"This Place Does What It Was Built For": Designing Digital Institutions for Participatory Change

SETH FREY, Department of Communication, University of California, Davis, United States
P. M. KRAFFT, Oxford Internet Institute, University of Oxford, United Kingdom
BRIAN C. KEEGAN, Department of Information Science, University of Colorado Boulder, United States

Whether we recognize it or not, the Internet is rife with exciting and original institutional forms that are transforming social organization on and offline. Governing these Internet platforms and other digital institutions has posed a challenge for engineers and managers, many of whom have little exposure to the relevant history or theory of institutional design. The dominant guiding practices for the design of digital institutions to date in human-computer interaction, computer-supported cooperative work, and the tech industry at large have been an incentive-focused behavioral engineering paradigm encompassing atheoretical approaches such as emulation, A/B-testing, engagement maximization, and piecemeal issue-driven engineering. One institutional analysis framework that has been useful in the study of traditional institutions comes from scholars of natural resource management, particularly that community of economists, anthropologists, and environmental and political scientists focused around the work of Elinor Ostrom, known collectively as the "Ostrom Workshop." A key finding from this community that has yet to be broadly incorporated into the design of many digital institutions is the importance of including participatory change mechanisms in what is called a "constitutional layer" of institutional design. The institutional rules that compose a constitutional layer facilitate stakeholder participation in the ongoing process of institutional design change. We explore to what extent consideration of constitutional layers is met or could be better met in three varied cases of digital institutions: cryptocurrencies, cannabis informatics, and amateur Minecraft server governance. Examining such highly varied cases allows us to demonstrate the broad relevance of constitutional layers in many different types of digital institutions.

CCS Concepts: • **Social and professional topics** → **Systems analysis and design**; • **Applied computing** → **IT governance**; Law, social and behavioral sciences; • **Human-centered computing** → *Collaborative and social computing systems and tools.*

Additional Key Words and Phrases: computational social science; institutional design; institutional analysis; resource management; knowledge commons; digital institutions; digital democracy

ACM Reference Format:
Seth Frey, P. M. Krafft, and Brian C. Keegan. 2019. "This Place Does What It Was Built For": Designing Digital Institutions for Participatory Change. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 32 (November 2019), 31 pages. https://doi.org/10.1145/3359134

*Authors contributed equally; Author order randomized.
Authors' addresses: Seth Frey, Department of Communication, University of California, Davis, 469 Kerr Hall, Davis, California, 95616, United States, sfrey@ucdavis.edu; P. M. Krafft, Oxford Internet Institute, University of Oxford, 1 St Giles, Oxford, Oxfordshire, OX1 3JS, United Kingdom, p.krafft@oii.ox.ac.uk; Brian C. Keegan, Department of Information Science, University of Colorado Boulder, ENVD 201 1060 18th St. Boulder, Colorado, 80309, United States, brian.keegan@colorado.edu.

2573-0142/2019/11-ART32 $15.00
https://doi.org/10.1145/3359134

"Because when this place does what it was built for, then we all get a little closer."
—"Here Together", 2018 Facebook advertisement

INTRODUCTION

In April 2018, Mark Zuckerberg, founder and CEO of a major online social networking platform, testified to members of the United States Senate and House of Representatives on the role of privacy in his platform's design. This Congressional testimony was noteworthy, not only because it generated global coverage, trending topics, and remixable memes, but because it placed representatives of two historically significant institutional experiments—Facebook and the U.S. Congress—in the same frame. The ensuing frenzy of posturing, apprehensions, and socio-technical suppositions highlighted the need for more robust conversations about how large, complicated *digital institutions* can and should be designed.

Institutions are the rules, norms, constraints, routines, roles, and other structures that people use to organize their social interactions [83, 87]. Institutions operate through regulative (rule-setting, monitoring, sanctioning, *etc.*), normative (obligation, expectation, morality, *etc*), and cultural-cognitive (taken-for-granted, comprehensible, shared understanding, *etc.*) mechanisms [102]. As digital institutions, whose rules and structures are encoded or enforced through software, come to play an out-sized role in mediating human affairs [14, 73], it becomes imperative to understand the considerations involved in their design and evolution. The quote in the epigraph, taken from a 2018 Facebook ad campaign in the aftermath of the Cambridge Analytica scandal, is an important admission that digital institutions like Facebook are not doing what they were "built for" and implies that (1) these systems can be designed to embody alternative values and (2) the design of these systems leaves them susceptible to defection, attack, and other failure modes.

Digital institutions present at least two unique challenges compared to non-digital forms. First, the inhuman rigidity, precision, and efficiency of code produces a new type of power asymmetry between platform owners and users. Second, the technical skills needed to build digital systems has led to a small, homogeneous group of *de facto* technocrats becoming the ruling class of the digital world. These twin challenges compound each other: a small and homogeneous group will have limited ability to predict misalignments between available services and the needs of a diverse user base. At the same time, the precision of software enables unprecedented levels of literalism in monitoring compliance to protocols, while its rigidity creates a fragility to unanticipated situations and malicious re-purposing. These characteristics of digital institutions exacerbate incongruities between the narrow intentions of designers and the unanticipated uses and misuses by users.

Within the relevant areas of academic computer science, one of the dominant theoretically-informed approach to the design of digital institutions is epitomized in the work of Kraut, Resnick, *et al.*'s *Building Successful Online Communities* [71]. This paradigm relies on an individualistic, incentive-based, and top-down "psychological engineering" approach to digital institution design. In the present work, we re-evaluate this dominant approach in light of the serious challenges digital institutions have confronted since 2016 around polarization, propaganda, and content moderation. We introduce working within an institutional analysis framework to answer past calls from the area of social computing to engage with literature from the social sciences [116]. Concentrating on the norms and rules, specifically the formal rules that define how communities operate, we draw upon the participatory view of institution design advanced by political scientist Elinor Ostrom.

Ostrom's research showed that participatory deliberation and decision-making are crucial for allowing institutions to adapt to local contexts and achieve greater sustainability than what is possible through centralized, top-down planning. By contrast, typical work in digital institution design, such as that of "nudges" [112], "optimization" approaches of A/B testing, or other engineering-oriented approaches, has focused on laying out the design space for digital institutions with negligible attention to who can make design decisions about the institution and how those decisions are adjudicated. Authors in the top-down engineering tradition—even in drawing upon Ostrom in building parts of their framework—see "a moral imperative to create online communities that work well", regardless of whether it is "social engineering" or "paternalism" [98, p. 9]. In contrast, Ostrom recognizes that choices about how to scaffold democratic participation are critical for designing responsive and resilient institutions.

Although related areas have some traction, institutional design and analysis perspectives are rare in computer and information science and engineering (CISE) fields: there are only 13 results in the ACM digital library and 19 results in the IEEE digital library for "institutional design" and 7 and 28 respectively for "institutional analysis". The absence of institutional design and analysis frameworks in CISE research conversations is conspicuous given that CISE practitioners design and implement digital institutions that have profound influence over the social, psychological, political, and economic lives of billions of people. These platforms' engineers, scientists, and managers are overwhelmingly drawn from science and engineering backgrounds, where most receive no formal training in institutional design and analysis, and are unaware of the implications of specific resource management configurations or the availability of more robust alternative designs.

In the remainder of this paper, we translate one component of Ostrom's institutional design and analysis framework to the problems of contemporary digital institutional design: the design of an institution's "constitutional layer" that defines how a system changes itself and who is formally included in change processes. Like all institutions, digital institutions change in response to a variety of factors—imitating other institutions' successes, adapting to changes in the availability of resources, and generating new rules and relationships [2, 103]. Engaging diverse stakeholder perspectives by incorporating a broad base of participants into "constitutional layer" processes ensures that these changes are more responsive, legitimate, and sensitive to diverse and local contexts. We use three case studies to illustrate how Ostrom's institutional analysis and design framework can be applied to digital institutions, and to argue for the importance of participatory constitutional design layers for creating more accountable and resilient digital institutions.

BACKGROUND

There is a long tradition in human-computer interaction and computer-supported cooperative work exploring how values are embedded in software and technological artifacts [39, 40, 119]. In close alignment with the related areas of infrastructure studies [25, 41, 45, 75, 80, 91] and critical platform studies [46], we build on this tradition by focusing on information technologies like software, protocols, and platforms as artifacts of institutional processes and structures. In this section, we offer a definition of digital institutions and trace out three distinct traditions within the human-computer interaction and information systems fields for designing institutions: participatory design, engineering, and the commons. We outline the features of the prevailing paradigm of top-down digital institutional design, introduce the Ostrom Workshop's concept of constitutional layers as an important component for designing sustainable institutions, and summarize previous work using Ostrom's frameworks to analyze digital institutions.

Designing a digital institution involves tasks such as defining the transfer and sharing of resources, how the user community will be built and sustained, and the action space of participation and interaction. During the rise of digital institutions over the last two decades, a handful of approaches to these questions have been pursued by computer and information scientists in industry. The pre-theoretical approach stereotypical of technology firms and digital platforms has been "build it and they will come", relying on the practical experience and the precedent of previous platforms to guide design decisions, as well as the intuitions and agility of managers to implement design choices. A related approach is to make design choices based on profit maximization, in the extreme case without regard to factors such as community sustainability or health. Research towards this goal focuses on atheoretical design methods such as A/B-testing for maximizing engagement, and machine learning for behavioral prediction. But these pre-theoretical and atheoretical approaches are increasingly unable to stay ahead of the accumulating socio-technical strains that are consequences of basic design choices made at the foundations of these digital institutions.

Digital institutions

While institutions are the rules, norms, constraints, routines, roles, and other structures that people use to organize their social interactions, digital institutions are institutions whose rules or structures are at least partially encoded or enforced through software. Modern institutions vary in the degree to which they are digitized; digitization is altering many historical institutions while also creating entirely new institutional forms [28]. Institutions like government departments use digital systems to implement or support existing legal or socio-technical infrastructure: a citizen submitting taxes online is participating in a digital component of a larger institutional framework, and the form of these digital structures creates affordances (such as new ways to submit taxes) and inequalities (such as through the "digital divide" [52]). Social media platforms like Twitter are more conspicuously identifiable as digital institutions because all actions and interactions are mediated through computers. However, even Twitter is not a wholly digital institution: Twitter-the-platform is managed by engineers employed by Twitter-the-corporation that is itself embedded within layers of other non-digital institutions like publicly-traded corporations and employment law.

It is difficult to think of any major institution today that does not have a digital component. A digital institution could be any company which uses software to shape participant behavior (e.g. enterprise resource planning systems, attendance control systems). Our definition of digital institutions is therefore quite broad. We inherit this definition from the literature on institutions, both digital and not. There is a large overlap between digital institutions and the type of systems studied in computed supported cooperative work or social computing—digital institutions certainly facilitate both cooperative work and social computing. The distinction is that compared to treating "cooperative work" as an object of study in computer supported cooperative work [51], or treating "computing" as an object of study in social computing [90], the focus of digital institutions is on institutional structures themselves. For example, thinking about "computing" is often oriented towards thinking about particular intended functions, whereas thinking about the infrastructure of a platform can facilitate thinking about many different uses and functions, intended or unintended. Similarly, thinking about "cooperative work" leads to questions such as how computers make certain collaborative tasks easier, whereas thinking about institutions leads to open-ended questions with institutions centered as the object of inquiry, such as what kinds of work do particular institutional choices facilitate, or to what extent do the rules or design choices of a particular digital institution yield a high-level property like adaptability of the institution or democratic participation within the institution?

One of our goals is to highlight the fault lines that exist when single agents like Twitter-the-corporation engage in digital institution design unilaterally to make Twitter-the-platform. Many of the largest and most influential digital institutions, including the platforms run by Facebook, Google, and Twitter, lack meaningful mechanisms for bottom-up input on their rule- and decision-making. Such participatory rule-making processes are widespread in the sustainable institutions studied by Ostrom and her colleagues [84].

While a social media user with grievances about fake news or spam may be able to flag content, these reports typically do not cause changes in the rules of the parent institution. Users of social media platforms have limited recourse for feedback and grievances in these digital institutions. Platforms operate under "intermediary immunity" rules like Section 230 that limit regulatory oversight of these platforms [27]. Platforms are also managed by for-profit corporations that prioritize increasing shareholder value or delivering reliable engagement to their advertiser customers. Change in the design of these digital institutions is typically caused through either (1) indirect and formal actions by management or shareholders, (2) indirect and coarse mechanisms like developing new product features to deliver better engagement for advertisers, or (3) direct but coarse mechanisms like user collective action, employee protests, or hacktivist disclosures. In the absence of reliable mechanisms for more direct and bottom-up input on rule-making, the design of many digital institutions have instead been governed through indirect and top-down means.

