

CRYPTOCURRENCY-BASED LAW

Michael Abramowicz*

Bitcoin is a protocol promoted as the first peer-to-peer institution—an alternative to a central bank. The decisions made through this protocol, however, involve no judgment. Could a peer-to-peer protocol underpin an institution that makes normative decisions? Indeed, an extension to the Bitcoin protocol could allow a cryptocurrency to make law. Tacit coordination games, in which players compete to identify consensus issue resolutions, would determine currency ownership. For example, an issue might be whether a cryptocurrency-based trust should disburse funds to a putative beneficiary, and the game's outcome would resolve the question and result in gains or losses for coordination game participants. A cryptocurrency can also be used to generate rules or other written codes. Peer-to-peer law might be useful when official decision-makers are corrupt or when agency or transactions costs are high. A modest starting point for cryptocurrency-based governance would be as a replacement for Bitcoin's centralized system for changing its source code. A cryptocurrency incorporating tacit coordination games could serve as a foundation for other projects requiring peer-to-peer governance, ranging from arbitration to business associations, which would enjoy inherent limited liability and would lack designated management.

* Professor of Law, George Washington University. I am grateful to David Abrams and Omri Marian for helpful comments and to participants in workshops at Boston College, George Washington University, and the University of Pennsylvania Law School. All errors are my own.

TABLE OF CONTENTS

INTRODUCTION.....	360
I. THE THREE CORNERSTONES OF PEER-TO-PEER GOVERNANCE	368
A. The Decentralized Ledger.....	371
B. The Decentralized Decision	374
C. The Decentralized Fisc.....	376
II. PEER-TO-PEER GOVERNANCE FOR CRYPTOCURRENCIES.....	384
A. Checkpointing.....	385
1. Resolution Without Tacit Coordination	386
2. Resolution with Tacit Coordination	390
B. Evolution of the Reference Code	396
C. Rewarding Institution-Promoting Activities	402
III. THE POSSIBILITIES AND PERILS OF PEER-TO-PEER GOVERNANCE	404
A. Peer-to-Peer Arbitration.....	405
B. A Peer-to-Peer Trust.....	408
C. A Peer-to-Peer Bank.....	411
D. A Peer-to-Peer Business Association.....	414
E. Peer-to-Peer Public Law.....	416
1. A Peer-to-Peer Central Bank.....	416
2. Other Public Institutions.....	418
CONCLUSION	419

INTRODUCTION

Bitcoin, described by its promoters as “an innovative payment network and a new kind of money,”¹ has attracted extraordinary attention as a financial innovation.² This attention results less from the functions that Bitcoin serves as a digital medium of exchange and store of value³ than from the decentralized nature

1. See BITCOIN, <http://bitcoin.org> (last visited Feb. 22, 2016). For a thorough introduction to the technology and economics of Bitcoin, see ARVIND NARAYANAN ET AL., BITCOIN AND CRYPTOCURRENCY TECHNOLOGIES (2016), http://d28rh4a8wq0iu5.cloudfront.net/bitcointech/readings/princeton_bitcoin_book.pdf.

2. See, e.g., Rob Wile, *10 Financial Innovations that Are Changing the World More than Bitcoin*, BUS. INSIDER (Jan. 30, 2014, 2:41 PM), <http://www.businessinsider.com/10-financial-innovations-more-exciting-than-bitcoin-2014-1> (identifying innovations in payment technology purportedly more important than Bitcoin, despite receiving far less publicity).

3. Currencies are generally thought to fulfill three functions: medium of exchange, store of value, and unit of account. See, e.g., N. GREGORY MANKIW, MACROECONOMICS 80 (7th ed. 2010). Critics maintain that Bitcoin has fulfilled the exchange and store of value functions poorly and has not served the unit-of-account function at all. See, e.g., David Yermack, *Is Bitcoin a Real Currency? An Economic Appraisal* (Nat’l Bureau of Econ. Research, Working Paper No. 19747, 2013), <http://www.nber.org/papers/w19747>. Its value is so volatile that few, if any, commercial

of Bitcoin transactions. Unlike traditional currency and financial instruments, Bitcoins are not issued by a central bank. Rather, anyone can attempt to “mine” Bitcoins by using computers programmed to guess answers to a computational puzzle.⁴ Bitcoin is thus neither a *commodity currency* (backed by gold or some other commodity) nor a *fiat currency* (used by convention as a result of a legal edict).⁵

