



Decentralizing Power Through Blockchains:

Key Research Questions Across Disciplines

Directors

Andrea Goldsmith

Faculty Co-Director

Jaswinder Pal Singh

Faculty Co-Director

Matt Weinberg

Faculty Associate Director

Mike Maizels

Executive Director

Steering Committee

Maria Apostolaki

Andrew Chignell

Sanjeev R. Kulkarni

Nolan McCarty

Prateek Mittal

Andrés Monroy-Hernández

Jacob N. Shapiro

Pramod Viswanath

David Wentzlaff

Wei Xiong

Executive Summary

Princeton University has established the Center for the Decentralization of Power Through Blockchain Technology (the “DeCenter”), with a research, educational, and programming mandate that spans **three key pillars: technology, applications, and implications for society**. Based on discussions among a cross-disciplinary group of DeCenter faculty and attendees of our events, we present here a set of critical questions across the three pillars for researchers and practitioners in the field to address.

The work of the DeCenter is grounded in two methodological commitments. First, we take a *first-principles approach* to *long-term questions* surrounding blockchains and the decentralization of power. While the technologies and applications are evolving quickly, our primary focus is on fundamental approaches, solutions, and principles that will remain durable through this evolution.

Second, we assume a primarily *interdisciplinary* lens. Resolving the questions posed in this white paper will certainly leverage disciplinary expertise, but we believe some of the most important challenges require interdisciplinary collaboration across computer science, engineering, economics, political science, ethics, and the humanities. Many relevant topics have also long been studied in other areas of technology, social sciences, and humanities: It is important that we leverage that knowledge.

Our goal is to help set proper intellectual foundations, goals, and frameworks to enable effective cross-disciplinary progress in these areas. **We hope these frameworks will support the advancement of knowledge and the development of meaningful applications that leverage the benefits of decentralized trust and help decentralize power in the service of humanity.**

We believe foundational progress in the area requires a deeply interdisciplinary approach—across technology, applications, and societal implications—to long-term questions surrounding blockchains and the decentralization of power.

Background

Many societal activities and mechanisms depend on trust, including money, assets, identity, communication, commerce, governance, safety, and data management. Such trust has historically been provided by centralized intermediaries; for example, strangers are able to work together and enter into contracts because of trust enabled by governments and corporations. By becoming essential to human cooperation and transaction, these centralized entities have accumulated tremendous power. The power is often abused, either in overt ways—such as through government repression or corporate censorship and manipulation—or in subtle ways, such as through reaping extraordinary profits by monetizing and manipulating user behavior.

Technologies that enable large-scale, trustful cooperation without centralized intermediaries are gaining increased interest. Permissionless blockchains, pioneered by the invention of Bitcoin in 2008, are prominent examples. These permissionless blockchains aim to decentralize trust: Instead of a centralized entity, trust among unfamiliar entities is provided by a decentralized protocol that is simultaneously technical, economic and social. While end users are not expected to care about decentralization in itself, the decentralization of trust is essential for the provision of key properties that may be important to users. These include permissionless participation, censorship resistance, transparency, immutability, self-custody (holding assets or data oneself rather than turning them over to a centralized exchange or custodian to hold), and difficulty of change or inflation.

Through these properties, the decentralization of trust has transformative potential in many application domains, ranging from finance and business to culture, human rights, public goods, and society. **These transformations could lead to the decentralization of power and hence the reimagining of societal structures.** They could also broaden access, circumvent repression and censorship, remove frictions, eliminate unnecessary overheads imposed by intermediaries and siloed systems, and distribute power in more democratic, transparent and resilient ways. Centralized entities are not always eliminated when power is decentralized; however, their role is, at least, reduced to places where they provide important benefits that cannot easily be obtained otherwise.

Success stories of these transformations have already emerged. Decentralized “money” (and hybrid money like stablecoins) allows families to avoid financial ruin due to the debasement of local fiat currencies by their governments, disadvantaged groups to increase their economic freedom within their society (e.g., women to be paid directly rather

than through men in Afghanistan), and individuals, activists, and human rights organizations to send and receive money across borders without censorship or interception. More generally, decentralization offers enormous promise in many domains: democratizing access to assets; allowing people to own their identities and monetize their own data rather than having them be owned, profited from, and used for manipulation by technology companies; providing transparency and provenance for AI-generated content and for the ownership and transfer of assets in general; and enabling low-friction, low-latency, and low-cost transfer of value across the Internet among people and among computer agents.

While many opportunities are clear (and many others have likely not yet been imagined), there are significant challenges in realizing them. First, centralization has benefits that decentralized systems have not yet overcome. Centralization clearly allows undesirable effects for ordinary citizens, such as currency debasement, censorship in the transfer of value, high rents and latencies imposed by intermediaries, and behavior modification by corporations owning people’s identity and data. Yet centralization has important benefits for those same users as well: in convenience, performance, streamlined decision making, and recourse against user mistakes. Second, incumbent centralized entities resist the decentralization of power. Since these entities include governments, financial institutions, and large corporations, the resistance as well as the desire to co-opt the technology for centralized control will be massive. Third, as scalability appears fundamentally limited in secure, decentralized base-layer blockchains like Bitcoin, and since true privacy is challenging from a regulatory perspective, attempts to address these issues often lead to centralization, at least in certain parts of systems. Fourth, as is often the case with new technologies, there have been mistakes, false starts, and fraud. This problem is accentuated for permissionless blockchains because central to their functioning are incentives managed through “tokens” which have financial value. The difficulty of valuing these tokens and the lack of proper regulatory oversight have invited substantial financial manipulation and fraud. Fifth, while permissionless participation provides many benefits, it can be difficult to identify and restrict bad actors until after the fact. Finally, while applications with directly positive societal value are increasingly being developed, early applications have been dominated by trading and speculation. To an extent, such speculation is essential to enable investment in new technologies; however, much of it is akin to gambling.

