that there is pent-up demand for better-quality internet platforms, and you can do well by doing good. YouTube is just one example—this analysis applies to most internet sites and platforms. The profit motive alone does not magically make the right thing happen. Competition is needed—the best outcome for advancing both free speech and truth is to have lots of sites with competing standards, and let people decide which ones to frequent.

In neoliberal theory, healthy competition should solve all problems. I wonder if that could work, if there really was strong competition. Regardless, it's clear that it doesn't work when actual competition never emerges. Productive competition is not naturally emerging in our current economic system. Our

current information space is dominated by a few huge players. We need public initiatives to foster the growth of healthy competition. Where that competition fails, some problems need to be changed with policy—passing laws that regulate what is not working.

New business models need to be *invented*. What do people value, and what are they willing to pay for? Targeted marketing is inextricably bound up with issues of privacy. Do people care about their privacy? Are people willing to pay more for privacy-preserving services? If so, then future entrepreneurs may invent new business models for privacy-conscious consumers. As of August 2020, the privacy-preserving search engine DuckDuckGo commands 0.5% of internet searches (Search Engine Market Share Worldwide n.d.). It will be interesting to see how their market share evolves over time, and whether more privacy-preserving options emerge for other types of software.

Each business model creates different incentives for the business and its users, and any business model can be leveraged in a more responsible or more exploitative way. The risk for users of some platforms is spending too much time or money there. This can damage the quality of people's lives. Analyzing in-game purchasing systems in video games, King et. al. note that "Some in-game purchasing systems may represent financial hazards that contribute to player over-commitment to gaming activities and increase risk of negative financial and psychological consequences" (King et al. 2019, 141). Platforms can decide to encourage as much use as possible, or to encourage reasonable use. My hope is that designers will recognize the spectrum of outcomes, and design systems that encourage balanced expenditures of both time and money.

In an ideal world, internet entrepreneurs would naturally do well by doing good—i.e. people who make less-toxic sites that enrich participants' lives will also make more money. Currently, that's often not true. The intriguing question is whether we can invent new business models that align these goals more effortlessly.

Business models seem to drive our design decisions and are driving us in unhealthy directions. We have two options to remedy this. One is to invent new business models that drive us in better directions. The other is to bravely ignore the bottom line and take a more nonprofit view—to put our values before our finances. I'm not sure if there really are magic business models that will make the internet a better place—that may be a naïve hope. But one thing is certain: *More of the internet in the future should be nonprofit.* A nonprofit business model and world view will let platforms put the needs of individuals and communities first.

There have been a variety of attempts at improving the internet through nonprofit means. The most notable success is Mozilla, which builds open-source tools that embody their vision "to ensure the Internet is a global public resource, open and accessible to all. An Internet that truly puts people first, where individuals can shape their own experience and are empowered, safe and independent" (Mozilla n.d.). Under this umbrella, they have created the Firefox browser (which is more privacy preserving than other browsers), open-source audio tools, virtual reality toolkits, and a host of other things. Part of why Mozilla has been so successful is because they have a business model which piggybacks off of for-profit models. Every time someone does a search in the Firefox browser, the search engine used pays a small fee to Mozilla. They are funded by

advertising indirectly, but in a way that doesn't change their basic values and mission.

Other attempts at nonprofit civic spaces have taken a low-budget approach. For example, the free, open-source software (OSS) Mastodon provides an alternative to Twitter and the OSS project Mattermost provides an alternative to Slack. It's tremendously hard for tools like these to achieve wide public adoption, because they can't look or function like commercial sites built by teams of thousands of full-time designers and engineers. Most people have never heard of them, and wouldn't find them usable or appealing if they tried them. Public funding of nonprofit social media efforts like these could enhance their value and give them a greater chance to win users and have impact. If we value social spaces that contribute to our civic well-being, then we need to publicly fund them as civic projects (like highways, railroads, and rural telephone and broadband access).

Although I do believe in the potential positive impact of open-source tools, they can still be misused. For example, the hate-speech site Gab has adapted ("forked" in OSS terms) the software for Mastodon ("Statement on Gab's fork of Mastodon" 2019). Tools are tools—a hammer can be used to build a house or to break windows. We need social and legal consequences for people who misuse them.

Theoretical Summary

The business models that support socio-technical systems shape what those systems become. If we understand the impact of financial forces, then we can

work to make finances *serve our larger goals*, rather than driving a system's design to unhealthy ends.

Practical Implications

Markets shape online behavior in three notable ways. First, people's willingness to pay for access to a site filters who participates. Filters on who participates can reduce bad content and improve content quality, but also can work against equity if some people are financially excluded.

Second, the cost of managing online behavior shapes decisions about what is allowed, and the resources they have (both human and technical) to enforce the standards they envision. Perfect moderation is possible in theory if financial resources are unlimited. With realistic resources, platforms need to make hard choices. To a surprising degree, the cost of managing behavior shapes what online sites are viable.

Third, the way a site is financed changes the site's priorities. Sites that prioritize financial results over all other factors often end up promoting content that does not improve the lives of members or the state of the world. Instead, sites need to articulate values for what sort of world they want to help create.