Bitcoin’s independence from central authorities helps explain the perception that it is a technological marvel. Bitcoin functions even though it is a protocol without a referee. Of course, other protocols operate with minimal supervision; the Internet does not require police officers to arrest those who violate the rules of TCP/IP.⁶ But what makes Bitcoin remarkable is that it settles the most controversial issue—who owns wealth—without need for a law enforcement apparatus. Bitcoin can be seen not just as a currency, but more grandly as an institution that creates and enforces property rights. It is an institution, however, that can resolve only one type of decision: whether purported transfers of Bitcoins will be validated and added to a list of approved transfers, known as the *block chain*.⁷ If this is libertarian nirvana, it may seem to expose the limits of what peer-to-peer transactions can accomplish. Governments necessarily make *normative* decisions—legislative, executive, and judicial—and Bitcoin transactions involve no judgment.

The most ambitious attempts to use cryptocurrencies as more general legal tools reveal the apparent limits of those strategies. Bitcoin includes a

parties would agree to denominate contracts in Bitcoins. Cf. William J. Luther & Lawrence H. White, *Can Bitcoin Become a Major Currency?* 6 (George Mason Univ. Dep’t of Econ., Working Paper No. 14-17, 2014), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2446604 (arguing that Bitcoin may succeed as a medium of exchange, even if not as a unit of account if it becomes easy to trade in and out of Bitcoin so that speculative risk can be decoupled from exchange use).

4. *Mining*, BITCOIN WIKI, <http://en.bitcoin.it/wiki/Mining> (last visited Feb. 22, 2016); see also NARAYANAN ET AL., *supra* note 1, at 131–64; *infra* Section I.C.

5. The conventional wisdom of economists is that fiat currency is more stable than commodity currency. See, e.g., Christopher Shea, *Survey: No Support for Gold Standard Among Top Economists*, WALL STREET J.: IDEAS MKT. BLOG (Jan. 23, 2012, 10:36 AM), <http://blogs.wsj.com/ideas-market/2012/01/23/survey-no-support-for-gold-standard-among-top-economists/>. Skeptics (often called “goldbugs”) argue that commodity currencies’ insulation from political decision-making makes them more stable. Some of these skeptics believe that Bitcoin’s insulation from politics may similarly allow for relative stability in the long term. See, e.g., DETLEV S. SCHLICHTER, *PAPER MONEY COLLAPSE: THE FOLLY OF ELASTIC MONEY* 15–16, 289–300 (2d ed. 2014).

6. TCP/IP is a protocol suite featuring the Transmission Control Protocol and the Internet protocol. *TCP/IP*, PCMAG, <http://www.pcmag.com/encyclopedia/term/52614/tcp-ip> (last visited Feb. 22, 2016).

7. There is no central repository for this list. It is maintained separately by participating nodes. See *Block Chain*, BITCOIN WIKI [hereinafter *Block Chain*], http://en.bitcoin.it/wiki/Block_chain (last visited Feb. 22, 2016). This is what makes Bitcoin peer-to-peer.

rudimentary scripting language⁸ that, in principle, permits a contract to be resolved by a third-party “oracle.” For example, a Bitcoin wiki suggests that a Bitcoin contract could allow money to be transferred to a third party only once the oracle gives confirmation that a named individual has died.⁹ In effect, the oracle serves as an escrow agent.¹⁰ Bitcoin is a weak substitute for conventional life insurance, however, as insurance involves much more than escrow.¹¹ The Ethereum Project could provide a closer substitute. It aims to create a cryptocurrency allowing Turing-complete computations—i.e., classical computer programs of arbitrary complexity.¹² So, it might be possible to aggregate insurance premiums into a fund and make payouts when specified conditions are met. A recent article explains how smart contracts might lead to a “Lex Cryptographia,”¹³ where smart contracts enable the creation of decentralized autonomous organizations. But until computer programs can exhibit general artificial intelligence, they will lack judgment. They will not, for example, be able to determine whether vague contract provisions have been satisfied. Cryptocurrencies cannot solve the problem of incomplete contracts,¹⁴ and as long as contracts are incomplete, humans will need to resolve ambiguities.

8. See *Script*, BITCOIN WIKI [hereinafter *Script*], <http://en.bitcoin.it/wiki/Script> (last visited Feb. 22, 2016); see also NARAYANAN ET AL., *supra* note 1, at 79–84.

