

Decentralizing Platform Power: A Design Space of Multi-Level Governance in Online Social Platforms

Shagun Jhaver¹ , Seth Frey², and Amy X. Zhang³

Abstract

Many have criticized the centralized and unaccountable governance of prominent online social platforms, leading to renewed interest in platform governance that incorporates multiple centers of power. Decentralization of power can arise horizontally, through parallel communities, each with local administration, and vertically, through multiple hierarchies of overlapping jurisdiction. Drawing from literature on federalism and polycentricity in analogous offline institutions, we scrutinize the landscape of existing platforms through the lens of *multi-level governance*. Our analysis describes how online platforms incorporate varying forms and degrees of decentralized governance. In particular, we propose a framework that characterizes the general design space and the various ways that *middle levels* of governance vary in how they can interact with a centralized governance system above and end users below. This focus provides a starting point for new lines of inquiry between platform- and community-governance scholarship. By engaging themes of decentralization, hierarchy, power, and responsibility, while discussing concrete examples, we connect designers and theorists of online spaces.

Keywords

online social platforms, platform governance, decentralization

Introduction

In the months following his purchase of Twitter, Elon Musk triggered an exodus of millions of users to competing platforms; for instance, an estimated 2M users, about 1% of its user base, made accounts on a federated competitor Mastodon (Peters, 2022). These departures were partly in response to controversial changes to Twitter policies (Huang, 2022). Twitter is a centrally governed platform, where moderation decisions are made by those in power at the company and carried out unilaterally to Twitter's millions of users. In contrast, several of the platforms to which Twitter users migrated (e.g., Mastodon, Bluesky) are designed to offer more decentralized governance. For instance, users on Mastodon can move between thousands of instances with different governance arrangements and moderation rules. This *middle level* of local administration formed through its federated architecture gives Mastodon a different approach to privacy, safety, growth, censorship, autonomy, management, and other fundamental governance properties (Rozenshtein, 2022).

Increasingly, users are weighing questions of governance design when deciding what platforms to join, in addition to common criteria such as affordances for posting, the user

base, or the platform's culture. As platforms experiment with forms of governance beyond simply centralized, partly fueling a broader fragmentation of the social media landscape (Chatterjee, 2023), more precise terminology to describe how a platform is designed to decentralize governance becomes necessary for differentiating between platforms, particularly as the term "decentralization" carries many meanings. As an example, while Mastodon may seem clearly different from Twitter in terms of having multiple centers of power, how does it compare to Reddit? Much like Mastodon, Reddit is organized into many communities with a variety of governance arrangements, moderators, and rules. However, Reddit still has a powerful *top level*, enabling it to take platform-wide actions, such as banning

¹Rutgers University, USA

²University of California, Davis, USA

³University of Washington, USA

Corresponding Author:

Shagun Jhaver, Huntington House 203, 184 College Avenue,
New Brunswick, NJ 08901, USA.

Email: sj917@rutgers.edu

X: @shagunjhaver



users and communities and meddling in local governance, which would be impossible on Mastodon. These differences are consequential, as was recently seen when the chief executive officer (CEO) of Reddit, Steve Huffman, pushed top-down changes to combat a protest among Reddit community moderators (Ingram, 2023).

The distinctions between different flavors of decentralized governance arrangements can be characterized using a lens of *multi-level governance*, where end users, community moderators, and platforms sit at different levels. Thus, in this research, we ask: What is the range of variation in how platforms organize multi-level governance systems, and what are the implications of this range for the value users receive from their platforms and the content they volunteer to contribute?

The increased attention to platform governance and power struggles at platforms we have seen in recent years mirrors the growing influence of online social platforms on society. Over the last three decades, much of social media activity has consolidated around a handful of platforms that host content for billions of people (DeNardis & Hackl, 2015). More recently, scholars and regulators have expressed alarm at the centralized and opaque nature of these major platforms, including their lack of procedural fairness and accountability (Chen et al., 2020; Fan & Zhang, 2020). Furthermore, some assert the futility of crafting a single set of rules that can be consistently applied over a large and diverse user population, invariably harming groups that do not fit a global standard (Jhaver et al., 2019; Seering, 2020).

In reaction, experts have called for greater decentralization of platform power to enable a plurality of moderation choices for users (Fukuyama et al., 2021; Rozenshtein, 2022). Despite, or perhaps because of its currency, the term decentralization is overloaded with multiple senses. For example, is a technically decentralized architecture with a clear charismatic central figure centralized? Is a strongly centralized architecture with stronger member political engagement decentralized? In this work, we are primarily concerned with differences in governance *design*, or the ways in which governance is intended to be carried out according to platform creators, designers, and implementers. Importantly, design encompasses technical architecture as well as intentional administrative processes, which can more readily change. For instance, a platform may be centrally hosted but choose not to exert top-level power over local communities. In some cases, this is bound by a set of bylaws, such as with the Wikimedia Foundation; in other cases, the boundary is undefined yet tacitly agreed upon and can be renegotiated, such as with Reddit.

We summarize the major ways in which multiple centers of power are introduced to a platform's design. One way to achieve greater decentralization of platform power is to introduce intermediate levels of local administration (Rajendra-Nicolucci & Zuckerman, 2021). Reddit offers semi-autonomous "subreddits" governed by volunteer users, while YouTube channel owners can moderate comments on

their videos. A second way is to have existing platforms enable application programming interfaces (APIs), plugins, or in-platform subscriptions to foster a marketplace of governance services (Frey & Sumner, 2019; Fukuyama et al., 2021; Schneider et al., 2021). This marketplace allows users to choose from a suite of possible governing methods. For instance, Twitter blocklisting tools allowed users to subscribe to the blocklist of their choice (Geiger, 2016; Jhaver et al., 2018), serving as a weak form of local administration that only governs account blocking. More recently, Bluesky introduced user-led moderation and curation tools, such as custom feeds and mutelists, within the platform (Graeber, 2023). Finally, decentralized governance can arise through the use of a shared technical specification of social exchange, such as peer-to-peer and federated protocols (e.g., Mastodon uses the ActivityPub protocol), or a shared set of data, such as blockchain technologies, including cryptocurrencies and decentralized autonomous organizations (DAOs) (Wright & De Filippi, 2015).