The participatory design tradition

Participatory design is a central framework within the fields human-computer interaction and information systems. Originating in Scandinavia and West Germany in the 1970s with strong traditions of industrial democracy [3, 6], the idea of participation was grounded in the right of workers to have an influence on matters that concern them, in this case, their work. Participatory design was a reaction to the design and deployment of information technological tools of managerial control that solicited little input from workers [67]. Gärtner and Wagner [43] offer three "arenas for design and participation", in which actors engage in system design. The first arena ("Arena A") is "designing work and systems": the (more or less) concrete artifacts of hardware, software, and processes where work is done. The second arena ("Arena B") is "designing organizational frameworks for action": the spaces where actors meet outside of the work tasks themselves and negotiate their respective interests. The third arena ("Arena C") is "designing industrial relations": the broader legal and political environment where relationships are defined by regulations in response to constituencies and agendas. Early research in participatory design attempted to link Gärtner and Wagner's different arenas together, but more recent efforts have largely ceded the second and third arenas to focus on the first arena [42].

A review of participatory design projects identified five requirements for substantive participation in design processes: (1) access to relevant information, (2) possibility of independent conclusions, (3) participation in decision-making, (4) availability of participatory development methods, and (5) room for alternative technical and/or organizational arrangements [20]. Participatory design projects have wide variation in the construction of participation [49, 67]. At one extreme, worker participation is limited to access to their experiences with a goal of surfacing their needs as end-users. These projects are initiated by managers or design professionals, and workers have little control over the design process, influence over the direction, or ability to propose alternatives. This situation approximates much of contemporary user experience research in digital institutions: evaluating already-engineered products or legitimizing one managerial decision over another. At the other extreme, worker participation is central to the success of a project through analysis of possibilities, evaluation and selection of components, prototyping of technologies, and involvement in implementation and deployment. This extreme of fully participatory engagement is rare in practice and most participatory design projects typically involve some level of sampling and delegated representation. Wikipedia's rule-making environment, in which even the most fundamental policies are subject to on-going revision by any editor, captures this extreme of participatory design in a contemporary digital institution [10, 54, 65, 72].

The sustainability of participatory design projects is also an on-going tension. They are often small-scale, isolated from other parts of the organization, and contingent on the presence of researchers and design professionals to establish and maintain participatory structures [67]. A participatory design for useful systems emphasizes responding to known issues while one for infrastructuring emphasizes discovering unknown issues [24]. Management often re-asserts its control over design processes once the project ends and resources are withdrawn or active engagement in participatory processes wanes [20]. These shortcomings often reflect the narrow "Arena A" scoping of participatory design projects rather than incorporating long-term viability and engagement as design goals alongside more robust "Arena B" and "Arena C" mechanisms. Gärtner and Wager argue,

> "taking politics seriously in systems design requires designers not only to analyze existing actor networks but ultimately to redesign them in ways that help establish and maintain participatory structures." [43, p. 212]

To address these sustainability concerns, participatory design scholars emphasize the need to shift from "design-for-use" to "design-for-future-use" and creating the infrastructures of technical practice as well as social support to sustain on-going and future participatory design processes [33, 106]. Meaningful participatory design of "things" (Arena A) often requires robust participation in the design of the infrastructures and institutions around those things (Arenas B and C). This is the space of institution design.

What is the role of participatory design in systems like online communities and other digital institutions where concepts like "management" and "labor" are less clear? A shift towards thinking about participatory design and the public provides a new set of strategies for engaging in settings, participants, and authority dynamics that lack clear institutionalized divisions [11, 24]. Filtered through the lens of any given contemporary debacle around privacy or content moderation, the "Arena A" would be collaboratively exploring alternative technical infrastructure like newsfeed ranking or API access, the "Arena B" would emphasize settings and processes for employees and users to negotiate how to translate grievances into releases, and "Arena C" would emphasize the development of national policies and regulations requiring platforms create processes for responding to user grievances. Any of these alternative institutional configurations admittedly seems far-fetched given the significant gaps in all five of the participatory design requirements, not least because users lack any substantive bargaining power over platform governors outside of coarse employee or consumer feedback mechanisms like exit, voice, loyalty, and neglect [57, 120]. What would a substantive participatory design project of digital institutions look like and how could it be sustained? Ostrom's concept of constitutional layers for institutional design, which we define in detail later on, provides an under-appeciated interface with this participatory design tradition [79].

The engineering tradition

Today, the most dominant and successful paradigm in the design of digital institutions is what we term the "behavior change" or "psychological engineering" approach. Psychological engineering represents the design of social systems with a top-down goal-oriented analytic procedure that uses psychology and economics to shape user behavior with carefully structured incentives, and subtle informational nudges that leverage human cognitive biases to steer users toward preferred outcomes. Design choices in this paradigm can be made and implemented by a singular designer or administrator who has unilateral power over the institution and unquestioned legitimacy to exercise it. Participants are the wild cards in the institution, harder to control than the rest of its structure, and the goal of a design is generally to structure their decision environment so as to orchestrate user behavior toward some collective end. Psychological engineering approaches institutional change through the lens of iterative design: an institution is designed to satisfy a goal, its proximity to that goal is measurable, and it merits change to the extent that it is deviating from the goal. So after deploying a design, the designer measures and analyzes deviations, tweaks relevant portions of the decision environment, and deploys a refinement that more sensitively harnesses psychological and economic insights in service of the institution's design goals.

Challenge	Example design claim
Motivations	"Requests from high-status people in the community lead to more contribution than anonymous requests or requests from low-status members."
Commitment	"Highlighting a community's purpose and successes... can translate members' commitment... into normative commitment to the community."
Newcomers	"Entry barriers for newcomers may cause those who join to be more committed to the group and contribute more to it."
Regulation	"Moderation decided by people who are members of the community, are impartial, and have limited or rotating power will be perceived as more legitimate and thus be more effective."
Founding	"Ambiguity of scope for the community creates opportunities for adjustment and member ownership."

Table 1. Examples of design claims and their corresponding design challenges from [71].

Although certainly not the only representative of this approach, one of the most prominent examples within HCI of the psychological engineering tradition is the book *Building Successful Online Communities* (*BSOC*) [71] by Kraut, Resnick, Kiesler, Riedl, and their colleagues [8, 34, 96, 97]. Drawing on social psychological theories about motivation, commitment, and identity, the *BSOC* approach has directly informed the design of social media platforms like Facebook. *BSOC* seeks to "identify a wide variety of levers of change, features of online communities that can be deliberately and strategically chosen" [71]. We go into depth analyzing this work as a grounded example of how a behavioral or psychological engineering approach has been used in HCI and CSCW, even though many other systems-oriented or engineering-oriented top-down approaches have similar blind spots. We provide details on the philosophy, method, and recommendations of *BSOC* in order to contrast its design implications with those of work in the Ostrom tradition.

Within the psychological engineering approach, the central governor of the digital institution decomposes effective communities into aggregations of socio-technical functionality for engineering different dimensions of desirable community member behavior. *BSOC* describes eight "levers of change" that are social or technical configurations reflecting design decisions made by managers, designers, and members of online communities [98, p.6–8]:

Community structure. *How is the community organized?* Size, homogeneity, subgroup structure, and recruitment of members.

Content, tasks, and activities. *What kinds of activities does the community support?* Self-disclosed, imported, volunteered, professional content; interdependent vs. dependent tasks; social vs. immersive activities.

Selection, sorting, highlighting. *How can members find the information that is best for them?* Dividing community into spaces, highlighting good content, removing inappropriate content, and feeds or recommendations of relevant/related content.

External communication. *How can members communicate beyond the community?* Sharing content, migrating identities, and importing relationships embeds communities.

Feedback, rewards, sanctions. *How do community members receive feedback about their behavior?* Ratings, rewards, and sanctions provide informal or formal changes in status.

Governance. *How do online communities employ social roles, rules, and procedures to govern member behavior?* Handling newcomers and conflicts; rules for behavior by position; procedures for decision-making.

Access controls. *What controls can be imposed on its members?* Limits on membership and actions; selection and permissions of moderators.

Presentation and framing. *How does the community use examples to compare behavior?* Privileging good vs. hiding bad behavior; emphasizing similarities to other communities.

This is a generative framework for enumerating the types of digital institutional design decisions that need to be made, but is oblique about *who* makes these decisions or *how* these decisions are made. In response to critiques that online communities are not easily designed or controlled, Kraut and Resnick concede that people "cannot be shaped or programmed in the way physical materials or software can" but offer that "online communities can be designed and managed to achieve the goals that their owners, managers, or members desire" through a combination of social and technical configurations [98, p.6]. Typically, the default is implementing institutional designs without the time, cost, and risk of soliciting community members' input. Table 1 provides representative examples of Kraut, *et al.*'s [71] design claims.

Responding to the criticism about designing online communities as forms of "social engineering," Kraut and Resnick offer a one-page exposition on the "Morality of Design" weighing institutional design as a mechanism to "elicit individual behavior that benefits the community" against creating online communities that "make the communities more attractive for their members or more productive" [98, p.9]. Drawing on the concept of choice architectures introduced by Thaler and Sunstein's "libertarian paternalism" [109, 111, 112], Kraut and Resnick argue that encouraging compliance with rules through psychological and economic incentives—behavior-change mechanisms such as default design decisions—are sufficient to drive collective behavior within a digital institutions toward a community's goals.

This engineering approach is silent on participatory institutional design and has only a narrow conception of the importance of designing processes that include input from the members of the community affected by these design decisions. In the chapter "Regulating Behavior in Online Communities" the *BSOC* authors consider democratic recommendations like the need for community participation in rule-making, and explicitly invoke Ostrom to do so. They offer an overview of the reported benefits of democratic institutions and participatory mechanisms, with direct mappings of Ostrom's prescriptions to their own. But despite their consideration of democratic means, their imagined audience is nevertheless a central governor who is empowered to unilaterally implement designs. *BSOC*'s most prominent argument for fostering participation is fundamentally instrumental: participatory mechanisms raise perceptions of legitimacy by increasing compliance. They endorse participation here more to help administrators maintain stature than to sincerely incorporate user perspectives into iterations on their community's design. Because the legitimacy-increasing benefits of participation can be achieved with superficially participatory mechanisms—suggestion boxes, ticket trackers, and non-binding polls—the psychological engineering tradition's instrumental argument for participation can only motivate a superficial commitment to it.

There is a pernicious slippage that happens in the space between this kind of engineering approach's arguments about the morality of design defaults and the morality of readers. Kraut and Resnick explicitly "leave moral judgments—about which goals are worth designing for—to our readers" [98, p.9]. Throughout *Building Successful Online Communities*, the community designer is assumed to be a singular actor who can act unilaterally to implement design choices, not as a collective making a legitimate decision (*i.e.*, participatory democracy) or a governor representing its constituents' interests (*i.e.*, representative democracy). In the behavioral engineering framing, the moral judgments of the designer have ultimate authority. Leaving moral judgments about design goals as an exercise for the reader is itself a default design decision that not only defines substantive democratic participation out of the framework, but left a decade of digital institution design, about interactions involving billions of people, in the hands of digital governors who were rewarded for using newsfeeds to deliver behavioral change and engagement maximization. Why aren't moral judgments about the design of the community left to community members? The instrumentalist approach arguably contributed to the crises of computational propaganda in 2016 and afterward because the behavior change and engagement maximization mechanisms at the core of these platforms' design were not fully in control of either the users or their governors, but could be hijacked by malicious agents. Facebook, Twitter, YouTube, and other platforms are only beginning to re-evaluate their commitments to this engineering-based approach through more aggressive content moderation and a shift towards interpersonal and ephemeral "living room" rather than "public forum" messaging [62].

There are several potential rebuttals to our critique. The first is a matter of scope: the online communities that works such as *BSOC* serve are relatively small, voluntary groups sharing some common identity or bond, whose low stakes and modest scale make the "benevolent dictatorship" of a single designer/administrator a much lower-risk governance model [95]. Perhaps researchers in psychological engineering should not be held accountable for applications of their approach beyond the scales they imagined. But the adoption of psychological engineering by the largest and most consequential digital institutions suggests the need for researchers to tread especially carefuly when their empirical recommendations can be so quickly scaled to interactions among billions of people. The second rebuttal is ecological: because the costs of migrating between communities is so small, users can self-select, and the most popular online communities successfully reflect the morality of their members. The designers of these online communities are simply implementing the preferences elicited from their users to secure their engagement and prevent their defection to competitors. Two of our three cases consider ecosystems that pit communities against each other, and find varying levels of effectiveness for the competitive mechanism for ensuring subservience to users. The third rebuttal is anthropic: the difficulty of designing truly participatory institutions means that it is easier to advance the science of centrally engineered digital institutions, design guidance from which is harder to transfer to more participatory institutions. Overcoming the shortcomings of current participatory frameworks is one goal of research in the Ostrom Workshop.