9. *Contract: Example 4: Using External State*, BITCOIN WIKI, http://en.bitcoin.it/wiki/Contracts#Example_4:_Using_external_state (last visited Feb. 22, 2016). If the oracle determines that a condition is met, it produces a digital signature needed to complete the transaction. *Id.*

10. The wiki also gives a separate example of an escrow transaction in which a client’s funds are placed in escrow under terms such that the money can be sent to the merchant if both the client and merchant agree (completing the purchase), to the client if both agree (refunding the amount, perhaps because of a problem with delivery), or to the merchant if the merchant and the mediator agree. *Contract: Example 2: Escrow and Dispute Mediation*, BITCOIN WIKI [hereinafter *Contract: Example 2: Escrow and Dispute Mediation*], http://en.bitcoin.it/wiki/Contracts#Example_2:_Escrow_and_dispute_mediation (last visited Feb. 22, 2016).

11. See generally KENNETH S. ABRAHAM, *DISTRIBUTING RISK: INSURANCE, LEGAL THEORY, AND PUBLIC POLICY* (1986).

12. *White Paper: A Next-Generation Smart Contract and Decentralized Application Platform*, GITHUB [hereinafter *White Paper*], <http://www.github.com/ethereum/wiki/wiki/White-Paper> (last visited Feb. 22, 2016).

13. Aaron Wright & Primavera De Filippi, *Decentralized Blockchain Technology and the Rise of Lex Cryptographia* (Mar. 10, 2015) (unpublished manuscript), <http://www.ssrn.com/abstract=2580664> (defining the term “Lex Cryptographia” as a subset of law in which rules are “administered through self-executing smart contracts and decentralized (autonomous) organizations”).

14. Contracts are incomplete in part because some contingencies are not anticipated, and also because parties leave them deliberately incomplete, either because the contracts are self-enforcing or because people believe that norms of fairness will help resolve disputes. See Robert E. Scott, *A Theory of Self-Enforcing Indefinite Agreements*, 103 COLUM. L. REV. 1641 (2003). A cryptocurrency that can exercise normative judgment can be seen as a mechanism that makes a contract self-enforcing or as a mechanism that avoids judicial enforcement through peer-to-peer decision-making.

This Article, however, shows that cryptocurrency protocols can be used to aggregate human judgment and make legal decisions. Just as a cryptocurrency need not identify a central banker who maintains transaction records, a cryptocurrency also need not identify specific people responsible for making a decision. In other words, cryptocurrencies can crowdsource decision-making.¹⁵ Crude mechanisms for crowdsourcing decision-making already exist—consider polls conducted on the Internet. A cryptocurrency relying on counting votes would be similarly unreliable, as the principle of “one person, one vote”¹⁶ cannot be implemented with any currency that allows anonymous ownership. An alternative would be voting based on Bitcoin interest.¹⁷ But such a system, or one allowing vote buying,¹⁸ would give greater influence to those with more Bitcoins. A better approach is to design a system in which the cryptocurrency protocol implements what game theorists call a “tacit coordination game.”¹⁹

In Thomas Schelling’s famous tacit coordination game experiment, a subject must attempt to meet another subject in New York City the following day without advance coordination of time and place.²⁰ Schelling’s survey indicated that most subjects would meet at the Grand Central Terminal information booth at noon.²¹ A similar tacit coordination game could give each participant the goal of answering a question in the same way as later participants will answer the same question. Participants would seek focal point solutions, much like the prospective rendezvousers in New York. The answer to the question posed is the most logical focal point. For example, imagine asking someone on the street whether it is “cold” or “hot” outside, and informing her that she will receive \$10 if the next person to whom you ask the same question (with the same deal) answers in the same way. Reporting her true evaluation of the weather—or, better yet, what she expects would be the average person’s evaluation of the average person’s evaluation of the weather—is a better strategy than answering at random.

This Article connects two literatures on tacit coordination games. First, a prior article recognizes the possibility of using tacit coordination games to address normative questions.²² But that article imagines that some central authority has

15. See generally JEFF HOWE, CROWDSOURCING: WHY THE POWER OF THE CROWD IS DRIVING THE FUTURE OF BUSINESS (2009) (providing other examples of crowdsourcing by businesses).