There have been some spectacular realizations of these risks. Many of the most prominent—such as the collapse of FTX, Celsius, and Three Arrows Capital—were the result of abuse by *centralized* actors rather than failures of decentralized protocols. Others, such as the failure

of Terra stablecoin and associated Luna token, were due to systems with inherently faulty designs being sold to the public as innovative decentralized systems by their creators. In addition to these centralized failures, there have also been instances of malicious actors exploiting the properties of decentralized systems working as intended (e.g. the accused money laundering by North Korea using the Tornado Cash ‘tumbler’ on Ethereum).

The abuses to date point to a major policy question that must be resolved if we are to unlock the potential for decentralization: Key properties enabled by decentralization, such as permissionless participation and censorship resistance, are at once essential for human rights and other critical applications and yet—due to lack of easy recourse in decentralized systems and of a clear regulatory landscape—can be exploited for harmful behavior. Self-custody, which also requires decentralization, is attractive, especially in less trustworthy environments, but is currently difficult to manage and can lead to loss of funds. As in other areas of rapid technical advancement, such as artificial intelligence, opportunities and risks abound. In fact, history is replete with “bubbles” preceding actual technological booms, including with radio and the World Wide Web. In the case of blockchains, there is the potential to significantly benefit humanity through the decentralization of power; but there is also much to be understood, built, and iterated and many important policy, ethical and technical considerations to resolve. The goal should be to incentivize innovation that provides long-term, positive impact for society while minimizing fraud and manipulation, though the path to clear and durable solutions is complex.

We believe finding good solutions to these challenges requires a concerted interdisciplinary approach across technology, applications, economics, policy, ethics, and societal implications. For example, application needs, technical requirements, regulatory frameworks, business models, and social considerations must be taken into account in the design of successful blockchain architectures, even more so since their protocols are not merely technical but also economic and social. The success or failure of a decentralized ecosystem should be measured in outcomes that people actually care about (such as cost relative to value; autonomy, censorship resistance or transparency achieved; difficulty of changing important promises made; quality of service provided; freedom of choice and exit; etc.), not merely in “units of decentralization,” which are themselves difficult to define and measure. And in the end, success will depend on whether decentralized applications that provide clear positive value to society are built and achieve traction (just as happened with the Internet and Web), so a focus on such applications is critical at this juncture. The applications will inform the research that is needed in key

enabling technologies as well as explorations of the policy, governance and ethical considerations needed to optimize benefits and mitigate risks for different groups of people.

Every one of the three pillars—technology, applications, and societal implications—should be developed with adequate understanding of the others, and advances made in concert. Indeed, the DeCenter was established to foster this interdisciplinary approach, both across academia and in partnership with industry, open-source ecosystems, policymakers, and regulators. These partnerships are essential to achieve our aspiration of helping to set the proper intellectual foundations, goals, and frameworks across disciplines. We hope these frameworks will support the development of meaningful applications that leverage the benefits of decentralized trust and power in the service of humanity.

The remainder of this white paper outlines key questions across the three pillars where progress is much needed. It is structured in sections that are primarily about technology, applications, or societal implications, though the questions have many interdisciplinary connections.

Technology Questions

What does a first-principles approach to decentralizing power look like from a technological perspective?

While blockchains have been successfully used to create new monetary assets (Bitcoin) and application platforms (Ethereum), important technical considerations and tradeoffs in performance, capacity, privacy, and programmability remain open questions. Moreover, the ways in which the developments so far are good and bad for society, and how they can best be managed from the policy and governance perspectives, remain topics of discussion.

A first-principles approach to developing systems and technologies that decentralize power must begin by defining clear design objectives within and across domains. The needs of applications will help us define these. Many of the desired characteristics are similar to those previously applied to distributed systems, such as performance, scalability, capability, security, privacy, flexibility, and (inter-) operability. Some considerations are new, such as monetary properties, incentive goals, and social governance. Setting appropriate priorities and trade-offs among the objectives will guide the most important technological paths to pursue. Key questions towards a first-principles approach thus include the following. How do we provide the necessary system properties for distributed systems in the context of decentralization? How do we draw rigorous logical connections from the design choices made to the outcomes achieved? What can we learn about decentralized system design from the designs of other distributed systems and from economic mechanism design as used in other areas? Appropriate frameworks for analyzing the tradeoffs between design objectives in decentralized systems will go a long way towards developing systems that best serve the needs of the most important applications.

Along with key objectives, the key concepts in decentralized systems must also be defined. While the concepts commonly used in the context of blockchains—e.g., *trust-minimized*, *decentralized*, *permissionless*, *ensorship-resistant*, *transparent*—have intuitive meaning to people, the precise meanings vary across communities. Sociologists, economists, philosophers, and engineers define even basic concepts like *trust* and *power* in different ways. We must engage in cross-disciplinary efforts to clearly define and operationalize terms in the context of blockchains.

It is also essential to consider not just technological and economic mechanisms but also social ones, and how these interact, when designing decentralized systems. For example, it is not only the mechanisms for achieving consensus within a technical protocol that matter but also the social consensus that governs the software that participants run, the rules that the protocol employs, and how those rules can be changed. The technical and social consensus rules that govern Ethereum, the second largest blockchain by economic value, have changed multiple times since its inception, and changes desired by some to Bitcoin have led to separate chains like Bitcoin Cash via “hard forks.” Core questions associated with these social mechanisms include the following: If the decision to change the rules of an existing blockchain is made by a small number of core developers and major cryptocurrency holders then, regardless of the degree of decentralization of the blockchain’s technical trust mechanism, to what extent has it really decentralized power? If certain centralized entities have outsized economic power or other influence—such as stablecoin issuers, oracle providers, or rollup sequencers—how does that influence the decentralization of the underlying blockchain? And if participants democratically delegate social consensus to a single entity or group of entities, and can alter their delegation anytime, how centralized or decentralized is that solution? How should we design technical, economic, and social consensus mechanisms to best achieve goals while preserving decentralization?