In cases where platform governance incorporates decentralization, challenges still emerge. The different governance units sometimes operate in isolation (Caplan & Gillespie, 2020; Jhaver et al., 2019), resulting in redundant problem-solving and wasted effort. And without a central authority, local units face greater challenges addressing networked harassment like non-consensual sexual imagery (Marwick & Caplan, 2018; Masnick, 2019). The specific design details of how platforms carry out decentralization of governance are vital to achieving success and sustainability. In this work, we characterize the diverse landscape of decentralized governance in online social platforms. By cataloging contemporary platforms within a design space, we reveal the designs that remain to be explored.

Contributions

We contribute a set of five dimensions to characterize the design landscape of decentralized governance on existing online social platforms. We borrow primarily from the political science literature (Armitage, 2008; Bache & Flinders, 2004), which has used the frame of *multi-level governance* to describe decentralization in the form of authoritative decision-making dispersed vertically across multiple levels and horizontally over many local governance units (Hooghe & Marks, 2001).

In our characterization of multi-level governance in online social platforms, we focus our attributes primarily on the *middle levels* of governance. At middle levels, many interactions play out with end users at the lowest level of a platform and with centers of power, authority, or ownership that comprise the topmost level (Figure 1). We focus on middle levels because they are the main arena for low-level agents to collectively organize with or against each other, with or against top-level agents, and even to interact with other kinds of middle levels.




Figure 1. A depiction of multi-level governance in an online platform, with a top, middle, and bottom level and typical actors within each level.

Note. End users (denoted by the circles) are in the bottom level and are governed by one or more governance units in the middle level. All end users are additionally governed by the top level. Each unit in the middle level is also governed by the top level. Top levels can be concentrated and powerful (as in the cases of YouTube and Facebook) or offer a looser set of constraints (as in the cases of Wikipedia and Mastodon). Similarly, middle-level units can have a lot of governance capabilities (as in the cases of Minecraft servers and Reddit subreddits) or have only a few (as in the case of third party shared blocklists on Twitter).




Figure 2. Our four categories of design dimensions, including cross-level, within-level, within-unit, and system-level dimensions.

As it turns out, middle levels are ubiquitous in online social platforms. We draw upon contemporary examples in this work, including Facebook Groups, Twitter shared blocklists, Reddit subreddits, YouTube and Twitch channels, Mastodon nodes, Minecraft servers, WhatsApp Groups, and Wikipedia language editions toward an understanding of the general design space of multi-level social systems online.¹ As part of proposing dimensions for how middle levels can vary in structure and design, we identify the other parts of platform structures that a general framework for multi-level design would address. We describe not just the cross-level “vertical” dimensions that this manuscript focuses on but also the cross-unit “lateral” dimensions and within-unit and system-level dimensions (see Figure 2).

Given competition in an increasingly fragmented social media landscape (Chatterjee, 2023), we are entering an era where we expect a greater variety of governance structures will be attempted. Through our focus on multiple centers of power, we bring attention to how new designs can negotiate between the advantages of centralization versus decentralization toward more ethical, sustainable, and empowering online platforms. Beyond theoretical insights, we offer practical design implications for online communities, drawing on lessons learned from offline and online institutions. We also discuss how a substantive research agenda can build upon our characterization of multi-level governance.

Related Work

Regarding the governance of online social platforms, scholars have primarily distinguished between platform governance and community governance. While platform governance scholarship focuses on “macro”-scale policy, compliance, and legal questions, community-governance research focuses on the efforts of users and volunteers to engage in local community norm-setting (Gillespie, 2018; Grimmelmann, 2015; Seering, 2020). Although scholars in both communities acknowledge the importance of the other, the scholarship lacks a general framework for how they substitute, complement, or interact with each other. Our research offers early steps toward bridging the gap between platform and community governance scholarship by examining critical aspects of their interaction.

Though community-reliant platforms are decentralized, many incorporate a top level of governance comprised of platform operators, developers, or a global content moderation team overseeing the different communities. In addition, some platforms, such as Facebook, combine unitary and multi-level governance, with the Facebook News Feed governed via a centralized model and Facebook Groups employing a community-reliant model. The addition of a middle level allows end users to act as oftentimes volunteer community moderators. In this role, end-users can influence, monitor, and engage with platform operators (Chen et al., 2020). Decentralization also enables platform operators to leverage local information and innovations to improve the informational efficiency of their governance (Chen et al., 2020; E. Ostrom, 1990). Thus, with middle levels, platforms are more likely to attend to the welfare of all users.

While the literature on platform and online community governance spans multiple disciplines (Gillespie, 2018; Grimmelmann, 2015), the interplay between different power centers and jurisdiction levels has yet to be systematically characterized. However, many individual case studies have highlighted a growing interest in interrogating this interplay.

For instance, research on Reddit and Twitch content moderators has illuminated how they often govern multiple online communities, with tools and resources sometimes shared across communities (Jhaver et al., 2019; Matias, 2019). Others have examined how YouTubers, who govern channels at YouTube's middle level, collectively organized to pressure the platform to change its demonetization policy (Tait, 2016). Frey and Sumner (2019) show how player-run Minecraft communities self-organized around an emergent volunteer ecosystem of shared governance plugins. In other work, researchers have shown how top levels of governance regulate units in the middle level that they believe are governing poorly (Chandrasekharan et al., 2017, 2020). These cases suggest a broader pattern arising across different platforms regarding interactions between levels, one that may be informed by general scholarship on multi-level institutions from political science, which we next describe.

Drawing From Offline Governance Literature

Given the lack of online governance literature that systematically characterizes inter-level interactions, we turn to the offline governance literature and examine how its prevalent theories can map onto online platform governance toward developing our design space. We focus primarily on political science literature because its theories incorporate concepts such as conflict, competition, and contestation between and within levels.

There is no universally accepted governance theory in political science; the field has many overlapping theoretical discussions and debates (Ansell & Torfing, 2016; Tiihonen, 2004), but several have been developed for understanding multi-level governance systems. We draw from *federalism*, which focuses on how nations divide power between a central government and local states, and *polycentricity*, a general framework for how institutions with multiple centers of power compete and cooperate given overlapping jurisdictions. These frameworks provide an analytical structure for our study of multi-level governance and a means to challenge and strengthen our imagination beyond existing online examples (Aligica & Tarko, 2012; Scheuerman, 2004).