16. See *Gray v. Sanders*, 372 U.S. 368, 381 (1963) (“The conception of political equality from the Declaration of Independence, to Lincoln’s Gettysburg Address, to the Fifteenth, Seventeenth, and Nineteenth Amendments can mean only one thing—one person, one vote.”).

17. See *infra* note 151 and accompanying text (discussing the use of such voting by the cryptocurrency Nxt).

18. See *infra* notes 156–61 and accompanying text.

19. See, e.g., John B. Van Huyck et al., *Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure*, 80 AM. ECON. REV. 234 (1990).

20. THOMAS C. SCHELLING, *THE STRATEGY OF CONFLICT* 54–56 (1980). Schelling coined the phrase “tacit coordination game.” *Id.* at 54.

21. *Id.* at 55.

22. See Michael Abramowicz, *Cyberadjudication*, 86 IOWA L. REV. 533 (2001).

organized the tacit coordination game, performing tasks such as compensating the winners.²³ This poses a significant barrier to using a tacit coordination game for legal purposes, even in enforcing a voluntary contract. For example, contracting parties might agree to tacit coordination game dispute resolution, but courts might refuse to enforce such a contract because the game is not a recognized means of conducting arbitration,²⁴ and seems to be similar to gambling.²⁵ Second, two projects on the Internet have suggested using a cryptocurrency to build tacit coordination mechanisms.²⁶ But these projects have assumed that these mechanisms would work well only for questions with objective answers,²⁷ and the mechanisms they imagine require many players to announce the correct answer simultaneously.

This Article provides a different protocol for implementing tacit coordination, embracing subjective questions and allowing parties to indicate their opinions seriatim, updating the collective determination with each new opinion announcement. The protocol establishes the gains and losses of players, functioning as judges of the questions before them. The result of the game could determine the ownership of cryptocurrency, and it would not require judicial enforcement. A government wishing to interfere would need to resort to regulating

23. One commentator has proposed using a tacit coordination game for a particular purpose in Bitcoin, but this proposal cannot be extended to more general normative questions. See Ferdinando M. Ametrano, *Hayek Money: The Cryptocurrency Price Stability Solution* (Aug. 19, 2014) (unpublished manuscript), <http://www.ssrn.com/abstract=2425270> (proposing a mechanism using tacit coordination to achieve cryptocurrency price stability).

24. The Federal Arbitration Act generally requires agreements for mandatory arbitration to be enforced. See Federal Arbitration Act, 9 U.S.C. §§ 1–16 (2012); Elizabeth G. Thornburg, *Going Private: Technology, Due Process, and Internet Dispute Resolution*, 34 U.C. DAVIS L. REV. 151, 182 (2000). But an arbitration provision may not be enforced when “a waiver of judicial remedies inherently conflicts with the underlying purposes of that other statute.” *Rodriguez de Quijas v. Shearson/Am. Express, Inc.*, 490 U.S. 477, 483 (1989). Some courts have thus refused to enforce arbitration agreements where the agreement seemed unduly one-sided. See, e.g., *Hooters of Am., Inc. v. Phillips*, 173 F.3d 933 (4th Cir. 1999). A peer-to-peer arbitration provision might be voided if the courts are uncomfortable with it.

25. See Abramowicz, *supra* note 22, at 541–56. Such a game might be considered to be a “game of skill” and thus exempt from regulation. Cf. Steven D. Levitt et al., *Is Texas Hold ‘Em a Game of Chance? A Legal and Economic Analysis*, 101 GEO. L.J. 581 (2013) (criticizing the courts’ approach to distinguishing games of chance and skill).

26. See Vitalik Buterin, *SchellingCoin: A Minimal-Trust Universal Data Feed*, ETHEREUM: TECHNICAL BLOG (Mar. 28, 2014), <http://blog.ethereum.org/2014/03/28/schellingcoin-a-minimal-trust-universal-data-feed/>; Jack Peterson & Joseph Krug, *Augur: A Decentralized, Open-Source Platform for Prediction Markets*, AUGUR.NET (2016), <http://www.augur.link/augur.pdf>.