What are the appropriate technical, economic, and social architectures for public, permissionless blockchains, and how can policymakers, technologists, and humanists work together to define them?

Existing blockchain systems employ a variety of approaches to their technical, economic, and social designs. From a technical perspective, layering, partitioning, and interoperability are common methods used in system design. These methods enable abstraction, standardized interfaces, rapid development, performance, analysis, and other properties. The design choices made among these methods impact decentralization and security; for example, through “weakest link” effects, risks of capturing a partition more easily than the whole, and security holes in programs. A layering framework is already emerging for blockchains, encompassing the hardware, networking, consensus/settlement, data availability, execution, smart contract, and application layers. Some blockchains remain monolithic (e.g. Solana), some embrace layering for scalability (e.g. Ethereum), and some promote highly modularized and plug-and-play ecosystems (e.g. Cosmos), while

Most users do not care about decentralization in itself. But people care about, and applications benefit from, important properties that rely on decentralization, including permissionless participation, censorship resistance, transparency, immutability, and self-custody.

the largest blockchain (Bitcoin) focuses on simplicity and security and requires off-chain layers to provide meaningful programmability. These vastly differing properties of existing blockchain systems raise the following key question: How should system architectures for public permissionless blockchains be designed, and their consequences articulated, in a principled manner to achieve the objectives defined by a wide range of stakeholders?

A unique characteristic of blockchains is that they integrally include not just a technical architecture but also an economic design (e.g., monetary, marketplace, and incentive design) and a social governance design (how protocol upgrade decisions are made, and how attacks are responded to). As layers become separated, these economic and social aspects can apply independently to different layers. Designers must understand how incentives and social governance at one layer can influence behavior, including motivating malicious behavior, in another. For example, a centralized network provider can undermine an entire consensus protocol without violating any consensus-layer security assumptions. Protocols from one layer “enshrined” in another or dominating its activity can undermine the interests of other applications and layers. Consensus participants may behave strategically to reap rewards at the application level (via “Maximal Extractable Value”) without profiting via consensus-layer rewards. In the other direction, designers may build applications, such as restaking (in which economic stake used to secure one protocol, e.g. Ethereum, is reused to also secure other protocols, e.g. application services) or liquid staking (in which economic stake “locked” to secure a protocol is made available for other purposes at the same time), whose entire utility derives from increased rewards from other layers. This raises the question of how to design the economic, social, and technical interplay of modular, independently operated layers for maximum benefit and minimum risk.

Important capabilities will be driven by application needs. Permissionless blockchains were developed for peer-to-peer money, and a lot of the early applications are in finance. However, the properties they offer can be very valuable in many other applications. These include applications for which identity, transparency, and specific compliance are critical, such as public records and government applications, and those in which identity and data ownership are also central, such as healthcare and social networks. We must, therefore, ask what additional capabilities blockchains might need (e.g. privacy, proof of humanity, attestations, etc) to best support these application areas.

In fact, an important question regarding blockchain design in the context of application needs is sometimes framed as “money versus technology:” Is the focus of a blockchain and its ecosystem to support a peer-to-

peer decentralized asset or currency, or is it to provide a programmable technology platform for a variety of applications? For a permissionless system, even a technology platform requires a monetary asset, but the question is whether the primary focus is one or the other, or both. Architectures for the former would prioritize security, simplicity, and decentralization at the cost of programmability, speed and scalability, and would have a very high barrier to change. Bitcoin does this, leaving it to external, off-chain systems or layers to provide the latter capabilities. Most other blockchains, including Ethereum, are built with the latter focus, and they have been changing often, even though their plan is to largely “ossify” in the medium term. How should the architectures of blockchains evolve in the “money and technology” landscape, and can or should we aspire to achieve unified solutions? In considering these questions, it is worth noting that as Ethereum has grown it has often staked its claim to being a monetary asset as well, one that derives value through the usage of the chain (through its “burn” mechanism); and, at the same time, off-chain layers are attempting to bring scale and programmability to Bitcoin due to its economic power and robustness as an underlying settlement layer for transactions.

Some desired properties inherently conflict: a fully private system cannot also be fully transparent. Similarly, perfect censorship resistance may be incompatible with the needs to prevent behavior that may be harmful or even illegal in some jurisdictions. Decentralized systems aim to remove power from any single entity, but this also means there is no centralized entity that can provide recourse in certain situations (e.g. when someone’s private key is lost). No technology can circumvent these tradeoffs, so policymakers and citizens must prioritize among tradeoffs to best serve societal ends. However, it is important for policy discussions to be informed by an understanding of what technology-driven results are even possible. As one example, technologies such as proofs-of-reserves and zero-knowledge proofs can enable transparency regarding some properties of the system or application (e.g. solvency of an exchange) while respecting privacy of individuals (e.g. personal account balances). While users and policymakers must determine which properties are necessary for consumer protection, it is a technical question to determine what level of individual privacy can be supported within these constraints. Privacy is not all-or-nothing: It is possible to keep some information about transactions private, some public, and some available only to regulators or law enforcement. The privacy needs of individuals and of companies (such as exchanges) may also be different. These design considerations raise the following questions: How can disciplines work together to find the Pareto frontier of these conflicting goals, and to ensure appropriate design of the technical, economic and social architectures together, for different layers, especially in an interoperable world? How should these

architectures compose with one another to preserve good properties and be easy to reason about?