Federalism

One prominent theory relevant to our focus is federalism, a system of governance that divides a political territory into semi-autonomous states that share authority with one another and with a common central government (Aroney, 2016; Elazar, 2014). Each governing unit in this system can make laws directly affecting the citizens within its territorial purview (Watts, 1999). Political scientists (traditionally, primarily Western and American) have taken the federated organization of the United States as a paradigm

of multi-level governance. However, federations exist on a continuum. More peripheral systems have member states bound by loosely structured trade and defense alliances. The other extreme is “administrative decentralization,” where member states have little autonomy and behave more like the administrative units of an organization (Bednar, 2011; Elazar, 1994). In their idealized form, federal systems consist of neatly nested jurisdictions, with the top-level jurisdiction equal to the union of the non-overlapping jurisdictions of all lower level systems. Within this framework, parallels to the online realm are apparent. For instance, Reddit is composed of the union of its distinct subreddits, each having some freedom to implement governance, and there is platform-level governance over all subreddits.

Aroney (2016) offers a taxonomy of federalism that is particularly well-suited to describe the complexities of multi-level platform governance. His taxonomy details features describing (1) whether the system is a federation or confederation (*i.e.*, whether a unit can leave unilaterally or not), (2) whether it is formed by the aggregation of formerly independent governance units or the devolution of a formerly unitary government, (3) whether its different levels wield redundant or complementary powers over their members, and (4) whether its lower level governance units are symmetrical or asymmetrical in their powers, rights, and roles. The concepts of symmetric and asymmetric federations distinguish a body of equal states (as in the idealized United States) from a mix of states, commonwealths, districts, and territories with very different levels of autonomy and representation (as in the realized United States). In the next section, we use Aroney’s taxonomy to propose some dimensions of our design space.

While several concepts transfer well to our analysis, the generality of this literature has been hampered by its focus on federalism in nations, at the expense of smaller (and more numerous) multi-level governance, a gap that *polcentricity* was introduced to fill.

Polycentric Governance

In essence, polycentricity is an expression of self-governance capabilities that, over time, will produce a complex system of governance institutions (van Zeben, 2013). Although there is no clear consensus definition of polycentric governance, most scholars agree that it consists of: (1) multiple decision-making units with overlapping domains of responsibility, (2) that interact through a process of mutual adjustment in complex and ever-changing ways, and (3) generate a regularized pattern of overarching social order that captures efficiencies of scale at all levels of aggregation (Aligica & Tarko, 2012; McGinnis, 2016). Researchers have developed many different models of polycentric governance systems to build greater clarity and specificity around the concept and highlight its posited advantages (Aligica & Tarko, 2012; Carlisle & Gruby, 2019).

The concept of polycentricity is pivotal to the Ostrom school of institutional economics, pioneered by the work of Vincent and Elinor Ostrom (Aligica & Tarko, 2012). V. Ostrom adopted the term “polycentricity” to describe governmental fragmentation in US metropolitan areas (V. Ostrom et al., 1961). Later, E. Ostrom’s (1990, 2010) research on community-based collective management of natural resources became the best-known application of polycentricity to real-world settings. Since then, several scholars have explored polycentric governance for sustaining natural resource systems (Blomquist & Schlager, 2005; Marshall, 2015).

Federalism can be viewed as a type of polycentricity (Aligica & Tarko, 2012) since it incorporates many of the critical elements of the Ostroms’ theory, including multiple power centers and redundant jurisdictions (van Zeben, 2013). However, while federal systems consist of neatly nested jurisdictions under a single highest center of power, polycentric systems include crosscutting “issue-specific” jurisdictions and envisage an explicit role for autonomous private corporations, voluntary associations, and community-based organizations (McGinnis & V. Ostrom, 2012). This view is well-suited to inform the analysis of multi-level online platforms containing “nested quasi-autonomous decision-making units operating at multiple scales” (Folke et al., 2005). We find analogs to polycentric systems on sites like Twitch and YouTube, where different channels operate independently but account for each other through cooperation, competition, conflict, and conflict resolution (V. Ostrom, 1991).

A Design Space of Online Multi-Level Governance

This section presents a series of dimensions that characterize the different types of multi-level governance structures that exist within popular online social platforms. We curated this design space through an iterative exploratory process (Emmet, 1964), combining taxonomies from federalism and polycentricity literatures with insights gleaned from case studies and our own expertise with online governance. This process was iterative and involved interpretation, redefinition, and verification of emerging dimensions, their comparison with one another, and their relation to prior literature.

Our analysis centers on the *middle levels* of governance, including their interactions with the top and bottom levels. Table 1 lists examples of middle levels that our design space could characterize. We ensured that this selection of middle levels expressed a diversity of structural features (segregated, largely centralized, third-party, etc.) and represented popular platforms focused on a wide range of topics (e.g., gaming, peer production) and formats (e.g., streaming, bulletin-board).

We strengthened the practical utility of our emerging design space by applying it to middle levels in these 13 implementations (Table 1). For this, each author familiarized

themselves with each platform and its governance approach by installing its site (when necessary) and reviewing the site features and moderation documentation. Next, the co-authors independently considered how each platform’s governance structure can be characterized through our design space. Following this, we compared our characterizations, and especially focused on resolving disagreements. During this process, we updated the dimension definitions to further reduce ambiguities.

We identify four sources of design variation (Figure 2) in how online platforms structure multi-level governance:

1. *Cross-level* or “vertical” dimensions of how middle-level units interact with other levels;
2. *Within-level* or “lateral” dimensions of how middle-level units interact with one another;
3. *Within-unit* internal dimensions of the middle-level units; and
4. *System-level* dimensions of the whole platform.

In this work, we focus on cataloging cross-level dimensions in detail. However, by situating them in a broader design space, we give scholars an appreciation of the range of implications and tensions that any design decision has for other parts of the system. Next, we conceptualize the four sources of design variation before we focus our discussion on the cross-level dimensions.