27. Buterin’s sample data feeds include things like “the temperature in Berlin” or a currency exchange rate. Buterin, *supra* note 26. Peterson and Krug state that “truth-by-consensus works well only for outcomes which are easily and objectively determinable.” Peterson & Krug, *supra* note 26, at 9.

payments into or out of the cryptocurrency²⁸ or regulating contracting parties directly.²⁹

This Article's ambition is to describe the possibility of peer-to-peer law, not to argue that it is desirable. Traditional legal institutions have obvious advantages. Representative government is valuable both because political deliberation can improve decisions³⁰ and because democratic participation enhances legitimacy.³¹ A full analysis of the strengths and weaknesses of existing institutions is well beyond this Article's scope, but peer-to-peer law is most plausible where existing decision-making mechanisms are most flawed, for example where corruption is endemic. Peer-to-peer decision-makers would have incentives to combat self-interested decision-making. Similarly, peer-to-peer law could be helpful when agency costs are especially high, which may be the case with some corporate decision-making, or when decision-makers are relatively uninformed, or when bureaucracy or litigation impose unnecessary transaction costs on relatively simple decisions. Peer-to-peer decision-making could emerge in niche legal contexts, which could provide data and experience about the relative advantages and disadvantages of such decision-making in comparison to more conventional decision-making.

Peer-to-peer law is most plausible as a mechanism of voluntary private ordering. The strongest defense against the argument that Bitcoin is inherently worthless³² is that demand exists (or in the future may exist³³) for peer-to-peer

28. The government attacks Internet gambling in much the same way. *See* Unlawful Internet Gambling Enforcement Act of 2006, 31 U.S.C. § 5363 (2012).

29. The government has at times hesitated to enforce rules against voluntary contracting against the contracting parties. In the internet gambling context, for example, enforcement has been focused on the operators of gaming companies, not individual gamblers, though this has partly been because of ambiguity as to whether individuals commit illegal acts by gambling online. *See* Jason A. Miller, *Don't Bet on This Legislation: The Unlawful Internet Gambling Enforcement Act Places a Bigger Burden on Financial Institutions than Internet Gambling*, 12 N.C. BANKING INST. 185, 211–12 (2008).

30. *See, e.g.*, Joshua Cohen, *Deliberation and Democratic Legitimacy*, in *THE GOOD POLITY: NORMATIVE ANALYSIS OF THE STATE* 17–32 (Alan Hamlin & Philip Pettit eds., 1989) (discussing how ideal deliberation can enhance legitimacy). *But see* Cass R. Sunstein, *Deliberative Trouble? Why Groups Go to Extremes*, 110 YALE L.J. 71 (2000) (documenting that deliberation can sometimes promote polarization).

31. *See, e.g.*, James Weinstein, *Participatory Democracy as the Central Value of American Free Speech Doctrine*, 97 VA. L. REV. 491, 505–06 (2011) (discussing the significance of the norm of political participation).

32. *See* Alex Crippen, *Buffett Blasts Bitcoin as 'Mirage': 'Stay away!'*, CNBC, (Mar. 14, 2014, 10:29 AM), <http://www.cnbc.com/id/101494937> (reporting Warren Buffett's skepticism).

33. Anticipated future value is what makes Bitcoin valuable today. *See, e.g.*, Timothy B. Lee, *Why I'm Investing in Bitcoins (Updated)*, VOX (Sept. 5, 2014, 8:30 AM), <http://www.vox.com/2014/9/5/6086171/why-im-investing-in-bitcoins> (estimating Bitcoin's value based on the anticipated number of transactions in the future). The uncertainty about future value, however, contributes to Bitcoin's volatility. *See, e.g.*, Jon Southurst, *Bitcoin Price Continues to Fall, Breaks \$200 Mark*, COINDESK (Jan. 14, 2015, 9:58 AM),

transactions. Each element of this defense also suggests demand for peer-to-peer decision-making. First, government regulation imposes transaction costs, and a cryptocurrency could evade such regulation.³⁴ Similarly, peer-to-peer decision-making might reduce litigation costs. Second, some people have an *ideological* preference—based on some form of libertarianism, anti-corporatism, or anarchism—for using Bitcoin.³⁵ If it were just as easy to enter into transactions using Bitcoin or MasterCard,³⁶ then this preference could be cheaply indulged, as could a preference for nongovernmental decision-making. Third, a cryptocurrency could provide privacy protections, which both law-abiding citizens and criminals may value.³⁷ Currently, Bitcoin transactions can sometimes be traced,³⁸ though proposed changes to Bitcoin³⁹ and alternative cryptocurrencies⁴⁰ provide much stronger privacy protection. Peer-to-peer decision-making could provide anonymity to decision-makers and sometimes to litigants.