In the remainder of the paper we will explore how ecological factors and proper constraints can counteract the failures of the engineering approach in certain circumstances. Conflicts inevitably invite questions about rule-making: "who makes the rules?" and "how can the rules be changed?" It is imperative that we privilege more participatory design traditions, especially for the most influential digital institutions. While online communities may once have been fringe institutions, the post-2016 crises implicating social media platforms demonstrates how deeply embedded digital institutions have become within powerful geo-political systems. The initial conditions on which authors in the behavioral engineering tradition make their technocratic recommendations may be appropriate for new and small communities, but a phase transition happens somewhere in the course of digital institutions' growth where technocratic decision-making is unable to efficiently aggregate information, elicit representative preferences, or make legitimate decisions.

The commons tradition

Research in the Ostrom Workshop brings a more general and democratic view of institution design, one with increasingly clear potential to transform digital institutions [93, 105]. This body of work has worked to unify the findings of a community of scientists and practitioners across anthropology, political science, economics, sociology, and other branches of the social sciences under a common framework. Elinor Ostrom's book *Governing the Commons* [84] offers commonly-cited "design principles" — The Design Principles for Community-based Natural Resource Management [22, 84] — that have been taken up in diverse settings like lobster fisheries, forest management communities, USENET forums, and ancient systems of pasture, turbary, and estovers. One classic success of the framework was a large-scale comparative study of over 100 farmer- and state-managed irrigation works in rural Nepal, which revealed the increased ability of self-governing watersheds to respond to local material, environmental, and cultural conditions [86, 110]. The Design Principles distill governance lessons from ecologically diverse institutions around the world, some of which are centuries old, isolating elements of successful resource systems and providing criteria for diagnosing a community's resource or governance problems. They are the most well-known contribution of the Ostrom Workshop outside that community, and have been used by many to analyze digital institutions:

User boundaries. Clear boundaries between legitimate users and nonusers must be clearly defined.

Resource boundaries. Clear boundaries are present that define a resource system and separate it from the larger biophysical environment.

Congruence with local conditions. Appropriation and provision rules are congruent with local social and environmental conditions.

Appropriation and provision. The benefits obtained by users from a common-pool resource (CPR), as determined by appropriation rules, are proportional to the amount of inputs required in the form of labor, material, or money, as determined by provision rules.

Collective-choice arrangements. Most individuals affected by the operational rules can participate in modifying the operational rules.

Monitoring users. Designated monitors who are accountable to the users monitor the appropriation and provision levels of the users.

Monitoring the resource. Designated monitors who are accountable to the users monitor the condition of the resource.

Graduated sanctions. Appropriators who violate operational rules are likely to be assessed graduated sanctions (depending on the seriousness and the context of the offense) by other appropriators, by officials accountable to the appropriators, or by both.

Conflict-resolution mechanisms. Appropriators and their officials have rapid access to low-cost local arenas to resolve conflicts among or between appropriators and officials.

Minimal recognition of rights to organize. The rights of appropriators to devise their own institutions are not challenged by external governmental authorities.

Nested enterprises. Appropriation, provision, monitoring, enforcement, conflict resolution, and governance activities are organized in multiple layers of nested enterprises.

Although still not widely appreciated in the standard canon of human-computer interaction, the Ostrom Workshop's approach to institutional analysis has been productively applied to cases emphasizing self-government, governance and resource management, and designing participatory mechanism into online communities. Ever since Peter Kollock's analysis of USENET [68], resource management perspectives on digital institutions have had a small foothold. Major contributions have been *Knowledge Commons* by Hess and Ostrom [56] and *Internet Success* by Schweik and English [101]. A whole literature has emerged around institutional analyses of the archetypal knowledge commons, Wikipedia [35, 36, 54, 115], as well as peer production generally [9, 104], with a later contribution by Hess and Ostrom analyzing online bioengineering databases [55]. Others have used the Workshop's resource management perspective to investigate loot distribution norms in the game World of Warcraft [99, 107], self-hosted community servers [38], and online "dark institutions" like software pirate exchanges and hacker collectives [1, 53]. There is growing interest in the potential of Workshop-style institutional analysis to serve digital institution design [26, 92–94], but explicitly design-oriented work is rare, and even less research has focused on the role of constitution-level rulemaking in digital institutional design. Intentional design of the constitutional layer facilitates participatory design similar to how it is practiced in HCI [79], extended as an ongoing process. A focus on the constitutional layer emphasizes that digital institutions can be built with democratic values embedded in them, both in their governance and evolution, not just designed initially in a participatory manner.

DESIGNING FOR CHANGE

A key premise of institutional analysis within the Ostrom Workshop is that digital institutions must have policies and processes that define how the institution changes in order to be responsive to the shifting environments they inhabit. Critically, these processes must provide mechanisms for "lower-level" agents to participate in rule-making and enforcement. These agents are typically the first to note anomalies as well as the first to bear the brunt of consequences when institutions become de-coupled from their environments, so their input is critical for both the sustainability and legitimacy of new institutional arrangements.

Levels of choice in institutions

Rules and norms are difficult to classify, but successful classification schemes provide vital insights into an institution's structure. The basic component of an institution is the institutional statement, a linguistic description of the institution as it is used. Institutional statements can span several degrees of normativity, from suggestions, to unenforced expectations, to inviolable rules with consequences for noncompliance. They can be formal and informal, they can define and govern many different kinds of units of analysis, and they can apply to complex overlapping subsets of participants. For example, under the Ostrom Workshop's rule taxonomies, institutional statements can be classified by their "level" or scale of focus:

- *Operational rules* concern the lowest-level, most mundane behaviors taken by system members, as constrained by its collective choice processes. On the online marketplace Amazon, operational rules define the elementary actions that each type of user can perform, such as posting or purchasing a good.
- *Collective rules* concern the behaviors that the institution performs through the agents authorized to represent it, and functions mostly to define affordances at the operational level. Amazon's code defines the market context within which agents operate, and through which collective action processes drive the price mechanism.
- *Constitutional rules* concern the space of actions by which the collective choice level is changed and, in the broadest sense, the "meta" rules by which the system changes itself. Amazon is by no means a democracy, but like any large corporation defines internal research and review processes under which it evolves.

The constitutional layer defines how a system changes itself and who is formally involved in the process of change. For example, as we discuss in more detail below, constitutional layers that are beginning to appear in experiments in cryptocurrency governance focus on issues such as rules for off-chain dispute resolution and processes for adjusting those rules. These levels help define institutional structure in terms of agent capabilities. For example, on Wikipedia, there are operational rules determining how edits should be performed and evaluated, there are collective action rules defining how conflicts should be resolved, and there are constitutional rules outlining the encyclopedia's extensive body of policy for specifying how policy changes [10, 15, 16, 65, 72, 78].

Case	Architecture	Environment	Domain	Participants	Location
Cryptocurrencies	Decentralized	Online	Economic	Mixed	International
METRC	Centralized	Physical	Economic	Professionals	Colorado, United States
Minecraft	Mixed	Online	Social	Amateurs	International, United States

Table 2. A table showing salient dimensions along which our cases vary. By choosing a varied set of cases we emphasize the relevance of constitutional layers to many different types of digital institutions.

A system without a constitutional level—without formalized change processes—will either not change at all or will be susceptible to unstructured or informal drift-like change that, by reducing its similarity to the formal rules that describe it, undermines its accountability to all but the most powerful [37].

The need for a constitutional layer

Elinor Ostrom's early thought explicitly ties institutional self-modification to cybernetic theories of system fit and responsiveness to the outside environment [85]. In the general terms first offered by W. Ross Ashby, a system must be able to change to match the degree of change of its environment [4]. The idea of constitution-layer rule-making articulates this in the context of policy with explicit arguments, taken from Ostrom's *Design Principles* [22, 84], that successful self-governing institutions are able to maintain environmental fit, maintain structure at several scales, and include the full range of stakeholders in decision making and meta-decision making.

Defining institutional design so that it includes a constitutional layer has the side effect of expanding the domain of the theory to include participatory systems in which important governance-related decision-making incorporates agents beyond the platform creator and its appointed moderators. A related benefit of a constitutional layer of policy is that it provides a natural place for designers to express values and commit to ethics. We argue that, for an institution to be participatory in a meaningful sense, it must provide all agents some avenue for constitution-level action. Ticket systems, polls, and other operational feedback schemes—such as those endorsed by the behavioral engineering approach—do not faithfully implement participation unless they are provided an explicit formalized role in constitutional change processes. The system is most participatory that endows all agents with unmediated access to constitution-level choice.

With these considerations in mind, the central arguments of our work are that digital institutions should explicitly define their change processes, that they should define those processes to give several types of stakeholder a stake in meta decisions, and that major problems faced by prominent digital institutions are due to their failures to enact these recommendations.

CASE STUDIES

In the remainder of this paper, we set out to illustrate the importance of constitutional rules in digital institutions through three case studies: cryptocurrencies, a U.S. state-level cannabis monitoring system, and the Minecraft server ecosystem. We selected these narrative cases based on the areas of expertise of the authors of the present paper, after identifying common conceptual ground between the disparate institutional forms underlying each. We draw on a high-level qualitative understanding of these systems resulting from our own prior research in each domain [38, 63, 64, 66, 70]. We avoid more familiar cases such as Facebook or Twitter intentionally: (1) to highlight how ubiquitous and varied different forms of digital institutions are, (2) to avoid reinforcing the dominance of a few corporate actors in conversations about the design and regulation of empowering digital institutions, and (3) to demonstrate the vast potential of constitutional design frameworks to benefit

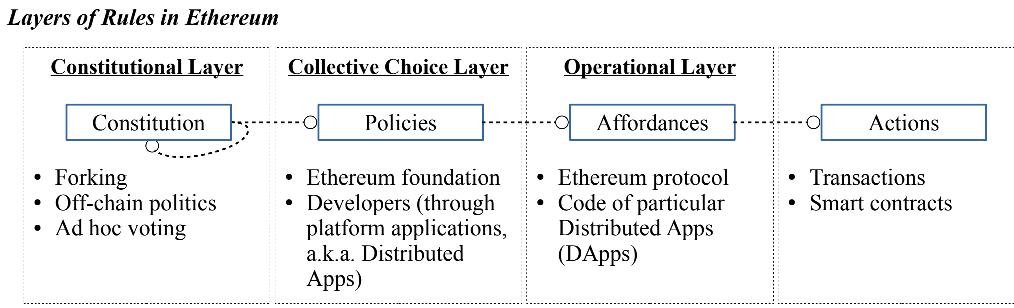

Fig. 1. Examples of layers of rules in the cryptocurrency Ethereum. Constitutional mechanisms are not developed explicitly in Ethereum, so the processes of change in collective choice and operational rules are coarse and chaotic.

other emerging sectors of digital society. The figures in each case provide broad context for the cases in the format of the Ostrom Workshop's institutional analysis and design framework, emphasizing where the constitutional layer plays a role in the broader institutional landscape of each case. We focus on issues of constitutional crisis, which are crises in an institution that relate to rules (or lack thereof) at the constitutional layer. These systems are unified by virtue of all being contemporary digitally mediated institutions, but they vary along other dimensions (see Table 2).

In our first case, we see how cryptocurrencies have struggled with protocols that are incapable of making changes to themselves that are proportional to changes in their environment. We focus in particular on the few most prominent currencies developed in the anarcho-capitalist tradition of libertarian thought, which holds that healthy markets can emerge spontaneously with an absolute minimum of institutional structure, and that markets are sufficient to provide all public goods and services. In our second case, we see ad hoc constitution-level participatory change in an unwieldy tool for monitoring compliance with cannabis regulations. In our final case we see a relatively successful market implementation of "participation-like" governance in the ecosystem of amateur Minecraft servers. After describing each case, we introduce a constitutional crisis that it faced: an event or events beyond what its constitution designed it to adapt to. We then discuss for each case available theoretical insights from relevant bodies of theory, focusing on the theories of the Ostrom Workshop. We accompany each case with a figure decomposing its institutional structure into the categories of the three levels of choice.

On a methodological and definitional note, institutions are not always neatly bounded. Institutional structures are in fact invariably multilayered, polycentric, and intertwined with each other. Our approach to bounding the digital institutions we studied to scope our analyses involved selecting pieces of software—cryptocurrency code, METRC code, Minecraft code—and then conducting analyses that emanated from there. In each case there were more rules at play than just the code involved, and those additional institutional structures came into our analysis as bearing upon the immediate functions and uses of the code our analyses centered upon.