Finally, the history of decentralization has often been one of “recentralization” or “centralized value capture” elsewhere. This has been observed historically in the social sciences, and it is also true of the Internet itself, where the protocols underlying the Internet (like HTTP, SMTP, etc.) are decentralized but power has accumulated to centralized corporations at the application layer. Some argue that the incentive mechanisms of blockchains—in particular, the ability to capture economic value at the protocol layer—will help preserve decentralization. At the same time, developments like stablecoins, exchange-traded funds (ETFs), ownership concentration, and centralized cloud infrastructures for blockchains are already pointing toward centralization. A key question for decentralized system design is thus: What additional mechanisms can be put in place to ensure that re-centralization of power in other layers or entities is not the ultimate outcome?

What are the most effective consensus protocols to truly decentralize trust and power?

Decentralized consensus protocols, in which independent, permissionless participants reach agreement on the content of a shared ledger without the help of a centralized organization, are the key technical backbone underlying modern permissionless blockchain technology and the key source of decentralization (at the cost of performance). Starting with Bitcoin’s proof-of-work consensus in 2008, many consensus protocols have been designed, including those that use the proof-of-stake mechanism such as the one adopted by Ethereum in 2022.

Numerous paradigms exist for the design of consensus protocols, and further study is necessary to understand the downstream implications of each. Some comparisons between the proof-of-work and proof-of-stake approaches are generally accepted: proof-of-work consumes significantly more energy than proof-of-stake, but results in simpler protocols with a smaller and well-understood attack surface. Nevertheless, strong opinions abound, and ambiguity remains over which of these paradigms most effectively enables decentralization. Key questions in this domain include the following: Do miners with access to cheap electricity gain outsized influence in proof-of-work protocols? What about players with greater access to capital in proof-of-stake protocols? To what extent do mining pool operators or dominant producers of mining hardware accumulate power in the former paradigm, and “liquid-staking” providers or dominant centralized exchanges in the latter? Which paradigm composes better with other necessary layers, and are there hybrids or

variants that should be used in those contexts? Are there characteristics other than external work and internally denominated stake that can serve as better or complementary bases for consensus protocols? Are the protocols needed to support a monetary asset (like Bitcoin) different from those needed in platforms designed to build applications other than the assets themselves—such as decentralized finance, asset tokenization, provenance, artificial intelligence, or social networks—or are there ways to compose protocols in layers that can support both money and other applications well? What about the trend toward reusing crypto-economic security or proofs from one blockchain to secure other chains, layers, or applications? How much power can dominant applications (which may be centralized) exert indirectly over the underlying protocols of different types? These questions are crucial to understanding the extent to which blockchains truly decentralize trust and/or power.

A related but vital question is the age old one: Who decides which consensus protocols will succeed and fail? Designers, users, and government officials (at the state, federal, or international level) all will likely have a say. But the relative power of these groups to determine the future of this technology, the timing of interventions, and the power-sharing arrangements among them are crucial questions.

How do we measure decentralization and its impact on systems?

Measuring the right quantities appropriately with regard to decentralization is also a major challenge. For example, many proponents care deeply about any user with a laptop and an Internet connection having the ability to run a node, although running an effective Bitcoin miner requires customized hardware (that is, the ability to create blocks is effectively proportional to the amount of customized hardware one operates), and running an Ethereum validator requires a large amount of capital to be locked up in the protocol (that is, voting power is effectively proportional to the amount of ETH one stakes). This raises the following questions. How do we measure how decentralized a protocol is, and hence compare systems in this regard? How do straight-forward assessments (such as numbers of miners/nodes/validators, distribution of voting power, geographic diversity, etc.) relate to meaningful outcomes (such as resources required to double-spend a coin or to recover from an attack)? How do we measure how decentralization composes across different technical, economic, and social protocols that constitute a system?

Another important question is the following: How do we assess where and in which layers the degree of decentralization truly matters? Different protocols consume different levels of resources, e.g. electricity,

network capacity, and other capital, with different characteristics (i.e. up-front versus ongoing cost). For instance, critics of Bitcoin point to its high energy consumption, while advocates suggest that the incentive for miners to use the cheapest forms of energy actually promotes the development of green energy. Permissionless participation and individual incentives makes resource consumption complex to measure and reason about, resulting in a battle of qualitative argument and opinions. Proper measurement is essential to reach correct conclusions and develop the best designs. It also requires expertise from multiple disciplines: to design and run randomized controlled trials, to evaluate the influence of incentives, to consider various ethical implications, and to measure complex protocols. More generally, it is important to understand where progress in blockchains can leverage past work in existing technical, social science, and humanistic domains, and where the challenges are fundamentally new ones.

Application Questions

What are the properties of decentralized systems and applications that have a real impact on users? What exactly is decentralization good for, when, and for whom?

Decentralization as a standalone goal is rarely of interest to users, who are typically much more concerned with the properties that decentralized architectures might support: permissionless participation, censorship resistance, transparency, lower rent or fees, etc. However, decentralization comes with potential drawbacks for users, including performance, scalability, ease of use, and a lack of certain types of convenience and recourse (e.g. upon mistakes or theft). Other technologies aim to address performance and scalability, but they compromise aspects of security and censorship resistance. For instance, rollups with centralized sequencers may offer comparable transparency and automation, but they lower the barrier to censoring transactions.

Given these tradeoffs, it is important to rigorously consider which properties applications really need, and what the costs and benefits of various technologies, such as blockchains, are to different classes of applications. What, for example, are the consequences to users of no single entity controlling an application or its data? Will prices be lower, due to the lack of a powerful monopolist? Or higher, due to technology overhead? Will innovation be better incentivized, due to the lack of a

strong incumbent, or less so due to lack of agency by any particular entity? What users *really* want may be best fulfilled by a thoughtful, clearly identified mix of decentralized and centralized components, together with a transparent understanding of exactly what properties of interest the system possesses. Frameworks for articulating and assessing these properties, in light of user and application needs, are daunting but critical. The questions apply to decentralized systems beyond blockchains, though the latter are the focus of this paper. As noted above, a meta question here is who should decide these questions, and when.