By *cross-level*, we mean variation in how middle-level units can interact with their platform administration above or member individuals below. We introduce and elaborate upon the following five cross-level dimensions in the next section: (1) overlap of jurisdiction, (2) cross-cutting membership, (3) degree of autonomy from above or below, (4) degree of authority over above or below, and (5) degree of support by above or below.

By *within-level*, we mean variation in middle-level units or how they relate to or interact with one another. For instance, Slack workspaces have software support to connect via shared channels, as do Wikipedia articles for linking across language editions, while Mastodon nodes can federate with or block other nodes. Variation also arises from “transit costs” between units: one can easily subscribe and unsubscribe to subreddits, while in WoW “lateral exit” is more costly (Kurild-Klitgaard, 2010). Multi-level architectures can also differ in whether their middle-level units are of the same “type”: While subreddits are mostly symmetric in the formal powers available to them, YouTube provisions its Channels differently depending on whether they are run by amateurs, legacy media organizations, or contracted producers (Caplan & Gillespie, 2020).

By *within-unit*, we mean cross-platform differences in the internal characteristics of middle-level units with implications for their multi-level architecture. One example is whether the middle-level units are part of the formal architecture of the platform or emergent from users. Twitch

Table I. Examples From Online Social Platforms of Middle Levels Within a Multi-Level Governance Structure.

Platform	Middle level	Description
Facebook	Groups	User-created communities that are moderated by volunteer members. Activity within all groups is also moderated by Facebook's global content moderation carrying out Facebook's site-wide community guidelines.
Reddit	Subreddits	User-created and managed communities, each containing its own community guidelines and automated moderation configurations. Subreddits may be banned or quarantined by the centralized Reddit administration.
Twitch	Channels	Channel owners live stream their videos, and each channel has its own set of moderators and settings for automated moderation of comments. The Twitch platform also enforces site-wide community guidelines.
YouTube	Channels	Channel owners post videos and control the comments posted on their videos. YouTube's centralized moderation may take down videos that violate YouTube's community guidelines. It also has a site-wide filter that automatically moderates inappropriate comments.
Wikipedia	Language editions	Each language edition has its own processes for governing article content in that language. These editions are supported by committees, developers, and operators within the Wikimedia Foundation.
WhatsApp	Groups	Each WhatsApp group has admins who moderate that group. WhatsApp platform operators can make platform-wide decisions, such as limits on forwarding, but are limited in their power to moderate due to end-to-end encryption.
World of Warcraft	Guilds	Guilds are created by players, and can be banned by the platform, but are left to govern their own affairs. They vary greatly in their governance, membership, goals, and membership requirements.
World of Warcraft	Communities	Communities are created by players, and can be banned by the platform. They replace the user governance of Guilds with automated mechanisms for unacquainted players to form pick-up groups and approximate closer-knit guilds. Communities are now the dominant middle level of WoW.
Twitter	Shared blocklists	Third-party moderation tools relying on the Twitter application programming interface (API) to create lists of block-worthy users. The lists enabled subscribers to block accounts on the list automatically. Shared blocklists were either manually curated by a small staff of volunteers or automatically curated using algorithms.
Bluesky	Custom feeds	Users can browse a marketplace of feed algorithms that are created by other users and then add custom feeds to their home view. Custom feeds typically require third-party infrastructure to ingest, analyze, and rank content in real-time, and they can be updated by the feed maintainer(s) at any time.
Mastodon	Nodes	Each self-hosted Mastodon node has server administrators who decide moderation policies and what other nodes to federate with. User accounts are tied to a specific node. While Mastodon as an open-source project has developers, each node can decide the code running on its server so long as it follows the shared ActivityPub protocol.
BitTorrent	Filesharing communities	Self-hosted, emergent communities for filesharing, usually running forum software. Communities' requirements for membership tend to focus around management of free-riding. Communities have complete independence and can even use different filesharing clients, except that they all must use the BitTorrent protocol.
Minecraft	Servers	Although game owner Microsoft controls validation of user accounts, Minecraft worlds tend to be privately hosted. Administrators select plugins, world content, and have control over the membership and goals of their server, though the code itself is closed source.

The first 8 incorporate decentralization through the platform's design while being centrally hosted. Next, two examples (Twitter shared blocklists and Bluesky custom feeds) incorporate third-party tools to enable decentralization, while the last three examples involve the use of a shared technical standard connecting different servers.

channels, Mastodon nodes, and most other middle levels we have encountered belong to the former type. Examples of the latter type include Twitter shared blocklists and BitTorrent communities, which were created through collective community action with the help of API hacks and third-party

services. Another example of unit variation across platforms is the transparency of a middle-level unit's governance: Reddit communities have publicly viewable rules and rule enforcement actions by default, while these governance features are private by default in Facebook Groups.




Figure 3. Cross-level design dimensions.

Note. These dimensions characterize the relationships and interactions between middle-level units and the top level or between middle-level units and end users.

By *system-level*, we mean high-level design dimensions that influence multi-level architecture. One example is the depth of nesting of middle levels in each platform: Subreddits are the only formal level between Reddit and its users, whereas Wikipedia users have organized middle levels below the level of the language edition, such as WikiProjects. Another example of system-level variation is the support for multiple orthogonal types of middle levels: WoW has two forms of middle levels, the Guild and the Community.

Cross-Level Design Dimensions

We now conduct a deep dive into a key part of the general framework, the cross-level dimensions (see Figure 3) that characterize the relationships and interactions (1) between middle-level units and the top-level administration and (2) between middle-level units and end users in the bottom level.

Overlap of Jurisdiction. This dimension considers the degree of overlap between the middle and the top level regarding their areas of responsibility: What are each level's governance actions, and how much do those overlap? Overlap of jurisdictions quantifies the redundancy in governance actions between different levels, and it is a crucial feature of polycentricity (Aligica & Tarko, 2012; Carlisle & Gruby, 2019). V. Ostrom (1973) described governance systems with overlapping jurisdictions as "highly federalized" political systems. As an example, YouTube content creators can remove comments on their videos, just as YouTube platform operators can also remove comments on any video on YouTube. Here, the responsibilities and actions afforded to the middle level are a large subset of what is afforded to the top level, leading to a high overlap in the style of a nested federation.