Case Study 1: Cryptocurrency governance

Cryptocurrencies like Bitcoin and Ethereum are at the frontier of digital institution design, and governance is a central part of the cryptocurrency conversation [5, 13, 32]. A cryptocurrency is a digital system that maintains a "distributed ledger"—in practical terms, a database—of financial transactions. Publicly recording transactions and tracking who owns what helps give cryptocurrencies the properties of a currency, making them viable competitors to familiar currencies such as the dollar, euro, and yen. Advocates imagine replacing fiat currencies with cryptocurrencies as a digital, decentralized alternative to government-controlled treasuries or banks. Cryptocurrencies are decentralized in the sense that the record of cryptocurrency transactions does not reside in any single location, and updates to the log of transactions occur through protocols that involve computation distributed across many computers owned by different people and organizations. Cryptocurrency transactions and computations are pseudoanonymous in the sense that cryptocurrency accounts generally do not have identities explicitly associated with them, although de-anonymization is possible [82]. Cryptocurrencies face the institutional challenges that confront any monetary system, along with the additional challenges of governing a digital, decentralized, largely anonymous system. We view each cryptocurrency as a separate circumscribed digital institution, although often with similar institutional challenges, and apply our proposed institutional analysis framework in this capacity. We review two design challenges confronting this emerging class of digital institution: the "scaling debate" and the idea of a proof of stake, both described in more detail in the following. The way these issues have unfolded challenges the techno-libertarian narrative that cryptocurrencies offer "trustless" mechanisms—mechanisms that do not rely on trusting a power centralized authority—that can operate without institutional scaffolding and without constitutional mechanisms beyond market incentives, a narrative that has been enabled by the engineering approach to digital institution design. After introducing governance problems typical of today's cryptocurrencies, we discuss the role that constitutional-level structure plays in perpetuating these problems.

Constitutional crisis. Fundamental questions about cryptocurrency protocol standards have in important cases been left to the sociopolitical rather than the market realm. Since the sociopolitical realm has in many cases been assumed away by cryptocurrency designers, the systems they engineer are resistant to adaptive changes, whether minor and prescient or major and existential. Two representative design challenges are Bitcoin's scaling debate and the debates around proof of stake. The difficulties around how to resolve these debates demonstrate the need for institutions around each coin that, in many ways, undermine the trustlessness that is supposed to distinguish cryptographic currencies from fiat currencies [114, 117]. We deem these challenges "constitutional crises" because of our assessment that they result in part from a lack of attention to the constitutional layer of institution design.

The scaling debate centers around the size of Bitcoin blocks. Each Bitcoin block is a collection of transactions, and the total number of transactions per block is limited by the size of the blocks. The current maximum allowed block size is just a few megabytes. As several simultaneous transactions are proposed, their order of execution is computed according to a bidding system, and Bitcoin miners' profits from these auctions are drawn from these bids. Miners are agents in the currency ecosystem who perform the costly work behind adding each block of transactions to the blockchain, and so they have a profit incentive to keep block sizes small, because that keeps transaction fees high. However, small block sizes are bad for buyers and sellers because new blocks are only created on the blockchain on average every 10 minutes. If many people are trying to transact, congestion of transactions in the blockchain can quickly occur, which in turn leads to people being unable to get transactions processed, potentially for days or more. These conflicting incentives between consumers and miners have created an enormous debate around this issue, and miners' successes at keeping the block size small have severely reduced Bitcoin's competitiveness in the face of newer and more owner-friendly currencies. The Bitcoin scaling debate highlights how multiple actors with different incentives in the cryptocurrency ecosystem create institutional tensions.

A second design challenge faced in cryptocurrencies centers around how to securely verify blocks. Most contemporary coins operate with a "proof of work" model in which miners allocate computing cycles to find hashes that satisfy a criterion. This mechanism was originally argued to be fair in principle because anyone with a commodity laptop can be a miner and contribute to and benefit from the currency. In practice, the escalating costs of energy and the emergence of specialized hardware for cryptocurrency mining have dramatically reduced the profitability of mining to all but a concentrated handful of specialist agents. The people who now have the ability to become profitable miners are those who have the existing capital and mobility that allows them to establish large computing centers in places with relatively cheap land and electricity. These problems have driven the development of "proof of stake" mechanisms, which use monetary bids as an alternative to computing cycles. The idea of proof of stake is that an "ante" buys you into a block. If you cheat by creating a false block, you risk losing the money you have bid. Proof of stake relies on this financial incentive for the correctness of the protocol. The debate around proof of stake is what the minimum bid should be. If the bid is too high, fewer people will have the resources to participate in validating blocks, which undermines the vision of decentralization. If the bid is too low, people have no incentive to be honest.

In both, a finite valuable resource strains agents, but the system prevents them from effectively organizing at the collective level to implement system changes. These problems should not be insurmountable, but a consequence of designing for trustlessness is that currencies are essentially immutable after release: the code supporting a coin has no constitutional layer. Each change to a currency's protocol, no matter how minor, requires a "fork" of an existing currency, in which the code from the old coin is copied, edited, and released as a new coin. Holders of the old coin are then all encouraged to simultaneously divest from the old into the new. If a coin's community cannot reach consensus on the necessity of a fork, it bifurcates, and the coin has failed to fulfill the primary function of money: to provide a standard unit of exchange.

Absent formal constitutional mechanisms, what have these types of debates looked like in practice? A third debate provides a dramatic example. In 2016, a hacker diverted millions of USD from The DAO (Decentralized Autonomous Organization) project to an unknown agent's account. In response, after weeks of "off-chain" debate through channels such as blogs and web forums, the Ethereum Foundation actually called coinholders to vote to undo the theft's damage by "reversing time" in the blockchain ledger of Ethereum through a "hard fork" of the system [31, 118]. In other words, the response to the crisis of the hack was an *ad hoc* coordinated election process that occurred outside the Ethereum protocol, and was in fact coordinated by high-prestige founding members representing an intervening central authority. The community solved their problem through precisely the dynamics that blockchain institutions position themselves in opposition to. In the process, the debate's resolution revealed the existence of an elite group with disproportionate influence over the system's governance. Facing competing pressures from different parts of the community, such as miners holding a large amount of power, the Foundation ultimately acknowledged the need for, and their power to implement, constitution-level institutional structure. Throughout the debate around this issue, a purist techno-libertarian wing opposed the hard fork on the grounds that the code of the system, compromised as it was, should be treated as the ultimate authority; the code was agreed to by all coinholders prior to investing, the bug was in the code, so accepting the code's flaws is a necessary condition of respecting trustlessness and decentralization. In an effort to rebut that argument, the leaders of the voting initiative countered that this one-time top-down recourse to democracy was not a threat to the ideals of the community, as future attempts to rely on such a political solution would generally be ineffective at resolving the more quotidian, fine-grained policy issues that, as we argue, an institution must be able to resolve (such as properly parameterizing block size and proof of stake): "Imagine how hard it would be to get a patch approved, pushed out to mining pools and to get them to reach consensus about a less clear-cut issue. It's just not happening in most circumstances" [113]. While cryptocurrencies have shown an ability to change in extreme circumstances, their reliance on the heavy-handed forking mechanism for even the most mundane design decisions translates to a functional lack of constitution-layer rules, and corresponding adaptivity. Without "meta" policies for making granular course-corrections to existing policies, individual cryptocurrencies are unable to meet the demands of a changing world, overcome technical or sociotechnical system failures, or keep up with innovations introduced by competing currencies.

Comparing theoretical guidance. What are the aspects of the structure of cryptocurrencies, as digital institutions, that are creating the conditions for the long, grueling debates that have haunted blockchain communities? How might the engineering tradition analyze these issues, and how might the participatory design tradition?

The behavior engineering tradition focuses on nudges or incentives as solutions. As such, cryptocurrencies stand out as a success of the paradigm. A behavioral engineering analysis might suggest incentivizing stakeholders to converge on a block size or minimum bid. But the *status quo* in each case is itself an unexpected outcome of perverse incentives put in place by the engineering mindset, the same mindset that designed these systems to be immutable. Without a theory of how institutions change, behavioral engineering is doomed to react to change, rather than plan for it.

In contrast, Ostrom's participatory lens allows us to diagnose a larger issue at play. This perspective encourages us to attend to the lack of rules for resolving debate and making amendments. Because there are generally not clear cryptocurrency protocols for resolving disputes around cryptocurrency protocols, questions around these issues end up bogged down in open-ended debates, ad hoc damage control, extralegal maneuvers, and power plays.

Under the "levels of choice" framework, the elements of a cryptocurrency's structure can be assigned different levels of analysis, as we do for Ethereum in Fig. 1. The basic limited resource being managed is the cryptographic coin, and the institution is defined around creating the idea of a coin, making coins valuable, tracking their flow through the economy, and ultimately making them ownable and exchangeable. At the operational level the protocol defines, in code, the possible actions that different kinds of agents can perform on or around coins: how transactions are made, bundled, and processed. The collective choice layer is the heart of a cryptocurrency, defining how individual actions will be aggregated to implement the price mechanism of microeconomics and bring a currency into existence. At the third, constitutional layer, formal mechanism suddenly becomes scarce. Protocols are fixed by design to adapt to only a carefully circumscribed range of situations, and mechanisms for changing the protocol have emerged only down the line, as communities have realized that their currencies must be more adaptable than they have been designed to be. Thus the protocol is generally not involved in the constitutional layer because it is not defined to have one. One exception to this generalization is forking, a coarse mechanism that is unsuited to regular course corrections, and that constantly invites the risk of fracturing a community. In practice, sudden internal and external perturbations have led communities to demand change outside what formal mechanisms could accommodate. Consequently processes for constitutional change have emerged through political processes outside the protocol, on mailing lists and forums, managed behind the scenes by economically or politically influential members of the currency, usually developers, and often with the blessings of miners and other powerful stakeholders.

Our discussion and analysis has revealed how several different types of users and stakeholders within the cryptocurrency ecosystem have differing degrees of access to levers of institutional power. A constitution does not only specify change processes, it can also lay out the values that a cryptocurrency is meant to represent, provide a basis for arguments to justify specific institutional changes, and define the specific powers of particular actors and institutional roles. Institutional structures for debating and facilitating change can be formulated in a constitutional layer. Without robust systems for resolving debates, currencies increase their exposure to collective action problems that endanger all stakeholders.

From the perspective of the democratic tradition of digital institution design, cryptocurrency communities are being forced to accept that it is inevitable that human trust or human power will have a role in complementing the shortcomings of formal governance protocols [5, 74]. The DAO symbolized an interest within the community of embedding governance structures within the protocols of cryptocurrencies themselves, having code functioning as law [31]. However, these efforts at governance have proven to not go far enough. Mechanisms of change could in part be implemented in code, but processes for facilitating human debate and interaction regarding these protocols must not be ignored. We propose that in a system that is robust, "off-chain" politicking will never cease to exist. Moreover, off-chain debates—whether they are public or clandestine—are best resolved through formalized processes that allow more nuance in collective outcomes than wholesale acceptance or rejection of an entire currency.

Of course, there is another way that cryptocurrencies can be seen as implementing a constitutional layer. From above the perspective of any individual currency, at the level of the ecosystem of competing currencies is the market for currencies, the libertarian implementation of constitutional rulemaking: currencies compete for coin holders, and coin holders vote for the best protocol "with their feet" by divesting from undesirable coins and investing in effective ones, such as occurred after The DAO hack. As technologies advance, and lessons are learned, new coins replace the old ones, and, at the ecosystem level, the change that constitutions provide occurs naturally. This is the key claim of anarchocapitalist thought: markets can implement the key functions of government. Following the democratic tradition, we dispute this claim. It is unlikely that "meta-market" pressures are sufficient to implement or guarantee constitutional-layer rules about rulemaking. Competition may be sufficient to find prices quickly and efficiently within a currency. But special features like high switching costs and network externalities impose formidable inefficiencies, posing a serious obstacle to the effectiveness with which market competition can implement constitutional change in the ecosystem of currencies [18, 76, 81]. Therefore, cryptocurrency developers must not just write code, but also detail the principles, values, and rules of off-chain political processes. These considerations are now at the forefront of recent cryptocurrency and blockchain research, and experiments in constitutions are occurring (e.g., [21, 60]). Based on the insights of the Ostrom Workshop, our perspective is that the most successful constitutions will be those that prioritize participatory change, and take other measures to support adaptiveness to local conditions.