The typical defense of a decision's legitimacy identifies the decision as an output of some recognized governmental or even private body, but some may perceive legitimacy to derive from an absence of individual control. As long as the output is recognizable, this can be seen as consistent with legal positivism.⁴¹ Even

<http://www.coindesk.com/bitcoin-price-continues-fall-breaks-200-mark/> (noting that Bitcoin had lost over 80% of its value in a year).

34. Cryptocurrency proponents argue that Bitcoin might someday serve as an effective system of micropayments. See Campbell R. Harvey, *Cryptofinance* (Jan. 14, 2016) (unpublished manuscript), <http://www.ssrn.com/abstract=2438299>; see also Daniel Cawrey, *The Promise of Bitcoin Micropayments: Corporations, Incentives and Altcoins*, COINDESK (Feb. 11, 2014, 4:55 PM), <http://www.coindesk.com/promise-bitcoin-micropayments-corporations-incentives-altcoins/>.

35. See Alan Feuer, *The Bitcoin Ideology*, N.Y. TIMES, Dec. 15, 2013, at SR12.

36. See Rick Falkvinge, *Bitcoin's Four Hurdles: Part One—Usability*, FALKVINGE & CO. ON INFO. POL'Y (June 4, 2011), <http://www.falkvinge.net/2011/06/04/bitcoins-four-hurdles-part-one-usability>.

37. See Daniel J. Solove, *"I've Got Nothing to Hide" and Other Misunderstandings of Privacy*, 44 SAN DIEGO L. REV. 745 (2007).

38. See Tom Simonite, *Mapping the Bitcoin Economy Could Reveal Users' Identities*, MIT TECH. REV. (Sept. 5, 2013), <http://www.technologyreview.com/news/518816/mapping-the-bitcoin-economy-could-reveal-users-identities/>.

39. See, e.g., Tim Ruffing et al., *CoinShuffle: Practical Decentralized Coin Mixing for Bitcoin*, MMCI, SAARLAND UNIV. (Aug. 14, 2014), <http://crypsys.mmci.uni-saarland.de/projects/CoinShuffle/coinshuffle.pdf>.

40. See, e.g., *Darkcoin*, DARKCOIN WIKI, http://wiki.darkcoin.eu/wiki/Main_Page (last visited Nov. 9, 2015) (providing an overview of a cryptocurrency called DarkCoin, which relies on a technology called DarkSend that prevents transactions from appearing in a public block chain).

41. For a useful summary of positivism and its variants, see Brian Leiter, *Positivism, Formalism, Realism*, 99 COLUM. L. REV. 1138, 1140–44 (1999) (reviewing ANTHONY SEBOK, *LEGAL POSITISM IN AMERICAN JURISPRUDENCE* (1998)). Positivist theory indicates that law's content is a social fact. *Id.* at 1141. Ordinarily, the relevant social fact might be whether a particular institution (such as the legislature) has made a particular decision, but the occurrence of a peer-to-peer decision also could be social fact.

social customs can serve as an authoritative source of law, at least under many versions of positivism.⁴² But the possibility of peer-to-peer decision-making challenges the conventional assumption that centralized institutions (such as legislatures and courts) are needed to produce law of sufficient clarity to be workable. Justice Holmes insisted that “[t]he common law is not a brooding omnipresence in the sky, but the articulate voice of some sovereign or quasi-sovereign that can be identified.”⁴³ This Article, however, suggests that law can be produced by nonsovereigns competing to discern the brooding omnipresence of the best answers to normative legal questions.

The dawn of cryptocurrency-based law is not near. There are serious obstacles to its emergence, including the need for experimentation with tacit coordination games to establish that participants will seek to address the normative questions posed.⁴⁴ There is, however, a natural test case for cryptocurrency-based tacit coordination games. They could be used to make (or merely recommend) decisions necessary for effective operation of Bitcoin or another cryptocurrency. The Bitcoin protocol, ironically, is coordinated in the same centralized manner as other open source projects.⁴⁵ A few people decide whether to accept pull requests on the source code.⁴⁶ It is sometimes said that Bitcoin is decentralized because anyone can fork⁴⁷ the Bitcoin code and create a new cryptocurrency.⁴⁸ But this is a bit like saying that colonial governments were not centralized because anyone could move to the wilderness and form their own governments. Open source is not inherently peer-to-peer. A cryptocurrency is a natural testing ground for peer-to-peer decision-making because the existence of centralized decision-making is at odds with the broader goals of the alternative currency movement. The Bitcoin

42. Whether this is in tension with positivism, however, depends on the particular version of positivism. *See, e.g.*, Henry E. Smith, *Custom in American Property Law: A Vanishing Act*, 48 TEX. INT’L L.J. 507, 519 (2013) (noting that “there is no reason in many versions of positivism why custom could not be a source of law,” though “narrow Austinian-style positivism that identifies law with commands of a sovereign does not naturally look at custom as a source of the law”).