In the cases of low overlap, we typically observe a "thinner" top level that has access to or chooses to exert few actions compared to the middle level. The middle level does not share the responsibilities the top level oversees. For instance, while Mastodon node administrators are in charge of almost every aspect of governance in each node, the top level mainly is involved in code development.

Cross-Cutting Membership. This dimension indicates whether the membership in each unit is exclusive—that is, whether two units can govern the same user. In WhatsApp, groups have inclusive membership—users can simultaneously join more than one group. In contrast, the WoW guilds are middle-level units with exclusive membership—characters can be part of only one guild at a time. Cross-cutting memberships distinguish the neat hierarchical nesting of federalism from the more unconstrained notion of polycentricity, which permits multiple middle-level units to assert jurisdiction over the same user.

When a platform has cross-cutting membership, a user's activity typically still is required to have a location in one unit or another that then determines the applicable local government. However, users may be banned from a unit due to their activity in another unit that is externally visible, so they may still experience some compounding constraints from multiple memberships. An example of this is the sometimes cascading effect of being banned from one Mastodon node or of being added to one oft-replicated Twitter shared blocklist, as units may look to other units' governance actions for guidance (Jhaver et al., 2018). Another example is bulk bans on Reddit, where moderators in one subreddit have at times banned all users who have posted on another subreddit to guard against brigades (Datta & Adar, 2019).

Degree of Autonomy From Above or Below. This dimension measures the extent to which middle-level units can operate independently or have accountability to another level that constrains their autonomy. It maps directly to a central concern of polycentricity and federalism scholars: The degree of freedom from control by other jurisdiction levels (Loughlin et al., 2013). Many platforms have middle levels that possess moderately high autonomy. For example, moderators of subreddits independently make a vast majority of subreddit-related decisions; they are not closely monitored by Reddit administrators.

As Marshall (2015) notes, *de facto* autonomy may matter as much as formal autonomy. While platforms may promise greater autonomy to middle-level units, they may still

employ governing-at-a-distance strategies to exercise greater control over outcomes, for example, by imposing reporting and compliance requirements (Carlisle & Gruby, 2019; Marshall, 2015). Another strategy is through the threat of sanctions. Middle levels have less autonomy when they are held to account by platform operators who can ban middle-level units that do not adhere to platform policies.² In some cases, technical affordances restrict this assertion of power. For example, in the case of WhatsApp, end-to-end encryption limits platform visibility into the actions of middle-level administrators.

We also conceptualize the degree of autonomy of a middle level from those below it. In the “implicit feudalism” model inherent on many social platforms (Schneider, 2021), a community’s founder has absolute autonomy over that community by default and is not directly responsible to users. Often, the only accountability is indirect, through the threat that users will exit for other more accountable or otherwise desirable communities (Frey & Schneider, 2020). The more rare and interesting case is a middle-level formally accountable to those below it. On Wikipedia, the community has the power to withdraw an administrator’s special privileges in cases of abuse of authority, restrict their use of certain functions, and place them on probation.³ Another case is when users can report abusive community moderators to the platform, at which point the platform uses its power to provide accountability; this is currently supported on Reddit.

Degree of Authority Over Levels Above or Below. A related but distinguishable dimension to a middle-level unit’s autonomy *from* is authority *over*, which refers to the ability to regulate or sanction. The question of authority is as central to federalism and polycentricity as autonomy (Morrison et al., 2019). Indeed, a middle-level unit having autonomy implies it has authority over something. However, piecing out the direction of its authority, like its autonomy, helps to distinguish different instances of multi-level governance online.

For example, admins of Facebook Groups have authority over Group members because they have the power to sanction members or their posts. These admins also make governance decisions autonomously, as they are usually not elected by Group members. On the other hand, admins of a Wikipedia language edition have authority over the governance of Wikipedia pages in that language. However, they do not operate autonomously since community members can strip away their administrator rights.

Middle levels usually have significant authority over low levels, with some work demonstrating the emergence of administrator oligarchies (Shaw & Hill, 2014). One exception we have found is Twitter shared blocklists (Geiger, 2016): As a third-party service with limited functionality, these blocklists could not strongly sanction their subscribers.

Middle levels can also have formal authority over higher levels, though most examples involve informal input. The closest example may be the Wikimedia Foundation, which,

likely due to its mission, has informal mechanisms for granting authority to its language editions. However, the most evident and concrete mechanism—the Foundation’s partially member-elected board—encodes a direct democracy that empowers bottom-level users rather than a representative democracy that might empower middle-level units. As another example, the Reddit Mod Council⁴ enables a select group of subreddit moderators to informally advise platform operators; however, there is no mechanism for sanctioning the top level.

Support by Levels Above or Below. On some platforms, middle levels are supported by higher levels through various means, such as receiving technical help or getting access to moderation resources. For example, Twitch provides streamers on each channel moderator tools that let them set rules to remove inappropriate content automatically. Discord launched a Moderator Academy for educating server mods.⁵ Middle levels can also be supported by lower levels, such as the private servers of Mastodon, Minecraft, or WoW, which sometimes depend on users’ financial support.

Relevant indicators of levels of support from above could include whether a middle level is self-hosted or hosted on a central platform and whether it receives funding, human staff support, or special access to developer tools from the top level. In addition, levels of support from below could be quantified based on whether a middle level is wholly or partially user-run and user-funded and has tools or culture supporting voluntary contributions.

Design Implications

As can be seen from our characterization of the landscape, multi-level governance can take many forms online. Decentralization has the potential to create governance more attuned to community needs. However, it can also mean more points of failure and more significant overhead and inefficiency (Bednar, 2008). We point to design implications of our analysis of multi-level governance for platforms and communities.