Case Study 2: Cannabis informatics

For our second case study we examine the digital infrastructure involved in the legalized cannabis markets within the United States. Recognition of the medicinal benefits, economic potential, and racial disparities in law enforcement have all contributed to a dramatic shift in attitudes and policy towards legalizing cannabis [44, 100]. Information technologies are playing a central role in the regulation of emerging recreational cannabis markets. While several states in the U.S. have operated regulated marketplaces for recreational cannabis since 2014, the production, distribution, and consumption of cannabis remain federal crimes in 2019 that carry significant penalties. The legal rationale against federal intervention in these state markets is the "closed loop theory" that requires demonstrating cannabis is neither being diverted into the black market nor crossing state lines [19, 59]. A central component of these new legalized markets are government-controlled

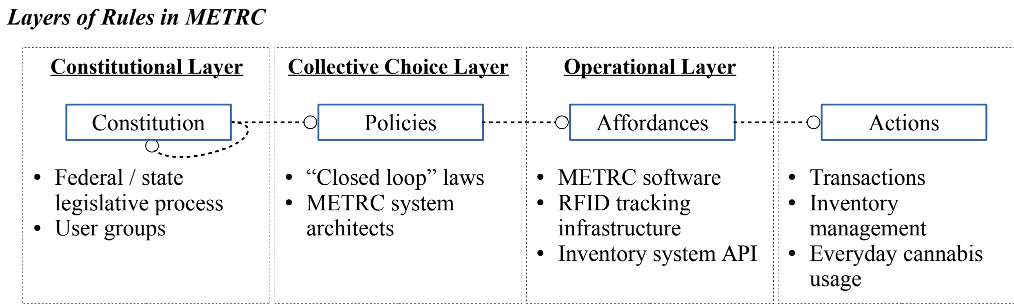

Fig. 2. Examples of layers of rules in the METRC cannabis informatics system. Overarching change processes involve surrounding U.S. legal infrastructure and guidance from user groups. The introduction of user groups provided a bottom-up constitutional mechanism that allowed METRC to address many unanticipated problems.

"seed-to-sale" inventory tracking systems that collect and store detailed data about each transaction in the supply chain from planting to retail sale. The creation of an entirely new market commodity—recreational cannabis—at a time of vast digitization has produced a peculiar institution, an ideal case for the place of participation in digital institution design. The history of this system, and its responses to market and cultural shifts, demonstrates the importance of designing digital institutions with adaptability and wide-ranging participation in mind.

The Marijuana Enforcement Tracking Reporting Compliance (METRC) system plays a central role in Colorado's $1.5 billion annual cannabis industry. METRC is a "regulatory compliance system" licensed by the State of Colorado's Marijuana Enforcement Division (MED) from Franwell, Inc., a privately-held, supply chain technology services company that offers "internet-deployed applications and services" like RFID technologies, technologies it originally developed to improve "track and trace visibility" for perishable food and pharmaceutical firms. Under the legal framework decriminalizing recreational cannabis, all medical and retail marijuana businesses in Colorado are required "to use METRC as the primary inventory tracking system of record". METRC was built "by regulators specifically for oversight" to provide "the necessary visibility for adherence to rules, regulations and statutes" to fulfill the legal demands of the "closed loop theory" [77].

METRC has two sides: an industry side that is "used to report the required events and information" and a regulatory side "used for enforcement and compliance monitoring". The industry side consists of RFID-enabled tags and barcodes attached to every plant and derivative product that exists, for tracking "from seed to sale". Regulatory compliance begins with tagging seedlings and clones, vegetating and flowering plants, trimmings, business-to-business transfers and distribution, and final sale at a retail location. The inventory and status is entered with specialized tag scanners and web user interfaces into a central database. METRC allows regulators to perform live audits to trace the provenance of every single product in the state-wide marketplace. They can do so using METRC's own built-in reporting tools, or secondary and tertiary data analysis systems that draw from METRC's APIs.

Constitutional crisis. By reducing compliance to a problem of keeping all activity in-state, the designers of this digital institution assumed that regulating the recreational cannabis market was isomorphic with the supply chain management problems faced by the manufacturers of perishable foods or pharmaceuticals—the industries Franwell previously served. This assumption ultimately drove a wedge between the capabilities of the system and the unique demands of a complex, statewide market. "Fit to local conditions" is one of Ostrom's principles of successful self-organizing resource management institutions. Another principle is "fit to the needs of stakeholders." METRC did not account for the market's many types of stakeholder (*e.g.*, growers, distributors, producers, testing facilities, retailers, *etc.*), the unusual regulatory constraints (*e.g.*, traditional testing facilities like professional laboratories and research universities withheld their services to avoid jeopardizing their federal accreditations), alternative business models, and the cannabis industry's difficulties securing basic economic protections (*e.g.*, the inability to access FDIC-regulated banks or SEC-regulated capital markets).

Consequently, METRC was initially ill-prepared to fulfill the market's demands. There was an avalanche of problems [108]: the MED regulatory agency mismanaged the roll-out of a predecessor system for medical cannabis [47], there were bottlenecks in manufacturing and distributing enough RFID tags for growers and producers [61], the central database had poor usability and uptime, preventing users from uploading their inventories as required by law [48], and innovative new products and business models could not be made to fit into METRC's models for logging transactions or ontologies for classifying products [17, 121]. As METRC was introduced to regulate a legally precarious marketplace, its fumbled launch jeopardized the larger institutional experiment of legalizing recreational cannabis in Colorado. At the core of this crisis were decisions about the design of a new digital institution and the absence of mechanisms for formally including stakeholders on the ground into its change processes.

In a remarkable turn for government information technology deployment, Colorado's MED was eventually confronted with the scope and stakes of its failures, and committed to a participatory process for revising the design of the system following its launch. Growers, retailers, regulators, and technologists participated in a series of "User Group" meetings to identify the fault lines between the legislatively-mandated affordances, and emerging practices in the market. A significant development was the introduction of a more robust API that allowed METRC to fulfill its statutory obligations to be a central inventory tracking system while also providing producers and retailers greater flexibility to develop alternative approaches for entering, representing, and retrieving data with secondary and tertiary data systems. At present, this API has end-points for employment, facilities, harvest, items, lab tests, packages, patients, plant batches, plants, rooms, sales, strains, and transfers: an ontological panopoly not recognizably akin to within-firm supply chains for perishable foods or pharmaceuticals, but necessary for the management of a complex ecosystem of stakeholders and resources. METRC stabilized following the development of these user group feedback mechanisms, enough that it has become the regulator-mandated seed-to-sale tracking platform in twelve other states' legal medical or recreational cannabis markets.

Comparing theoretical guidance. The structure of the METRC system, as a governance institution, can be seen as attending to several layers within the Ostrom Workshop's levels of choice (Fig. 2). At the operational level are the actions that METRC's client software provides producers and regulators for registering and tracking cannabis products through the system, along with the associated RFID infrastructure. Being a bureaucratically managed system, METRC's collective action layer is thin, as anything analogous to collective action is going to take the form of action by a central planner, imposed by strongly-enforced regulation, as informed by information aggregated from producers and other low-level agents in the system. As such, the collective choice level is largely concentrated in the components of the software that link operational-level actions and present them in aggregate form to regulators and law enforcement. Just as METRC's bureaucratic nature predisposes its collective action layer to being thin, its roots in a regulation influence the structure of its constitutional layer. Existing in the institutional context of Colorado's representative democracy, any regulation is subject to the legislative process, which by design provides for a well-defined, if unwieldy process for including all stakeholders in the process by which METRC changes. In principle, regulations may include finer-grained domain-specific constitutional processes, and these may be designed to include input from the range of stakeholders. This is not the case for METRC, where the user groups remain firmly in the space of “Arena A” to use terminology from participatory design: designing the artifacts and systems of work. Colorado’s Marijuana Enforcement Division granted Franwell considerable discretion to translate aspirational regulatory language into concrete interfaces, devices, and ontologies, and to evolve that implementation over time, within the constraints of the law. Consequently, METRC’s constitutional layer is not participatory by default, although our case reveals how, in times of crisis, METRC was able to gain input and build legitimacy by voluntarily bringing a wider range of stakeholders to inform a major design change. Furthermore, it institutionalized the user group model to continue to represent stakeholder voices in changes to its operations. To the extent that the system is and will remain stable, this may be sufficient. However, future shocks are certain as factors like technology and consumer preferences change. This assumption is the basis of our motivating thesis that digital institutions should design constitutional change processes and design a wide range of stakeholders into those processes.

METRC’s crisis at launch can be traced back to decisions to design the digital component of this institution with little input from its users. The top-down design approach focused on the superficial similarities between the products moving through cannabis, agricultural, and pharmaceutical supply chains, and overlooked how a within-firm inventory tracking system migrated into a profoundly different institutional context. In traditional supply chains, actors can enforce compliance with their rules and practices through legal contracts, technical standards, professional norms, and prices. These instruments typically perform the functions that the Ostrom Workshop identifies with successful governing: defining boundaries and rights, prescribing mechanisms for monitoring and sanctioning, and organizing information around the system. Cannabis products are clearly the resource that the institution orients itself around, and METRC creates information flows about this resource that must satisfy such a complex set of criteria—from inventory tracking to compliance monitoring and enforcement—as to be beyond the means of a traditional top-down product design process. METRC’s state-mandated inventory tracking system was the bureaucratic implementation of a legislative directive lacking many of these feedback instruments, compliance with which could be legally enforced. In the absence of substantive feedback mechanisms, the technical system and its bureaucratic managers could not anticipate demand for tags, sustain uptime, or accommodate new business models and product categories. That it was mandated by a democratically governed state did not redeem it, as a legislature representing millions of citizens necessarily operates at timescales much larger than the emerging market required for policy adjustments.

Although METRC is not the kind of online community envisioned by *BSOC* specifically, it is a digital institution of the sort designed by others within the behavioral engineering tradition, such as in work on nudges [112]; METRC is a government-operated database that collects the daily inventory and transactions of every licensed cannabis operator across Colorado. Yet seen through the lens of digital institution design, METRC was confronted with challenges even recognizable in CSCW work like *BSOC*, such as cold-starting a new system, socializing newcomers, and regulating behavior. In the case of starting a new system, as a state-sanctioned monopoly METRC did not face the same challenges of carving out a niche, defending the niche from competitors, or building a critical mass of participation. But this institution faced similar design choices around defining the scope of the institution, compatibility with other systems, and organizing information and interactions [71, p.232–233].

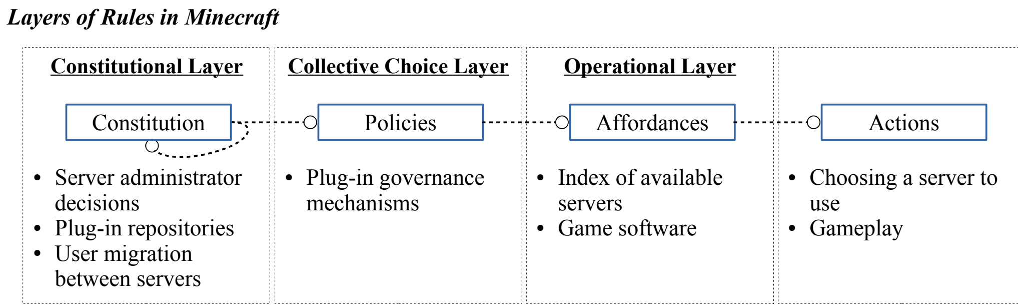

Fig. 3. Examples of layers of rules in servers of the multiplayer game Minecraft. Server administrators can choose from a suite of plugins that facilitate choices in server governance, such as rules of in-game markets or chat capabilities. The rich landscape of plugins provides constitutional layer flexibility that enables a thriving ecosystem of varied and successful communities.

Much of the guidance the behavioral engineering tradition offers comes from an assumption that participation in a digital institution is voluntary. Design claims like "People are more likely to comply with requests the more they like the requester" [71, p.32] or "Face-saving ways to correct norm violations increases compliance" [71, p.153] are deeply distorted in non-voluntary institutional contexts like METRC. Compliance with requests from the MED, through systems like METRC, is incentivized through escalating sanctions like coercion, confiscation, and imprisonment rather than charisma, motivation, or commitment. METRC is not an isolated case: institutions like Transportation Security Administration watch lists, financial and social credit scores, and mobile location traces share similar features of less-than-voluntary participation. How can digital institutions with less-than-voluntary membership be designed to be more accountable, successful, or sustainable?

The "choice architecture," from the nudge paradigm of psychological engineering for compliant behavior, would suggest that technocrats who fully apprehend the decision space could provide better defaults, reduce overload, and improved comparability for agents to make ideal decisions. But the decision space of newly legal cannabis markets could not be apprehended because there were no precedents, little information, and complex interactions; there had never been a legal recreational cannabis market, and there were no empirical priors from which to draw institutional design guidance. A complex regulatory apparatus governing a legally precarious market with a complex set of incentives and stakeholders nevertheless shows the importance of participatory mechanisms for feedback and institutional design. Here, our approach goes beyond what a human-centric or participatory design approach would prescribe by attending to power. METRC ended up adopting a strategy akin to a participatory design approach with user groups that enabled it to be more responsiveness to end-user needs, but these still lack formal regulatory recognition as negotiating or deliberative bodies. Our focus on constitutional layers is meant to provide a continuous opportunity for suggestions and modification by users.