What are the greatest opportunities and challenges for applications that can decentralize power?

The Internet and the World Wide Web enabled applications that changed the way we live and work. They did this by transforming communication and access. Decentralized technologies seek to transform trust, with the hope that this can empower people, advance autonomy, and reduce the power and control that have accrued to centralized entities. This raises important questions such as: What applications truly benefit from the decentralization of trust and provide lasting, high value to society? That is, what exactly is this decentralization good for, when, and for whom?

There are contexts in which the opportunities for decentralization of power are straightforward. For example, censorship-resistance and permissionless participation (without “know your customer” requirements) are sometimes first-order imperatives. They can provide human-rights value in countries with authoritarian regimes by allowing self-custody and permissionless, peer-to-peer transfer of money or assets within and across national borders, reducing the power that governments wield over citizens’ assets. Especially in autocratic or dysfunctional regimes, these benefits can be life-altering for people who otherwise might be vulnerable to capricious confiscation or financial ruin. And they can greatly reduce latency, cross-silo friction, and fees taken by intermediaries in the movement and settlement of assets, especially across borders, reducing the power of banks and enabling 24/7 global transactions. Moreover, permissionless participation, even while producing uneven results, may raise universal living standards because anyone can participate and have access. As the Web and Wikipedia elevated public access to knowledge, cryptocurrencies hold the potential to increase financial inclusion in areas badly served by present regimes. But, aside from these applications of basic financial infrastructure for those who wouldn’t otherwise have it, how do we identify other opportunities for decentralization of power to have significant positive impact?

What applications truly benefit from the decentralization of trust and provide lasting, high value to society?

Decentralizing power through permissionless blockchains also brings a unique set of challenges. The lack of a centralized intermediary creates challenges in providing users with recourse to mitigate the risk of loss, theft, or other problems. How, and at what layers of the system and application stacks, do we address these issues? It is also true that while self-custody has a high technical barrier and requires significant risk tolerance, without self-custody assets are given to a centralized actor to hold and control, subjecting them to similar hacks, theft, and censorship as existing, non-blockchain solutions (and depending on the surrounding legal landscape, perhaps even more so). Another challenge is that cryptocurrency assets are difficult to value at this stage and tend to be less regulated, creating opportunities for bad actors and malicious behavior. Such opportunities also exist and are exploited in traditional, centralized finance, but centralized finance has greater controls and many time-tested, stable applications. Decentralized Finance (DeFi), the most prominent area of application for blockchains so far, has the potential to disrupt and improve finance, but it is commonly criticized for being “self-referential” within cryptocurrency assets and of a risky, gambling nature.

This ethical and legal ambiguity accentuates the question of developing applications that provide genuine long-term value. A lot of the early activity in the “crypto” market has been about the trading of new assets and has led to financial harm to many retail users through greed and information asymmetry. More substantive applications are being pursued: to make the largest assets productive, create new financial products, tokenize and democratize “real-world” assets, record provenance, refactor governance, and create new business models to replace targeted advertising. There is not yet a “killer app,” although stablecoins (tokenized dollars or other fiat currencies) have shown important success, especially in countries with unstable currencies, and undoubtedly there will be many false starts and mistakes as new applications are developed and tried out in the market. While many areas require permissioned and not decentralized systems, we should explore what role decentralized blockchains might have even in centrally controlled systems, such as tokenized sovereign debt and government assets, trade finance, and digital fiat currencies.

Many other promising areas of investigation have had slower uptake so far. Individuals owning and controlling their identities and online data can diminish the power of centralized corporations like large technology and social media companies. However, decentralized identity, control of data, and social networks have so far been slow to displace more convenient, centralized incumbents, perhaps due to the strong network effects of the latter. Decentralized games, with interoperable NFT (non-fungible token) assets, are often slow, and both “content NFTs” and “utility NFTs” have

had volatile swings in popularity. Decentralized autonomous organizations (DAOs) are still quite nascent.

At the same time, important new areas of application continue to emerge. For example, artificial intelligence, especially the use of blockchains to establish provenance of machine- and human-generated assets and models, to promote trust and safety, and to provide machine-to-machine payments for AI agents, is an exciting emerging area of application. As in the early phases of the Internet and Web, many important blockchain-based applications are likely yet to be imagined. Some argue that the major applications will come when the infrastructure is easy to use and inexpensive, others believe that compelling applications are essential to motivate the need for the technology, while still others believe the growth of speculative applications will fuel the development of more lasting ones. Regardless of timing and sequencing, a focus on applications that can have a positive, long-term impact on society, and on their implications for the technology as well as for policy, societal and ethical issues, is critical to fundamental progress in the field.

What do the characteristics of applications imply for the levels of decentralization needed, both in the application itself and in the underlying infrastructure?

As applications are developed to serve a diversity of use cases, care is needed to understand what properties resulting from decentralization are most critical to which applications, and in what aspects of the applications (e.g., the data, the transaction records, the execution, the user interface) decentralization is important. This helps us design better systems, and, in the long term, applications decentralized for the sake of using blockchains are unlikely to be successful.

For example, if an application primarily relies on data availability rather than timely and properly-ordered updates, a user might be fully comfortable trusting a more efficient Layer-2 blockchain that periodically “rolls up” application data to a robust Layer-1, even if a centralized entity controls the sequencing of transactions on that Layer-2 chain. By contrast, the same decision might be a dealbreaker for an application that truly values decentralized sequencing and transaction ordering. How should application developers reason about decisions like these and clarify to their users exactly what they are getting? How should system designers learn from these assessments, across a range of applications, how best to design the layers of their systems?