Enabling Innovation, Adaptation, and Healthy Competition

One benefit of decentralization is the ability of governance units to innovate, as emphasized in the phrase “laboratories of democracy” used to describe state governments in the United States. As communities evolve, their survival partly depends on their willingness to evolve their governance in response to events. The polycentricity literature emphasizes continual adaptation in changing environments (Carlisle & Gruby, 2019; E. Ostrom, 1999), which depends on centers having the capability to continually experiment with rules (Carlisle & Gruby, 2019; E. Ostrom, 1999). Thus, to enable innovation and adaptation, online social platforms should

provide autonomy and otherwise support local administrators' experiments with community guidelines, sanctioning criteria, or automation settings (Jhaver et al., 2019; Matias & Mou, 2018).

Competition for users is another outgrowth of increased decentralization that can lead to better user conditions. New communities often form when a faction of an existing community is dissatisfied with its moderation (Matias, 2019; McGillicuddy et al., 2016). According to polycentricity and federalism scholars, competition between different units induces self-regulating tendencies as it compels units to demonstrate their value to members (McGinnis & Ostrom, 2012; V. Ostrom et al., 1961; Thiel et al., 2019). However, competition can be detrimental for users when it leads to cross-community conflict or significant consolidation. Platforms with cross-cutting memberships may be able to avoid excessive competition, as joining one unit does not imply exiting another. This permits units to be complementary despite the appearance of competitive overlap.

Facilitating Social Learning

Many governance scholars identify social learning and building social capital as essential conditions for increasing resilience and sustainability (Berkes, 2010; Folke et al., 2005; Gelcich, 2014), without which each center must arrive at an optimum arrangement through trial and error (E. Ostrom, 1999). However, many platforms do not offer formal avenues for administrators in different centers to communicate with one another, leading to inefficiencies such as the building of redundant tools (Jhaver et al., 2019). Allowing for cross-cutting membership permits members to be in multiple communities and enables knowledge transfer. In their analysis of Wikia, Zhu et al. (2014) found that when users participate in multiple communities, their membership helps the survival of each due to the transfer of best practices.

Separately, the top level can also facilitate social learning. Given some overlap in the jurisdiction, the top level could promote specific middle-level solutions to additional middle-level units. For example, after many subreddits voluntarily used Reddit Automod, an automated moderation tool, Reddit realized its potential and incorporated it as a default mechanism on all subreddits (Jhaver et al., 2019). Another example is the Discord Moderator Academy, which included lessons, an exam, and an online community of fellow Discord moderators before being shut down in 2022.⁶

Establishing Accountability

Holding decision-makers accountable for poor performance is crucial for the proper functioning of governance systems (Carlisle & Gruby, 2019). Surprisingly, we find little evidence of top levels actively evaluating the moderation practices of middle-level units, except in the most public or egregious cases. One rare example is Twitch's proposed

"Brand Safety Score,"⁷ which attempts to score streamers based on aspects including how they govern their stream. However, it seems primarily for the benefit of advertisers and not accountability to members. Accountability can also come from the levels below. Processes for users to report, elect, or otherwise exert pressure on local administrators would exert higher accountability to lower levels (Frey & Schneider, 2020). A middle level supported from below can work autonomously from the level above and would be obliged to hold greater accountability to the level below for its decisions.

Limitations and Future Work

These dimensions are far from the only insights that scholarship on offline multi-level institutions can offer to online platforms. We aim to start a conversation that more closely links political science and public administration scholars to the platform governance community. Our dimensions are not measurable indicators but qualitative descriptions sufficiently general enough to be contextualized to each platform. We call for further studies to precisely operationalize these dimensions and carefully develop the broader design space's within-level, within-unit, and system-level dimensions. While we focus on middle levels in this work, examining units in other levels can also reveal valuable insights. For example, Jhaver et al. (2023) show how end users deploy platform-offered artificial intelligence (AI) tools to personalize the governance they experience.

We hope that a substantive research agenda can build upon the design dimensions we have presented here. Some research questions that represent important next steps include: How can platforms support governance units in their experiments with a diversity of governance strategies? Which design mechanisms and policies can help platforms foster healthy competition between different governance units? What initiatives, technical means, and tools can platforms develop to facilitate social learning? How can platform levels with overlapping jurisdiction collaborate effectively to ensure that no governance unit free rides on others' efforts or oversteps clearly defined bounds? How do different incentives and sanctions levied on governance units affect their operations in the long run?

Conclusion

Social systems are often complex, interdependent, and have a hierarchical structure (Simon, 1991). Understanding their complexity requires studying multi-level design, which has been explored in multiple disciplines. Multi-level frameworks are particularly important for examining large-scale social platforms, as evident in the ubiquity and diversity of middle levels that we find. In this work, we wrangle the many manifestations of multi-level online platforms into a generative design space. Drawing upon multi-level offline

governance theories, we examine our framework's design implications for more empowering and accountable governance. Unlike offline institutions, online platforms are an ideal laboratory for exploring the design space of potential multi-level governance architectures as designers can iterate quickly, base decisions on mass data, and compare communities directly.

Authors' Statement

All the authors have agreed to this submission. This article is not currently being considered for publication by any other print or electronic journal.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iD

Shagun Jhaver  <https://orcid.org/0000-0002-6728-7101>

Notes

1. For simplicity, we use the term “platforms” to describe these examples as a whole, although some of these are not traditionally considered platforms.
2. https://www.facebook.com/policies/pages_groups_events/
3. <https://en.wikipedia.org/wiki/Wikipedia:Administrators>
4. <https://mods.reddithelp.com/hc/en-us/articles/4415446939917-Reddit-Mod-Council>
5. <https://discord.com/moderation>
6. <https://discord.com/blog/announcing-the-discord-moderator-academy-exam>
7. <https://www.digitalinformationworld.com/2021/03/twitch-scores-users-through-brand.html>

References

Aligica, P. D., & Tarko, V. (2012). Polycentricity: From Polanyi to Ostrom, and beyond. *Governance*, 25(2), 237–262.

Ansell, C., & Torfing, J. (2016). Introduction: Theories of governance. In C. Ansell & J. Torfing (Eds.), *Handbook on theories of governance* (pp. 1–17). Edward Elgar Publishing.

Armitage, D. (2008). Governance and the commons in a multi-level world. *International Journal of the Commons*, 2(1), 7–32.

Aroney, N. (2016). *Types of federalism*. Max Planck Encyclopedia of Comparative Constitutional Law.