Case Study 3: Amateur game server governance

For our final case study we examine governance in the popular game Minecraft. Governance in digital games has become an important arena for innovative institutional design [12, 23, 38, 69], especially multiplayer games that follow a self-hosting model, by which fans in the community personally host publicly accessible instances of a game. Fans who self-host assume the difficulties of governance, while also gaining access to technological innovations for addressing those difficulties.

We focus on one particularly popular self-hosted game, Minecraft. Minecraft is an open-world ("sandbox") construction game that is notable for a culture of amateur-run self-hosted multiplayer servers that players shop between in search of engaging community. Servers must compete with each other for users while struggling to manage several resource provisioning problems, including the needs to manage endemic vandalism by the game's relatively young player base, to provide sufficient RAM, CPU, and network bandwidth for players, and to solve in-game resource problems. Preventing vandalism may mean controlling the availability or destructive effects of in-game phenomena like fire, magma, and TNT that the majority of players might enjoy and use responsibly, but that an inevitable minority abuse. Providing sufficient physical resources is surprisingly difficult: each additional player to this resource-intensive, lag-sensitive game increases the recommended server requirements by 1GB of RAM, 1–3 Mbits/s of up- and down- bandwidth, and 5GB of cache. A failure to adequately provide any of these complex sociotechnical services can lead within minutes to users abandoning a server for competing hosts. Contrary to mythologies about the "cloud" abolishing marginal cost economics, Minecraft provides a case where finite and exhaustible digital resources imply the need for robust resource management strategies [38].

Fortunately, administrators can address these forbidding governance challenges by drawing upon several tools developed by Minecraft's large informal developer ecosystem. There is a large collection of plugins that automate many dimensions of governance, including peer monitoring, resource monitoring, rule enforcement, trade, vandalism, decision-making, information transmission, and communication, as well as plugins that define complete market and property rights institutions: systems of property rights, shops, social hierarchies, and group allegiances. The richness, granularity, and modularity of plugin governance has reduced many of the challenges of amateur governance to the flipping of switches, and has made possible an exciting diversity of governance styles. The library of governance plugins available to an administrator leaves it entirely to their discretion what governance capabilities their server will have. And administrators take full advantage of this flexibility. Among the diverse amateur-run digital institutions we have observed, and well beyond the professionalized for-profit servers which exist in their own ecosystem, are a groups of friends who install no rules but punish infractions by chiding each other in person, a pristine server that solves its key resource problem by preventing visitors from changing anything, and an explicitly anarcho-capitalist server with market-based law enforcement, in which users secure justice for themselves by placing rich bounties for the virtual heads of those that have wronged them. There are active, long-lived communities with as few as 2 users and as many as 2,000. Administrators implement their naïve folk theories of cooperation, incentive design, and resource management, and iterate. The lessons learned by these governance autodidacts are then reified into the plugins that are refined, repropagated, and refined again.

Constitutional crisis. What governance challenges do Minecraft server operators face, especially small-scale operators with fewer resources? How should operators design their digital institutions to manage these resource challenges? And, specifically, what is the role of servers' constitutional choice mechanisms in managing or exacerbating these problems? Minecraft servers fail when they stop being visited, something that can happen when their administrators fail to manage the many finite resources their server needs to function, or when they fail to provide users with a stable and secure environment safe from vandalism and harassment. A server's sensitivity to these threats increases with server size. But because they are all relatively small, defined by default by a single administrator, and exist entirely in software, individual servers can change themselves easily.

Of course, the existence of constitutional-level rules and widespread participation in those rules are different. Minecraft servers are not democratic by default. Administrators often make themselves available to user feedback either directly or through ticket or forum systems, but it is rare for these mechanisms to be substantively participatory, or formally defined into the processes of constitution-level change. The number of people entitled to participate at a server's constitutional level is rarely greater than one.

Zooming out from the individual server, to view the system of competing servers as a digital institution, another mechanism for broad-based involvement in constitutional change becomes clear, one very much like that available in cryptocurrencies. Because servers compete with each other for users, users can, again, vote with their feet, by departing servers that do not explicitly or incidentally represent their needs, and populating or even reproducing servers that do. Because of this structure, Minecraft servers fare well as a community. Collectively they seem able to solve the resource management problems that individual servers so often fail to overcome.

This success raises a question: why does the market approach to constitutional change seem to work in Minecraft but not in cryptocurrency? The answer may be that, compared to cryptocurrencies, Minecraft servers have much lower switching costs, largely because fewer people have to coordinate on a Minecraft server than on a currency for objects there to have value (fewer network externalities). With these ingredients in place, Minecraft servers that are individually autocratic seem as a collective of competitors to implement meaningful participatory change processes typical of the democratic institutions that the Ostrom Workshop describes.

Comparing theoretical guidance. The amateur Minecraft community builder seems especially suited to adopt the engineering approach to digital institution design. These relatively small-scale online communities have one self-appointed leader with complete control over the structure of their virtual world, a leader highly incentivized to overcome the obstacles to successful community. As such, the elements identified by the psychological engineering tradition map directly to many of the problems that Minecraft server administrators face. Amateur Minecraft administrators struggle to attract, socialize, and retain new members, and motivate rule compliance. The engineering approach offers solutions to these problems that administrators can benefit from directly, because they have the power to redesign the system in a way that implements "best practices" extracted from previous social engineers who have faced the same problems.

However, the behavioral manipulations of the engineering paradigm miss important dimensions of governing a self-hosted game server, and important mechanisms for ensuring its success. For example, *BSOC* cases tend to consider communities in isolation, largely ignoring the influence of competing communities that pursue the same mission in a possibly superior, inferior, or just different way. By failing to give servers' sociological and ecological contexts a central role in the framework, they miss factors such as the strategic dimension that competition adds to an administrator's reasoning, and the strong incentives administrators have to make decisions in the community's interest. The Ostrom Workshop, by contrast, is as much influenced by ecology as economics, and has several multi-scale frameworks — including the Institutional Analysis and Design (IAD) [87], the Socio-Ecological Systems framework (SES) [88], and the action situation [89] — for analyzing institutions with respect to their socio-ecological setting.

Applying the resource management perspective begins with identifying the limited resources around which collective action, and associated governance, structures itself. The threats to a game server include physical computational resources, which are common-pool and must be conserved, in-game resources, and vandalism (which must be minimized), and social capital, in the form of valuable contributors who must be retained. Governance structure around the resource management community can then be partitioned into three levels within the levels of choice outlined by the Ostrom Workshop (Fig. 3). At the operational level are the actions allowed by administrators, the software, and code in the greater ecosystem of the game. Administrators will take an active interest in what specific actions players can perform, shaping those possible operations by configuring the server to only allow certain actions. A restricted list of possible game actions can tacitly forbid or encourage certain types of behavior. There are much fewer constraints at the collective action layer, which is entirely the prerogative of the administrator, who is under no formal constraints on how to govern their server, only the practical constraints due to fear of flight by users. Administrators with a positive vision or goal for their community can use these capabilities to steer server and structure interactions between players to encourage aggregate behavior, or collective action, toward that goal. At the constitutional layer are the same constraints determining the collective action layer: the administrator is sovereign and has complete discretion over how and when the community's governance changes. The powers that players have over this process are the powers of voice and exit: to either sway a community's administrator or leave it for preferred one.

With its view of institutions largely framed around resource governance, the Ostrom Workshop's design principles are a valuable guide for the type of governance challenges facing a self-hosted game server. Applied to this case, our analysis highlights those principles that recognize the value of conflict as a symptom of the need for change, those that encourage institutional fit to the environment and, most importantly, those that recommend a role for users in how a digital institution changes.

DISCUSSION

In our case studies we applied the Ostrom Workshop's levels of choice framework to inspect three contemporary digital institutions, the challenges they face, and the role of formal participatory change processes (or lack thereof) in each institution's problems. The highly varied nature of these three cases highlights the broad applicability of institutional analysis and design tools to digital institutions. There were several commonalities in design recommendations through our institutional lens. In each of the cases, we highlight the need for more attention to scaffolding ongoing participatory design through constitutional mechanisms in digital institutions.

In the first case, we viewed each cryptocurrency coin itself as a digital institution, with the rules of each institution represented as not just each coin's codebase but also any other declared governance processes, laws pertaining to cryptocurrencies, and extralegal informal norms of social engagement around the coins. Our analysis of the cryptocurrency ecosystem belies its ideological commitment to purely technical (rather than sociotechnical) governance, and shows that the community's attempt to remove humans from the operation of human institutions has produced currencies that are inflexible. In many notable cases in this ecosystem, each currency's inevitable need for incremental policy change can be met only with blunt policy tools whose side effects undermine the very communities they are intended to serve.

In the second caset, we viewed the cannabis market in Colorado as a digital institution because of Franwell's METRC system that mediates all production, distribution, and sales activities. As with cryptocurrencies, there are also other institutional structures at play in this case besides the code of METRC, such as the state laws that regulate recreational cannabis. Our analysis of METRC showed how top-down management, incongruent ontologies, and ponderous change processes conspired to produce a digital supply chain surveillance tool that caused a crisis in a precarious industry. This case is particularly instructive for designers of digital institutions who resist theoretical or historical guidance because of their self-imagined exceptional conditions. Despite intense political and legal pressure, overwhelming popular attention, and clear market incentives, the roll-out of a legal recreational cannabis market was jeopardized by top-down digital institutional design thinking. The crisis abated only when stakeholders worked outside of established channels to amend the digital regulatory infrastructure to meet end-user demands. They used participatory design practices and have since institutionalized "User Groups" as a primary mechanism for evolving both the technology and regulations. Much as network effects lock users into social platforms, this case highlighted how monopolistic institutions insulated from "vote with your feet" mechanisms still must incorporate mechanisms for end-user feedback to ensure the institution remains responsive and resilient to changes in the broader environment.

In the third case, on multiplayer activity in the sandbox video game Minecraft, a population of volunteer amateur server administrators has created hundreds of thousands of experiments in small-scale governance. In order for a server to succeed, its administrator must learn how to foster successful collective action "on the job", by formulating policy and installing code modules that implement different dimensions of governance. Where administrators fail to meet the demands of users or their environment, competition between administrators for traffic leverages market incentives to implement ecosystem-wide constitution-level change. Why "vote with your feet" works in Minecraft, where it seems to fail in the cryptocurrency case, seems due to several factors, including the immutable nature of individual currencies, the higher costs in them of exiting or performing most other actions, and a fundamental difference between the two types of platforms: convergence on a single standard/community is optional for effective game servers and essential for effective currencies.

The Ostrom Workshop's resource management frameworks emphatically do not make prescriptive ("if X then Y") or normative ("A is better than B") recommendations about how to *do* institutional design. Rather, they provide a vocabulary that abstracts the immediate components of a system into more generalized components that can be used for comparison across seemingly disparate classes of digital institutions, and emulation of successes. This is not to frame institutional design thinking as mere theoretical abstraction: political scientists, economists, and ecologists have spent decades applying this framework to policy analysis and accumulating empirical evidence that the most effective management of common-pool resources is often through democratically designed and managed institutions [84, 86]. This evidence repeatedly demonstrates that institutional structures ignoring local contexts or lacking mechanisms for local stakeholder-driven change are vulnerable to failure and collapse. The application of the resource management and institutional analysis perspectives to these cases enables an engineer, scientist, or manager engaged in digital institution design to abstract away from the proximate pressures of "building an alternative financial system", "designing a surveillance system for regulators", or "running a sandbox game", and attend to more general challenges of sustainable, participatory resource management.

Robust digital institutions require rules about rule-making that substantively involve diverse stakeholders. Being intentional about the constitutional level enables digital institutions to anticipate disruptions and accommodate innovations, which can make them more adaptable and humane. We have emphasized two arguments for formal participatory change processes in digital institutions. First, participatory design and institutional resilience are fundamentally coupled: institutions (digital or otherwise) that forsake substantive participatory mechanisms run the risk of stagnation, drift, and capture. Top-down decision-making may have superficial benefits of fidelity and scalability when implementing new digital institutional designs, but it can also lead to information bottlenecks, founder syndrome, challenges to legitimacy, and problems adapting to a changing environment. The three case studies illustrate different strains that digital institutions undergo when they lack robust constitutional layers. Second, privileging participatory processes in digital institutional design generates emergent outcomes: the design process should not be an exercise in choosing among finite options or legitimizing a pre-ordained outcome. A common frustration when working with institutional design frameworks like the Ostrom Workshop is their lack of clear prescriptive guidance about how to design the "best" institution: "if X problem then institutional design Y." One of Ostrom's major conclusions was that there is no "best" design for a resource management problem: that it is possible for many kinds of structures to sensitively adapt to local conditions, and, further, that the most enduring institutions exhibit "institutional diversity": they pragmatically combine superficially incompatible governance strategies under one institutional umbrella. The goal of the Ostrom Workshop's frameworks is to provide a vocabulary for abstracting from the specifics of a single case to comparing abstracted cases and identifying general principles of success and failure.