How can applications and users be made more clearly aware of the decentralization implications of complex underlying infrastructures? What can centralized applications safely inherit from underlying decentralized infrastructure?

Seemingly mundane implementation decisions can have a huge impact on where power lies. For instance, Layer-2 chains are used to provide increased performance and lower fees on underlying Layer-1 chains like Ethereum (and to provide smart contract programmability itself on Bitcoin). Smart contract code for applications executes on the Layer-2 chain, and transactions finally settle on the underlying Layer-1 chain. In a system like Ethereum with Layer-2s, the native smart contracts of applications could run entirely on the Layer-2 or could use a Layer-1 native smart contract on Ethereum as well. The difference in arrangements is subtle, but it can have decisive implications for censorship. In the former setup, the application cedes to the Layer-2 system the power to forever censor its interactions with its smart contract. In the latter, it does cede the power, but it can also circumvent the Layer-2 and resist censorship (by finding another Layer-2 to run on or by posting transactions directly to the Layer-1). In other words, the specifics of how systems are built and how they are used by applications can have major implications for the extent to which the desirable properties emanating from decentralization are achieved in practice. Especially as systems become more complex and layered, the question becomes more pressing: How should we develop frameworks to systematically and clearly highlight such centralization implications?

Decentralization in a blockchain protocol (technical or even social) does not imply decentralization throughout a system from the user perspective. Consider centralized exchanges and marketplaces for decentralized assets such as cryptocurrencies and NFTs. On the one hand, these allow ordinary users to easily outsource custody and transaction settlement for these assets. On the other hand, these centralized entities have substantial power to commit classical forms of insider malfeasance: Assets may be stolen, lent out without permission, or subject to arbitrary corporate censorship. How should applications decide which aspects of their service to retain centralized control over, and which to decentralize? Similarly, blockchains can inherit centralization properties from underlying networking infrastructure due to the centralization in entities like border gateway protocol (BGP) routers. How can the impacts on power of such inheritance be made transparent to users?

The interactions among applications, systems, governance, and technical and social consensus are complex and have real implications for what

users might expect and what they actually get. This is especially true as systems become more layered and complex, because of how the decentralization, trust, and power allocation properties of different layers interact. Analyzing and articulating these effects from a user perspective is a difficult but very important undertaking.

Social, Ethical, and Policy Questions

How does a “first-principles” approach both raise and help answer questions about the role of trust, permission, hope, decentralization, power, value, and transparency in specific social and ethical contexts?

Philosophers and social scientists have generated working analyses of many of our key concepts (e.g. “trust,” “decentralization” and “power”). Do these analyses fare well when applied to decentralized or blockchain contexts, or do these contexts require us to amend widely-held assumptions about how trust, value, and power work in society? In a similar vein, do traditional ethical frameworks unproblematically apply to the kinds of interactions that are enabled by blockchain technologies, or do blockchains generate novel use cases that necessitate revisions to those frameworks? What kinds of hopes does the potential for decentralization of trust and power tend to engender in people, and to what extent are such hopes rational or actionable?

Does the emergence of cryptocurrencies, NFTs, and incentive-driven systems provide new insights into the nature of value? Or are the concepts of value similar to those obtained in equities, market dynamics and exchange, or perception and speculation? Is there value for an artist to having their work tokenized beyond that derived from direct economic consequences (such as potential for increased demand or ongoing royalty payments)? Is it easier for investors to “hodl” cryptocurrencies and NFTs through violent market gyrations and even catastrophic losses (measured in fiat terms) than it is for them to hold other kinds of assets? If so, is this because they somehow measure the value or potential value of these token assets differently?

Finally, who or what is ethically responsible for the operations (and unforeseen consequences) of something that has no centralized governance or agency behind it? Can a blockchain (like a corporation)

Technical developments advance the state of the possible, but it is up to policy-makers and the people to determine how technologies best serve society.

sometimes be conceived as an ethical agent in its own right? Can a governing DAO be considered thus, and if so to whom does that ethical responsibility accrue? What would be the legal implications of one or another way of answering these questions?

What are the psychological, social, and ethical effects on people of interactions with something they perceive as permissionless, trustless, and decentralized?

The resilience of an ethical commitment is measured by the extent to which people tend to stick with it over time and apply it in new contexts, rather than defecting or deviating. How do human beings tend to react ethically to engagement with the blockchain? For example, do people who would normally be unwilling to exploit, defect, and steal from other people (or the government) retain these dispositions in web3 contexts, or does the pseudonymous, “wild west” character of NFT marketplaces, metaverse transactions, and crypto exchanges tend to bring out worse characteristics in people? On the other hand, do people stick with blockchains and cryptocurrencies longer than they would with other systems and financial instruments due to an abiding belief in their underlying value to society?

What are the psychological, ethical, and political results of state agencies encouraging (or mandating) the use of blockchain technologies? How do people react when they are compelled to trust a decentralized technology? For instance, what does El Salvador’s top-down decision to force businesses to accept Bitcoin as legal tender teach us about the benefits and risks of state encouragement of cryptocurrency (as opposed to the typical case where the state acts primarily as the source of constraint and regulation)?

Engaging questions like these require substantive interaction among people working in disciplines like philosophy, ethics, political theory, and law and those working in more empirical domains like psychology, political science, and engineering.

How do we best document and measure the impact of new, decentralized systems like blockchains on society, especially in different countries and on different segments of populations?