Bache, I., & Flinders, M. (2004). Multi-level governance and the study of the British state. *Public Policy and Administration*, 19(1), 31–51.

Bednar, J. (2008). *The robust federation: Principles of design*. Cambridge University Press.

Bednar, J. (2011). The political science of federalism. *Annual Review of Law and Social Science*, 7, 269–288. <https://doi.org/10.1146/annurevlawsocsci-102510-105522>

Berkes, F. (2010). Devolution of environment and resources governance: Trends and future. *Environmental Conservation*, 37, 489–500.

Blomquist, W., & Schlager, E. (2005). Political pitfalls of integrated watershed management. *Society and Natural Resources*, 18(2), 101–117.

Caplan, R., & Gillespie, T. (2020). Tiered governance and demonetization: The shifting terms of labor and compensation in the platform economy. *Social Media + Society*, 6(2), 2056305120936636.

Carlisle, K., & Gruby, R. L. (2019). Polycentric systems of governance: A theoretical model for the commons. *Policy Studies Journal*, 47(4), 927–952.

Chandrasekharan, E., Jhaver, S., Bruckman, A., & Gilbert, E. (2020). Quarantined! Examining the effects of a community-wide moderation intervention on Reddit. *ACM Transactions on Computer-Human Interaction*, 4, 1–26.

Chandrasekharan, E., Pavalanathan, U., Srinivasan, A., Glynn, A., Eisenstein, J., & Gilbert, E. (2017). You can't stay here: The efficacy of Reddit's 2015 ban examined through hate speech. *Proceedings of the ACM on Human-Computer Interaction*, 1(2), 1–22. <http://doi.acm.org/10.1145/3134666>

Chatterjee, M. (2023). *Social media's fragmented future*. <https://www.politico.com/newsletters/digital-future-daily/2023/06/22/social-medias-fragmented-future-00103238>

Chen, Y., Pereira, I., & Patel, P. C. (2020). Decentralized governance of digital platforms. *Journal of Management*, 47, 1305–1337.

Datta, S., & Adar, E. (2019). Extracting inter-community conflicts in Reddit. In *Proceedings of the international AAAI conference on web and social media* (Vol. 13, pp. 146–157). AAAI.

DeNardis, L., & Hackl, A. M. (2015). Internet governance by social media platforms. *Telecommunications Policy*, 39(9), 761–770.

Elazar, D. J. (1994). *Federal systems of the world: A handbook of federal, confederal and autonomy arrangements*. Longman Current Affairs.

Elazar, D. J. (2014). *Exploring federalism*. University of Alabama Press.

Emmet, E. R. (1964). *Learning to philosophize*. Penguin Books.

Fan, J., & Zhang, A. X. (2020). Digital juries: A civics-oriented approach to platform governance. In *Proceedings of the 2020 CHI conference on human factors in computing systems, CHI '20* (pp. 1–14). Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376293>

Folke, C., Hahn, T., Olsson, P., & Norberg, J. (2005). Adaptive governance of social-ecological systems. *Annual Review of Environment and Resources*, 30, 441–473.

Frey, S., & Schneider, N. (2020). Effective voice: Beyond exit and affect in online communities. *arXiv*.

Frey, S., & Sumner, R. W. (2019). Emergence of integrated institutions in a large population of self-governing communities. *PLOS ONE*, 14(7), Article e0216335.

Fukuyama, F., Richman, B., & Goel, A. (2021). How to save democracy from technology: Ending big tech's information monopoly. *Foreign Affairs*, 100, 98.

Geiger, R. S. (2016). Bot-based collective blocklists in Twitter: The counterpublic moderation of harassment in a networked public space. *Information, Communication & Society*, 19(6), 787–803. <https://doi.org/10.1080/1369118X.2016.1153700>

Gelcich, S. (2014). Towards polycentric governance of small-scale fisheries: Insights from the new “management plans” policy

in Chile. *Aquatic Conservation: Marine and Freshwater Ecosystems*, 24(5), 575–581.

Gillespie, T. (2018). *Custodians of the internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press.

Graeber, J. (2023). *Composable moderation*. <https://blueskyweb.xyz/blog/4-13-2023-moderation>

Grimmelmann, J. (2015). The virtues of moderation. *Yale Journal of Law & Technology*, 17, 42.

Hooghe, L., & Marks, G. (2001). Types of multi-level governance. *European Integration Online Papers (EIoP)*, 5(11), 1–32.

Huang, K. (2022). *What is mastodon and why are people leaving Twitter for it?* <https://www.nytimes.com/2022/11/07/technology/mastodon-twitter-elon-musk.html>

Ingram, D. (2023). *Reddit CEO slams protest leaders, saying he'll change rules that favor "landed gentry."* https://www.nbc-news.com/tech/tech-news/reddit_protest-blackout-ceo-steve-huffman-moderators_rena89544

Jhaver, S., Birman, I., Gilbert, E., & Bruckman, A. (2019). Human-machine collaboration for content regulation: The case of Reddit automoderator. *ACM Transactions on Computer-Human Interaction*, 26(5), 1–35. <https://doi.org/10.1145/3338243>

Jhaver, S., Ghoshal, S., Bruckman, A., & Gilbert, E. (2018). Online harassment and content moderation: The case of blocklists. *ACM Transactions on Computer-Human Interaction*, 25(2), 1–33. <https://doi.org/10.1145/3185593>

Jhaver, S., Zhang, A. Q., Chan, Q. Z., Natarajan, N., Wang, R., & Zhang, A. X. (2023). Personalizing content moderation on social media: User perspectives on moderation choices, interface design, and labor. *Proceedings of the ACM on Human-Computer Interaction*, 7, 1–33. <https://doi.org/10.1145/3610080>

Kurridl-Klitgaard, P. (2010). Exit, collective action and polycentric political systems. *Public Choice*, 143(3–4), 339–352.

Loughlin, J., Kincaid, J., & Swenden, W. (2013). *Routledge handbook of regionalism & federalism*. Routledge.

Marshall, G. R. (2015). Polycentricity and adaptive governance. In *Commons amidst complexity and change*, Working paper presented at the 15th biannual international conference of the international association for the study of the commons, Edmonton, Alberta. (pp. 1–19).