CONCLUSION

Despite the increasing importance of digital institutions, CSCW's strong traditional interests in participatory design and online communities have generated little synthesis. Our aim in this work is to appreciate the complexity of digital institution design, while still aiming to ultimately be engaging in it. The participatory tradition in digital institution design, kicked off in large part by Kollock and Smith [68], has lagged behind other, more convenient paradigms. Kraut and Resnick's *Building Successful Online Communities* represents one of the most influential and theoretically-grounded paradigms within the space of digital institution design. While its psychological engineering approach offers a useful slate of design tricks, such approaches have prioritized top-down thinking over designing institutions for substantive participation. As convenient as it is for platforms serving hundreds of millions of users, behavioral engineering design thinking fails to incorporate values fundamental to fair, resilient, and ethical digital institutions. The limitations of the engineering approach have become impossible to ignore since at least 2016: an emphasis on monetized engagement over end-user safety gave tools to outside agents who hijacked the behavioral change and engagement maximization capabilities of online platforms toward malicious ends. The result is online platforms that have lost control over their coupling to broader geo-political forces, and are failing to give users governance responsibilities, or fulfill their rights.

The fields of human computer interaction, social computing, and computer-supported cooperative work find their history in the design of interfaces to support individual user experiences, supporting group work, and understanding and designing for technology use in context. Dourish [29] and Grudin [50] have outlined how ethnographic methods came to be adopted within human-computer interaction in response to the need for understanding the social organization and contexts of technology use in "real world" settings beyond controlled laboratory environments. As these research communities have moved their attention to larger-scale settings, they have often brought the same limited image of the user's role [7]. Existing design processes have excluded individual users from substantive roles in system change, preferring instead to embed them into lower-level roles, and interfaces engineered to nudge behavior to suit design goals.

Frameworks for analyzing the design of institutions offer an additional step "upwards" and invite comparisons of these organizations and contexts. Where an experimental or evaluation perspective might ask which rules users perceive to be fair, or an archival or ethnographic perspective might analyze the unfolding interactions around implementing rules, an institutional design perspective asks what kinds of norms and rules (operational level), what implementation strategies (collective level), and what rule-making processes (constitutional level) are used, how these interact, and what alternatives exist. Dourish [30] argued for broadening HCI from the interactions between user and screen and orienting to considerations of scale for design engagements. The Ostrom Workshop's frameworks go a long way towards participatory design's concern for the "politics of design" as well as answering Dourish's call for a "design of politics" within HCI [30] by considering the design of institutions that are focused on governing a system's limited resources through participatory and democratic mechanisms. At a time when issues such as the disappearance of application programming interfaces (APIs) signal concentration of control in the hands of platform owners (cf. [58]), consideration of participation is of the utmost important.

The designers of digital institutions should look to resource management approaches as they confront the limits of their crumbling governance strategies. The Ostrom Workshop in particular has several successful frameworks for theorizing and classifying the technological, cultural, economic, and political strategies employed by institutions across a variety of contexts and scales. These frameworks make more facets of digital institution design legible to CSCW scholars. The engineers, scientists, and managers on major social platforms are engaged in digital institution design by making choices about roles, resources, rights, rules, and governance. However, research and teaching emphasizing resource management and institutional analysis are conspicuously absent from the computer and information science and engineering fields from which platforms' governing engineers, designers, and managers are recruited. With more general frameworks that permit more ambitious comparisons of more diverse institutions, we get closer to a science of digital institution design that is general enough to avoid the problem of asssuming away the future's most creative, innovative, and inspiring experiments in digital democracy.

ACKNOWLEDGMENTS

Special thanks to Matti Nelimarkka, Caitie Lustig, Andrew Hoffman, Steve Slota, Mike Katell, and Meg Young for feedback, and to Quinn DuPont for guidance on the cryptocurrency case. This work was supported in part by the Washington Research Foundation, by a Data Science Environments project award from the Gordon and Betty Moore Foundation (Award #2013-10-29) and the Alfred P. Sloan Foundation (Award #3835) to the University of Washington eScience Institute, and by the Neukom Institute for Computational Science at Dartmouth College. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsors.

REFERENCES

[1] Sadia Afroz, Vaibhav Garg, Damon McCoy, and Rachel Greenstadt. Honor among thieves: A common's analysis of cybercrime economies. In *2013 APWG eCrime Researchers Summit*, pages 1–11. IEEE, 2013.

[2] Howard Aldrich. *Organizations Evolving*. Sage, 1999.

[3] Peter M. Asaro. Transforming society by transforming technology: The science and politics of participatory design. *Accounting, Management and Information Technologies*, 10(4):257–290, 2000.

[4] WR Ashby. Requisite variety and its implications for the control of complex systems. *Cybernetica*, 1(2):1–17, 1958.

[5] Marcella Atzori. Blockchain technology and decentralized governance: Is the state still necessary? *Journal of Governance and Regulation*, 6(1):45–62, 2017.

[6] Liam Bannon, Kjeld Schmidt, and Ina Wagner. Lest we forget. In Susanne Bødker, Niels Olof Bouvin, Volker Wulf, Luigina Ciolfi, and Wayne Lutters, editors, *Proceedings of the 12th European Conference on Computer Supported Cooperative Work*, ECSCW '11, pages 213–232, London, 2011. Springer London.

[7] Eric P. S. Baumer and Jed R. Brubaker. Post-userism. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pages 6291–6303, New York, NY, USA, 2017. ACM.

[8] Gerard Beenen, Kimberly Ling, Xiaoqing Wang, Klarissa Chang, Dan Frankowski, Paul Resnick, and Robert E. Kraut. Using social psychology to motivate contributions to online communities. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work*, CSCW '04, pages 212–221, New York, NY, USA, 2004. ACM.

[9] Yochai Benkler. Coase's Penguin, or, Linux and "The Nature of the Firm". *The Yale Law Journal*, 112(3):369, December 2002.

[10] Ivan Beschastnikh, Travis Kriplean, and David W. McDonald. Wikipedian Self-Governance in Action: Motivating the Policy Lens. In *Proceedings of the International Conference on Web and Social Media*, ICWSM '08. AAAI, 2008.

[11] Erling Björgvinsson, Pelle Ehn, and Per-Anders Hillgren. Participatory design and "democratizing innovation". In *Proceedings of the 11th Biennial Participatory Design Conference*, PDC '10, pages 41–50, New York, NY, USA, 2010. ACM.

[12] Tom Boellstorff. *Coming of Age in Second Life: An Anthropologist Explores the Virtually Human*. Princeton University Press, 2015.

[13] Rainer Böhme, Nicolas Christin, Benjamin Edelman, and Tyler Moore. Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives*, 29(2):213–38, 2015.

[14] danah boyd and Kate Crawford. Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, Communication & Society*, 15(5):662–679, 2012.
[15] Susan L. Bryant, Andrea Forte, and Amy Bruckman. Becoming Wikipedian: Transformation of Participation in a Collaborative Online Encyclopedia. In *Proceedings of the 2005 International ACM SIGGROUP Conference on Supporting Group Work*, GROUP '05, pages 1–10, New York, NY, USA, 2005. ACM.
[16] Brian Butler, Elisabeth Joyce, and Jacqueline Pike. Don't look now, but we've created a bureaucracy: The nature and roles of policies and rules in Wikipedia. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pages 1101–1110, New York, NY, USA, 2008. ACM.
[17] John T. Carnevale, Raanan Kagan, Patrick J. Murphy, and Josh Esrick. A practical framework for regulating for-profit recreational marijuana in US states: Lessons from Colorado and Washington. *International Journal of Drug Policy*, 42:71–85, 2017.
[18] Carmelo Cennamo and Juan Santalo. Platform competition: Strategic trade-offs in platform markets. *Strategic Management Journal*, 34(11):1331–1350, 2013.
[19] Erwin Chemerinsky. Introduction: Marijuana Laws and Federalism. *Boston College Law Review*, 58(3):857–862, 2017.
[20] Andrew Clement and Peter Van den Besselaar. A retrospective look at PD projects. *Commun. ACM*, 36(6):29–37, June 1993.
[21] Commoneffort. EOS constitution. `https://github.com/EOSIO/eos/blob/5068823fbc8a8f7d29733309c0496438c339f7dc/constitution.md`, 2018.
[22] Michael Cox, Gwen Arnold, and Sergio Villamayor-Tomás. A review of design principles for community-based natural resource management. *Ecology and Society*, 15(4), 2010.
[23] Nicole Crenshaw, Jaclyn LaMorte, and Bonnie Nardi. "Something we loved that was taken away": Community and neoliberalism in World of Warcraft. In *Proceedings of the 50th Hawaii International Conference on System Sciences*, 2017.
[24] Christopher A Le Dantec and Carl DiSalvo. Infrastructuring and the formation of publics in participatory design. *Social Studies of Science*, 43(2):241–264, 2013.
[25] Laura DeNardis. *Protocol Politics: The Globalization of Internet Governance.* MIT Press, 2009.
[26] Ada Diaconescu and Jeremy Pitt. Holonic Institutions for Multi-scale Polycentric Self-governance. *International Workshop on Coordination, Organizations, Institutions, and Norms in Agent Systems*, pages 19–35, 2014.
[27] Gregory M Dickinson. An interpretive framework for narrower immunity under Section 230 of the Communications Decency Act. *Harvard Journal of Law & Public Policy*, 33:863, 2010.
[28] Paul DiMaggio, Eszter Hargittai, W Russell Neuman, and John P Robinson. Social implications of the Internet. *Annual Review of Sociology*, 27(1):307–336, 2001.
[29] Paul Dourish. Implications for design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, pages 541–550, New York, NY, USA, 2006. ACM.
[30] Paul Dourish. HCI and environmental sustainability: The politics of design and the design of politics. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems*, DIS '10, pages 1–10, New York, NY, USA, 2010. ACM.
[31] Quinn DuPont. Experiments in algorithmic governance: A history and ethnography of "The DAO," a failed decentralized autonomous organization. In *Bitcoin and Beyond*, pages 157–177. Routledge, 2017.
[32] Quinn DuPont. *Cryptocurrencies and Blockchains.* Polity, 2019.
[33] Pelle Ehn. Participation in design things. In *Proceedings of the Tenth Anniversary Conference on Participatory Design 2008*, PDC '08, pages 92–101, Indianapolis, IN, USA, 2008. Indiana University.
[34] Rosta Farzan, Laura A. Dabbish, Robert E. Kraut, and Tom Postmes. Increasing commitment to online communities by designing for social presence. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work*, CSCW '11, pages 321–330, New York, NY, USA, 2011. ACM.
[35] Andrea Forte and Amy Bruckman. Scaling consensus: Increasing decentralization in Wikipedia governance. In *Hawaii International Conference on System Sciences*, 2008.
[36] Andrea Forte, Vanessa Larco, and Amy Bruckman. Decentralization in Wikipedia governance. *Journal of Management Information Systems*, 26(1):49–72, 2009.
[37] Jo Freeman. The tyranny of structurelessness. *Berkeley Journal of Sociology*, pages 151–164, 1972.
[38] Seth Frey and Robert W Sumner. Emergence of integrated institutions in a large population of self-governing communities. *PloS one*, 14(7):e0216335, 2019.
[39] Batya Friedman, Peter H Kahn, and Alan Borning. Value sensitive design and information systems. *Human-computer interaction in management information systems: Foundations*, pages 348–372, 2006.
[40] Batya Friedman and Helen Nissenbaum. Bias in computer systems. *ACM Transactions on Information Systems (TOIS)*, 14(3):330–347, 1996.
[41] Alexander R Galloway. *Protocol: How Control Exists after Decentralization.* MIT press, 2004.
[42] Johannes Gärtner. Participatory design in consulting. *Computer Supported Cooperative Work*, 7(3):273–289, Sep 1998.

[43] Johannes Gärtner and Ina Wagner. Mapping actors and agendas: Political frameworks of systems design and participation. *Human-Computer Interaction*, 11(3):187–214, 1996.

[44] Abigail Geiger. About six-in-ten Americans support marijuana legalization. Technical report, Pew Research Center, 2018.

[45] Tarleton Gillespie. The relevance of algorithms. *Media Technologies: Essays on Communication, Materiality, and Society*, 167, 2014.

[46] Tarleton Gillespie. *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions that Shape Social Media*. Yale University Press, 2018.

[47] Eric Gorski. Colorado's vision of regulating medical pot dispensaries with high-tech checks and seed-to-sale tracking becomes a pipe dream. *The Denver Post*, page 1A, 2013.

[48] Eric Gorski. Regulators vow better oversight. *The Denver Post*, page 9A, 2014.

[49] Joan Greenbaum. Back to labor: Returning to labor process discussions in the study of work. In *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work*, CSCW '96, pages 229–237, New York, NY, USA, 1996. ACM.