The permissionless, cross-border, financial nature of blockchains suggest that applications built on top of them will have different impacts across different countries and population segments. For instance, people in countries with less stable governments or currencies will view the benefits differently than those in major Western democracies. Across nations, members of disenfranchised groups may see cryptocurrencies as a path to financial safety or improvement, while members of these groups less exposed to sophisticated financial systems may be at greater risk of fraud. Centralized power, whether in the form of repressive state governments or redlining by banks and the federal government, is associated with discrimination and abuse, while also creating programs to help out the less fortunate. Decentralization, as with all innovation, creates both opportunities and risks. These effects are not well studied, and appropriate frameworks and methods for analyzing these impacts will be critical to ensuring that utility is maximized and harm mitigated.

Of course, blockchains are not (nearly) the first technology with disparate impacts on different populations. Even recently, in the fast-growing fields of machine learning and artificial intelligence, bias in algorithms and data has been established and sub-disciplines including “AI Fairness, Accountability and Transparency” have arisen to understand these impacts and ensure public accountability. Disparate access to technology has also been studied, for example in the case of the Internet and the “digital divide.” Blockchains bring in financial incentives and social governance more explicitly: What impact will these factors have on blockchains and their adoption in different parts of the world? Questions of who is benefited and harmed by different distributions of power and authority are well studied in politics. Political scientists model the interactions of voters, legislators, and parties to predict how policy outcomes benefit different individuals and groups. Within these frameworks, they vary the authority of different actors to make decisions and trace how those reallocations affect policies and the social distribution of resources. This raises the important question: What can we borrow from established domains in studying blockchain systems, and what challenges may be unique to measuring the impact of decentralizing technologies? It is likely that in addition to quantitative and qualitative human-centered studies, proper technical measurement of systems and their governance will be important in evaluating the societal impact of decentralization. We must define and develop cross-disciplinary frameworks to enable this analysis.

Permissionless blockchains are inherently global. Regulation and policy tend to be country-specific, if not more local. Can we, and should we, develop global regulatory and policy frameworks for blockchains? Is there a way to factor the frameworks so they have globally and locally determined elements? Finally, what is the impact of culture: are there cultural barriers to adoption, how can they best be examined, and can blockchains as digital ledgers help in preserving the diversity of culture and heritage?

How do we choose among conflicting goals such as privacy and transparency, censorship-resistance and compliance, decentralization and recourse?

Technical developments advance the state of the possible, but it is up to policymakers and the populace to determine which technologies best serve society. Some regulatory and policy issues that can impact tradeoffs between innovation and safety—such as which cryptocurrencies or contracts constitute securities and which commodities or something else, and what types of transactions constitute securities transactions—have been publicly discussed, although not resolved to general satisfaction. Other issues have received less attention, so many important questions remain. For instance, how well do privacy-preserving auditing technologies address demands of both privacy advocates and regulatory bodies concerned about illicit activity? How should misbehavior by well-known companies or entities be treated differently from that by individuals, and how are the privacy considerations different for the two cases? How can we speed up certain types of recourse and ensure consumer protection more broadly in decentralized environments without relying only on the court system? Can we separate any inherently conflicting demands from those that better technology might plausibly address? For truly opposing tradeoffs, how do policymakers choose which side to prioritize?

An equally pressing challenge is for the industry to “get its house in order” and adopt technologies that seem to have broad support. Specific questions in this area include the following. How should technologies like “Proof of Reserves” and “Proof of Liabilities,” that gained support as ways for exchanges or stablecoin issuers to demonstrate that they have a certain amount of assets and liabilities, but then had interest in them wane, be comprehensively evaluated as potential solutions to avoid surprise insolvency? Once technologies are broadly supported, when is regulation necessary to ensure they are adopted, when should standards bodies be developed to encourage or essentially enforce them, and under what conditions do large, informed players have the

incentive to self-enforce them? How are the answers to these questions different for decentralized ecosystems or exchanges compared to traditional centralized corporations? To what extent should governments regulate the on-ramps and off-ramps to blockchains (such as exchanges, wallets, and custodians) versus the blockchain protocols, applications, or governance themselves?

Another important question is how we should think about the design of regulatory institutions and the interaction among existing ones, innovators, and the public. The answers are complex and the systems are different across countries. In the United States, turf wars rage regarding the jurisdiction of regulatory agencies to enforce or de facto set policies in a range of areas that will influence the development of decentralization technologies, and the positions of the legislative and executive branches on both policies and jurisdictions are unclear. Other countries have clearer policy and enforcement positions. It is essential for countries to have coherent policies on these questions, although erring on the side of caution for the sake of taking a position may end up stifling innovation or ceding it to other countries. A careful balance must be struck, and power properly allocated among various actors inside and outside the government.

How should governance be best provided for decentralized blockchains?

Regardless of how technically decentralized a system is, it has a governance layer above it. What should this governance layer look like? Open-source systems provide a foundation, as the major permissionless blockchains are open-source, as does experience building “credibly neutral” systems. We have discussed in a previous section the social consensus governing blockchain protocols. But who is in charge of these software projects, what are their rules and procedures (e.g. voting and change management rules), how do they make decisions, and how do they avoid the projects being captured by powerful centralized players and losing their decentralization in the future? A key challenge is creating the right incentives for people to want to keep contributing to open-source projects, as well as incentives for people to lean in and help govern them.

Decentralized autonomous systems (DAOs) are being studied as automated governance systems using a decentralized approach. But they also require active participation (e.g. voting), and it is unclear what can be automated and what must require human involvement. In general, while decentralized governance is attractive, and appropriate governance can enable decentralized systems to remain decentralized, the lack of ongoing participation by users in governance systems can easily lead

to regulatory or corporate capture. What kind of incentive systems, technical mechanisms, and regulatory frameworks are needed for good decentralized governance; for example, how should DAOs be regulated? Governance applies not only to Layer-1 blockchains but also to other layers and to individual applications, so it must be considered carefully at all levels and the composition of governance is itself a challenge. Governance is a rich and complex area that must be examined in an interdisciplinary manner. As we design better governance systems for blockchains and their applications, what can we learn about how to design better governance systems more generally for society in the future, including for other credibly-neutral, decentralized online platforms as well as offline governance systems?