Marwick, A. E., & Caplan, R. (2018). Drinking male tears: Language, the manosphere, and networked harassment. *Feminist Media Studies*, 18(4), 543–559.

Masnich, M. (2019). *Protocols, not platforms: A technological approach to free speech*. <https://knightcolumbia.org/content/protocols-not-platforms-a-technological-approach-to-free-speech>

Matias, J. N. (2019). The civic labor of volunteer moderators online. *Social Media + Society*, 5(2), 1–12. <https://doi.org/10.1177/2056305119836778>

Matias, J. N., & Mou, M. (2018). Civilservant: Community-led experiments in platform governance. In *Proceedings of the 2018 CHI conference on human factors in computing systems* (pp. 1–13). Association for Computing Machinery.

McGillicuddy, A., Bernard, J. G., & Cranefield, J. (2016). Controlling bad behavior in online communities: An examination of moderation work. In *ICIS 2016 proceedings*. <http://aisel.aisnet.org/icis2016/SocialMedia/Presentations/23>

McGinnis, M. D. (2016). *Polycentric governance in theory and practice: Dimensions of aspiration and practical limitations*. Indiana University Bloomington.

McGinnis, M. D., & Ostrom, E. (2012). Reflections on Vincent Ostrom, public administration, and polycentricity. *Public Administration Review*, 72(1), 15–25.

Morrison, T. H., Adger, W. N., Brown, K., Lemos, M. C., Huitema, D., Phelps, J., Evans, L., Cohen, P., Song, A., Turner, R., Quinn, T., & Hughes, T. P. (2019). The black box of power in polycentric environmental governance. *Global Environmental Change*, 57, 101934.

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.

Ostrom, E. (1999). Coping with tragedies of the commons. *Annual Review of Political Science*, 2(1), 493–535.

Ostrom, E. (2010). Beyond markets and states: Polycentric governance of complex economic systems. *American Economic Review*, 100(3), 641–672.

Ostrom, V. (1973). Can federalism make a difference? *Publius*, 3(2), 197–237.

Ostrom, V. (1991). *The meaning of American federalism: Constituting a self-governing society*. Institute for Contemporary Studies Press.

Ostrom, V., Tiebout, C. M., & Warren, R. (1961). The organization of government in metropolitan areas: A theoretical inquiry. *The American Political Science Review*, 55(4), 831–842.

Peters, J. (2022). *More than two million users have flocked to mastodon since Elon Musk took over Twitter*. https://ground.news/article/more-than-two-million-users-have-flocked-to-mastodon-since-elon-musk-took-over-twitter_4ba335

Rajendra-Nicolucci, C., & Zuckerman, E. (2021). *Forum logic: Topical, with a side of toxicity*. <https://knightcolumbia.org/blog/forum-logic-topical-with-a-side-of-toxicity>

Rozenshtein, A. Z. (2022). Moderating the fediverse: Content moderation on distributed social media. *Journal of Free Speech Law*, 3, 217.

Scheuerman, W. E. (2004). Democratic experimentalism or capitalist synchronization? Critical reflections on directly-deliberative polyarchy. *Canadian Journal of Law & Jurisprudence*, 17, 101–127.

Schneider, N. (2021). Admins, mods, and benevolent dictators for life: The implicit feudalism of online communities. *New Media & Society*, 24, 1965–1985.

Schneider, N., De Filippi, P., Frey, S., Tan, J., & Zhang, A. (2021). Modular politics: Toward a governance layer for online communities. *Proceedings of the ACM on Human-Computer Interaction*, 5, 1–26.

Seering, J. (2020). Reconsidering self-moderation: The role of research in supporting community-based models for online content moderation. *Proceedings of the ACM on Human-Computer Interaction*, 4, 1–28.

Shaw, A., & Hill, B. M. (2014). Laboratories of oligarchy? how the iron law extends to peer production. *Journal of Communication*, 64(2), 215–238.

Simon, H. A. (1991). The architecture of complexity. In *Facets of systems science* (pp. 457–476). Springer.

Tait, A. (2016, September 2). #YouTubeIsOverParty: Will YouTube's latest change destroy the platform? *New Statesman*. <https://www.newstatesman.com/science-tech/>

2016/09/youtubeisoverparty-will-youtube-s-latest-change-destroy-platform

Thiel, A., Garrick, D. E., & Blomquist, W. A. (2019). *Governing complexity: Analyzing and applying polycentricity*. Cambridge University Press.

Tiihonen, S. (2004). *From governing to governance*. Tampere University Press.

van Zeben, J. A. (2013). *Research agenda for a polycentric European Union*. SSRN.

Watts, R. L. (1999). *Comparing federal systems*. Institute of Intergovernmental Relations, Queen's University.

Wright, A., & De Filippi, P. (2015). *Decentralized blockchain technology and the rise of lex cryptographia*. SSRN.

Zhu, H., Kraut, R. E., & Kittur, A. (2014). The impact of membership overlap on the survival of online communities. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 281–290). ACM.

Author Biographies

Shagun Jhaver is a social computing scholar whose research focuses on improving content moderation on digital platforms. He is currently studying how internet platform design and moderation policies can address societal issues such as online harassment,

misinformation, and the rise of hate groups. Jhaver is an Assistant Professor at Rutgers University's School of Communication & Information in the Library & Information Science Department.

Seth Frey is a cognitive scientist and computational social scientist who studies human decision behavior in complex social environments. His expertise is in computational approaches to self-governance and the cognitive science of the commons. He is a Professor in Communication at the University of California Davis and an affiliate of the Ostrom Workshop at Indiana University. He was a behavioral economist at Disney Research Zurich in Walt Disney Imagineering, a Neukom Fellow at Dartmouth College's Neukom Institute, and a student at the New England Complex Systems Institute. He earned a Ph.D. in Cognitive Science and Informatics at Indiana University in 2013, and a B.A. in Cognitive Science from UC Berkeley.

Amy X. Zhang is a computer scientist and human–computer interaction researcher studying the design of online systems for social interaction. Her current research is focused on building tools for social media users and communities to have more agency in curating their online social spaces. Amy is an Assistant Professor at University of Washington's Allen School of Computer Science & Engineering.