[50] Jonathan Grudin. The computer reaches out: The historical continuity of interface design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '90, pages 261–268, New York, NY, USA, 1990. ACM.

[51] Jonathan Grudin. Computer-supported cooperative work: History and focus. *Computer*, 27(5):19–26, 1994.

[52] Eszter Hargittai. Second-level digital divide: Differences in people's online skills. *First Monday*, 7(4), 2002.

[53] Colin Harris. Institutional solutions to free-riding in peer-to-peer networks: a case study of online pirate communities. *Journal of Institutional Economics*, pages 1–24, 2018.

[54] Bradi Heaberlin and Simon DeDeo. The Evolution of Wikipedia's Norm Network. *Future Internet*, 8(2):14, April 2016.

[55] Charlotte Hess and Elinor Ostrom. A framework for analysing the microbiological commons. In *International Social Science Journal*, pages 335–349, 2006.

[56] Charlotte Hess and Elinor Ostrom. *Understanding Knowledge as a Commons: From Theory to Practice*. MIT Press, Cambridge, MA, 2007.

[57] Albert O Hirschman. *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*, volume 25. Harvard University Press, 1970.

[58] Bernie Hogan. Social media giveth, social media taketh away: Facebook, friendships, and APIs. *International Journal of Communication*, 2018.

[59] Kimberly A. Houser and Robert E. Rosacker. High times: A history of marijuana laws in the United States. *International Journal of Business & Public Administration*, 11(2):131–141, 2014.

[60] Matthew Iles. The civil constitution (Beta). https://blog.joincivil.com/the-civil-constitution-beta-64460a181e08, 2018.

[61] John Ingold. Pot tag tracking on hold. *The Denver Post*, page 5A, 2013.

[62] Mike Isaac. Facebook's Mark Zuckerberg says he'll shift focus to users' privacy. *The New York Times*, March 2019.

[63] Eaman Jahani, P. M. Krafft, Yoshihiko Suhara, Esteban Moro, and Alex Sandy Pentland. ScamCoins, S*** Posters, and the search for the Next Bitcoin™: Collective sensemaking in cryptocurrency discussions. *The 21st ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, 2018.

[64] Brian Keegan. Research agenda for cannabis informatics. In *2018 Institute of Cannabis Research Conference*. Colorado State University-Pueblo. Library, 2018.

[65] Brian Keegan and Casey Fiesler. The evolution and consequences of peer producing Wikipedia's rules. In *Eleventh International AAAI Conference on Web and Social Media*, May 2017.

[66] Brian C Keegan, Patricia Cavazos-Rehg, Anh Ngoc Nguyen, Saiph Savage, Jofish Kaye, Munmun De Choudhury, and Michael J Paul. CHI-nnabis: Implications of marijuana legalization for and from human-computer interaction. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pages 1312–1317. ACM, 2017.

[67] Finn Kensing and Jeanette Blomberg. Participatory design: Issues and concerns. *Computer Supported Cooperative Work (CSCW)*, 7(3):167–185, Sep 1998.

[68] Peter Kollock and Marc Smith. Managing the virtual commons. *Computer-mediated communication: Linguistic, social, and cross-cultural perspectives*, pages 109–128, 1996.

[69] Yubo Kou and Bonnie A Nardi. Governance in League of Legends: A hybrid system. In *Foundations of Digital Games*, 2014.

[70] P. M. Krafft, Nicolás Della Penna, and Alex Sandy Pentland. An experimental study of cryptocurrency market dynamics. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, page 605. ACM, 2018.

[71] Robert E. Kraut, Paul Resnick, Sara Kiesler, Yuqing Ren, Yan Chen, Moira Burke, Niki Kittur, John Riedl, and Joseph Konstan. *Building Successful Online Communities: Evidence-Based Social Design*. The MIT Press, 2012.

[72] Travis Kriplean, Ivan Beschastnikh, David W. McDonald, and Scott A. Golder. Community, consensus, coercion, control: CS*W or how policy mediates mass participation. In *Proceedings of the 2007 International ACM Conference on*

Supporting Group Work, GROUP '07, pages 167–176, New York, NY, USA, 2007. ACM.
[73] Lawrence Lessig. *Code and Other Laws of Cyberspace.* Basic Books, Inc., 1999.
[74] Caitlin Lustig and Bonnie Nardi. Algorithmic authority: The case of Bitcoin. In *2015 48th Hawaii International Conference on System Sciences*, pages 743–752. IEEE, 2015.
[75] Caitlin Lustig, Katie Pine, Bonnie Nardi, Lilly Irani, Min Kyung Lee, Dawn Nafus, and Christian Sandvig. Algorithmic authority: The ethics, politics, and economics of algorithms that interpret, decide, and manage. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pages 1057–1062. ACM, 2016.
[76] William J Luther. Cryptocurrencies, network effects, and switching costs. *Contemporary Economic Policy*, 34(3):553–571, 2016.
[77] Metrc. The System. https://www.metrc.com/the-system, 2016.
[78] Jonathan T. Morgan and Mark Zachry. Negotiating with angry mastodons: The Wikipedia policy environment as genre ecology. In *Proceedings of the 16th ACM International Conference on Supporting Group Work*, GROUP '10, pages 165–168, New York, NY, USA, 2010. ACM.
[79] Michael J Muller. Participatory design: The third space in HCI. *Human-computer interaction: Development process*, 4235:165–185, 2003.
[80] Francesca Musiani. Alternative technologies as alternative institutions: The case of the domain name system. In *The Turn to Infrastructure in Internet Governance*, pages 73–86. Springer, 2016.
[81] Malavika Nair and Nicolás Cachanosky. Bitcoin and entrepreneurship: Breaking the network effect. *The Review of Austrian Economics*, 30(3):263–275, 2017.
[82] Arvind Narayanan, Joseph Bonneau, Edward Felten, Andrew Miller, and Steven Goldfeder. *Bitcoin and cryptocurrency technologies: A comprehensive introduction.* Princeton University Press, 2016.
[83] Douglass C. North. Institutions. *The Journal of Economic Perspectives*, 5(1):97–112, 1991.
[84] Elinor Ostrom. *Governing the Commons.* Cambridge University Press, 1990.
[85] Elinor Ostrom. Designing complexity to govern complexity. In *Property Rights and the Environment: Social and Ecological Issues*, pages 30–43. The International Bank for Reconstruction and Development/The World Bank, 1995.
[86] Elinor Ostrom. Incentives, rules of the game, and development. In Michael Bruno and Boris Pleskovic, editors, *World Bank Conference on Development Economics*, pages 207–234. The World Bank, Washington DC, 1995.
[87] Elinor Ostrom. *Understanding Institutional Diversity.* Princeton University Press, 2005.
[88] Elinor Ostrom. A general framework for analyzing sustainability of social-ecological systems. *Science*, 325(5939):419–422, 2009.
[89] Elinor Ostrom. Beyond markets and states: Polycentric governance of complex economic systems. *American Economic Review*, 100(3):641–672, 2010.
[90] Manoj Parameswaran and Andrew B Whinston. Social computing: An overview. *Communications of the Association for Information Systems*, 19(1):37, 2007.
[91] Frank Pasquale. *The Black Box Society.* Harvard University Press, 2015.
[92] Jeremy Pitt, Dídac Busquets, and Sam Macbeth. Distributive Justice for Self-Organised Common-Pool Resource Management. *ACM Transactions on Autonomous and Adaptive Systems*, 9(3):1–39, October 2014.
[93] Jeremy Pitt and Ada Diaconescu. Structure and governance of communities for the digital society. In *2015 IEEE International Conference on Autonomic Computing*, pages 279–284. IEEE, 2015.
[94] Jeremy Pitt, Ada Diaconescu, and David Bollier. Technology for Collective Action. *IEEE Technology and Society Magazine*, 33(3):32–34, 2014.
[95] Joseph M. Reagle, Jr. Do As I Do: Authorial leadership in Wikipedia. In *Proceedings of the 2007 International Symposium on Wikis*, WikiSym '07, pages 143–156, New York, NY, USA, 2007. ACM.
[96] Yuqing Ren, F. Maxwell Harper, Sara Drenner, Loren Terveen, Sara Kiesler, John Riedl, and Robert E. Kraut. Building member attachment in online communities: Applying theories of group identity and interpersonal bonds. *MIS Quarterly*, 36(3):841–864, September 2012.
[97] Yuqing Ren, Robert Kraut, and Sara Kiesler. Applying common identity and bond theory to design of online communities. *Organization Studies*, 28(3):377–408, 2007.
[98] Paul Resnick and Robert E. Kraut. Introduction. In Robert E. Kraut and Paul Resnick, editors, *Building Successful Online Communities: Evidence-Based Social Design*, chapter 1, pages 1–19. MIT Press, Cambridge, MA, 2012.
[99] Travis L Ross and Lauren B Collister. A social scientific framework for social systems in online video games: Building a better looking for raid loot system in World of Warcraft. *Computers in Human Behavior*, 36:1–12, 2014.
[100] Landon Schnabel and Eric Sevell. Should Mary and Jane be legal? Americans' attitudes toward marijuana and same-sex marriage legalization, 1988–2014. *Public Opinion Quarterly*, 81(1):157–172, 2017.
[101] Charles M Schweik and Robert C English. *Internet Success: A Study of Open-Source Software Commons.* A Study of Open-source Software Commons. MIT Press, 2012.

[102] W. Richard Scott. Crafting an analytic framework: Three pillars of institutions. In *Institutions and Organizations*. SAGE Publications, Inc, fourth edition edition, 2013.
[103] W Richard Scott. *Organizations and Organizing: Rational, Natural and Open Systems Perspectives*. Routledge, 2015.
[104] Aaron Shaw and Benjamin M Hill. Laboratories of oligarchy? how the iron law extends to peer production. *Journal of Communication*, 64(2):215–238, 2014.
[105] M. Six Silberman. Reading Elinor Ostrom in Silicon Valley: Exploring institutional diversity on the Internet. In *Proceedings of the 19th International Conference on Supporting Group Work*, GROUP '16, pages 363–368, New York, NY, USA, 2016. ACM.
[106] Susan Leigh Star and Karen Ruhleder. Steps toward an ecology of infrastructure: Design and access for large information spaces. *Information Systems Research*, 7(1):111–134, 1996.
[107] Pontus Strimling and Seth Frey. Emergent cultural differences in online communities' norms of fairness. *Games and Culture*, 2018.
[108] Todd Subritzky, Simone Pettigrew, and Simon Lenton. Issues in the implementation and evolution of the commercial recreational cannabis market in Colorado. *International Journal of Drug Policy*, 27:1–12, 2016.
[109] Cass R. Sunstein and Richard H. Thaler. Libertarian paternalism is not an oxymoron. *The University of Chicago Law Review*, pages 1159–1202, 2003.
[110] S. Y. Tang. *Institutions and collective action: self-governance in irrigation*. ICS Press, 1992.
[111] Richard H. Thaler and Cass R. Sunstein. Libertarian paternalism. *American Economic Review*, 93(2):175–179, 2003.
[112] Richard H. Thaler and Cass R. Sunstein. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Penguin Books, New York, February 2009.
[113] Peter Vessenes. Point / counterpoint: Ethereum miners should blacklist TheDAO theft. https://vessenes.com/point-counterpoint-ethereum-miners-should-blacklist-thedao-theft/, 2016.
[114] Gili Vidan and Vili Lehdonvirta. Mine the gap: Bitcoin and the maintenance of trustlessness. *New Media & Society*, 21(1):42–59, 2019.
[115] Fernanda B Viegas, Martin Wattenberg, and Matthew M McKeon. The Hidden Order of Wikipedia. In D Schuler, editor, *Online Communities and Social Computing*, pages 445–454. Springer Verlag, Berlin, 2007.
[116] Fei-Yue Wang, Kathleen M Carley, Daniel Zeng, and Wenji Mao. Social computing: From social informatics to social intelligence. *IEEE Intelligent Systems*, 22(2):79–83, 2007.
[117] Kevin Werbach. *The Blockchain and the New Architecture of Trust*. MIT Press, 2018.
[118] Jeffrey Wilcke. To fork or not to fork. https://blog.ethereum.org/2016/07/15/to-fork-or-not-to-fork/, 2016.
[119] Langdon Winner. Do artifacts have politics? *Daedalus*, pages 121–136, 1980.
[120] Michael J Withey and William H Cooper. Predicting exit, voice, loyalty, and neglect. *Administrative Science Quarterly*, pages 521–539, 1989.
[121] Dave Yates and Jessica Speer. Over and under-regulation in the Colorado cannabis industry – A data-analytic perspective. *International Journal of Drug Policy*, 59:63–66, 2018.

Received April 2019; revised June 2019; accepted August 2019