What are the key legal questions that should be addressed?

As in all aspects of society, the legal system has an important role to play in mitigating harms in decentralized systems and applications. Such systems and applications present unique challenges, however, because in the pure case, there are no specific individuals or entities in charge of them. Some countries are using existing legal frameworks for blockchains and cryptocurrencies (e.g., much US enforcement follows from existing securities law), while others are developing novel frameworks (e.g., the EU developed the Markets in Crypto-Assets or MiCA regulations). To what extent can we leverage principled pre-existing frameworks, and what aspects require novel policy-making and regulation? What aspects are best addressed by ex-ante regulation (i.e. rules and standards) versus ex-post regulation (i.e. enforcement)?

An interesting question is how to assign responsibility and provide recourse when a decentralized system is used by harmful actors to disguise questionable activity (as was seen with the TornadoCash smart contract). How should legal accountability, liability, and recourse be distributed among the wide range of varying participants: developers and maintainers of the protocol itself, users running nodes that help maintain the system, builders of the blocks containing the harmful transactions, validators of the blockchain, token holders in DAOs who govern it, other users of the contract, and organizations that promote the contract? Accountability may increase the trust inherent in the system and spread losses in efficient ways, but it may also raise the costs of participation in ways that are prohibitive.

At the same time, where do we draw the lines of tolerable risk within the limits of the law? Is there an appropriate scale of ‘outlaw’ behavior that might be tolerable, or that might even be beneficial to

enable experimentation that is too risky for large-scale, mainstream systems? How should such experiments be conducted: Allowing truly permissionless participation in them may be important, but is there a way to contain such experimentation within large ecosystems without bleeding into the overall ecosystem at scale? Regulatory sandboxes in which otherwise suspect activities are permitted in order to learn are increasingly being used in other domains: Can they be valuable here?

How can we ensure that decentralized technologies have a positive impact on society?

Permissionless participation has advantages, but it also has the potential to favor those with an abundance of resources over those without, due to asymmetries in awareness and risk-taking ability, costs of participation in activities like mining and validation, and lack of external support to encourage level playing fields. What mechanisms can be put in place to promote the ethical use of decentralized technologies, especially those that mitigate inequality? While wealth and access inequality are first-order concerns, second-order effects can be important as well. For example, if proof-of-work consensus, used in Bitcoin, indeed has negative environmental consequences, climate change is already known to disproportionately impact already-vulnerable populations. Yet proof-of-work has many known positive characteristics and effects (the same can be said of other modern technologies, such as artificial intelligence). Are there ways to make blockchains more energy efficient without compromising the security and decentralization that proof-of-work can achieve? How can the potential predictable consumption of excess energy by proof-of-work blockchains be used to support environmental sustainability, to reduce or stabilize energy waste, or to incentivize green energy? How do we truly assess these effects and tradeoffs in comprehensive ways and come to meaningful conclusions about environmental impact?

Some questions are fundamental. How might decentralized trust and power, and the properties that result from them, impact social norms and human relationships? For example, what philosophical issues are raised by blockchains creating immutable records of human activity? While blockchains can afford pseudonymity, they also provide transparency: What are the ethical implications of potential mass surveillance enabled by the technology? Since power-seekers will always seek power, how do we prevent large actors from accumulating power over a seemingly decentralized protocol, whether this power accumulation is 'on-chain' (i.e. wealthy corporations purchasing significant fractions of mining/staking power or institutional investors owning significant fractions of initial token distributions) or 'off-chain' (i.e. controllers of large applications de facto

governing underlying decentralized infrastructure). Other questions lie more in the details; for example, what is the appropriate balance of on-chain governance (rigorously specified in a protocol) versus off-chain governance (sometimes referred to as "social consensus")?

These questions, as well as other fundamental ones, arise even when the systems behave exactly as intended. For example, are democracy and decentralization a natural fit or not? Decentralization, by design, creates a power vacuum. This can be exploited by malicious actors in ways that don't violate proper technical functioning of the system. For instance, anonymity or pseudonymity together with the lack of central oversight and constant monitoring may be used to more easily launder money or to fund terror organizations across national borders using decentralized systems (even though such efforts can also be tracked and sometimes caught). Lack of proper regulation can enable retail investors to be exploited through "pump-and-dump" schemes. These effects are not what most people want. Of course, money laundering, terror funding, and securities fraud are also possible, if not common, using cash, securities, gold, and countless other well-accepted tools. A key question in assessing the risks in decentralized systems should be: compared to what? For instance, modern societies use but do not constantly monitor cash because the net benefits are perceived to exceed the risks. In what ways can decentralized technologies like blockchains be best used to promote democracy, and what policies need to be in place to ensure that they don't undermine democratic institutions and principles?

Concluding Thoughts

Much work remains to examine the decentralization of trust and power through blockchains. The issues are complex, and addressing many of them requires highly interdisciplinary collaborations. The DeCenter was established to promote such collaborations and to explore fundamental, long-term issues from first principles, while leveraging what has already been learned in other areas of technology, the social sciences, and the humanities. We aim to find ways to properly evaluate decentralization for different types of applications and purposes, to pursue application opportunities of long-term value to society, and to mitigate risk. Our ultimate goal is to create—and aid the creation of—technologies, systems, social structures, and policies to leverage the benefits of decentralization and blockchains while managing their economic and other risks, and to help lay the cross-disciplinary foundations for long-term, first principles exploration in this field.

We are grateful to the many people who contributed to this document, and look forward to a rich journey ahead. To stay up to date on the DeCenter's activities, please visit our website at <http://decenter.princeton.edu> where you can find links to upcoming events, our mailing list and our social media channels.