43. *S. Pac. Co. v. Jensen*, 244 U.S. 205, 222 (1917) (Holmes, J., dissenting).

44. *See infra* Section II.A.2.

45. For a critique of Bitcoin as insufficiently decentralized, see Arthur Gervais et al., *Is Bitcoin a Decentralized Currency?*, IACR.ORG (2013), <http://eprint.iacr.org/2013/829.pdf>.

46. *See* Alec Liu, *Who’s Building Bitcoin? An Inside Look at Bitcoin’s Open Source Development*, MOTHERBOARD: VICE (May 7, 2013, 6:20 PM), http://www.motherboard.vice.com/en_au/blog/whos-building-bitcoin-an-inside-look-at-bitcoins-open-source-development (describing the development process and naming the developers with “the power to merge code into the official repository”).

47. For a description of forking and an argument that the possibility of forking constrains those supervising open-source projects to take into account community views, see Linus Nyman & Juho Lindman, *Code Forking, Governance, and Sustainability in Open Source Software*, TECH. INNOVATION MGMT. REV., Jan. 2013, at 7.

48. *See How to Fork Bitcoin and Build Own Cryptocurrency*, STACKEXCHANGE <http://www.bitcoin.stackexchange.com/questions/19287/how-to-fork-bitcoin-and-build-own-cryptocurrency> (last visited Feb. 23, 2016).

community has recognized the need for governance, but the only proposals so far have been for centralized governance.⁴⁹

Peer-to-peer decision-making could be used to determine whether to make changes to the Bitcoin reference code. This modest application of peer-to-peer law would allow the institution of Bitcoin to respond to the challenges it faces. An existential risk to Bitcoin is that some other cryptocurrency will emerge as dominant, and including peer-to-peer decision-making could bolster either Bitcoin or a competitor. Peer-to-peer decision-making also could be useful in conducting other currency-related activities. For example, it could be used to determine whether blocks of transactions should be added to the block chain. This is the central task performed by Bitcoin miners, and development of a reliable alternative system could save resources. These resources, in turn, might be used to reward activities that promote the currency, such as providing liquidity to stabilize the currency, lobbying, developing source code, or suggesting useful improvements to the cryptocurrency. Peer-to-peer decision-making could be used to decide whether to reward those engaging in such activities with newly minted currency.

Part I of this Article will introduce the concept of peer-to-peer governance by identifying its three critical components: (1) a decentralized ledger; (2) a decentralized decision; and (3) a decentralized fisc. Bitcoin has each of these, but its capacity to make decentralized decisions is limited, and its fiscal power is restricted to supporting mining activity. Part II explains how formal tacit coordination games could be played using transactions on the Bitcoin block chain and how the results of such games could transform Bitcoin into a genuine peer-to-peer institution, with a much more flexible decision-making apparatus. Finally, Part III examines the potential role for peer-to-peer decision-making in the legal system, focusing on private law (including voluntary arbitration and trusts), but also considering the possibility of public law institutions built on Bitcoin, most plausibly a central bank.

I. THE THREE CORNERSTONES OF PEER-TO-PEER GOVERNANCE

The *Oxford English Dictionary* defines “peer-to-peer” as “designating or relating to a network in which each computer can act as a server for the others, allowing shared access to files and other resources.”⁵⁰ The most familiar context, technological and legal, is peer-to-peer filesharing,⁵¹ where the absence of a central server eliminates the need for intermediaries to store files being shared and

49. See Eli Dourado, *How Should Bitcoin Be Governed*, PLAIN TEXT (July 30, 2015), <http://www.readplaintext.com/how-should-bitcoin-be-governed-680192fcd92b>.

50. *Peer-to-peer*, OXFORD ENGLISH DICTIONARY, <http://www.oed.com/view/Entry/139725?redirectedFrom=peer-to-peer#eid31476999> (last visited Feb. 23, 2016).

51. See, e.g., *MGM Studios Inc. v. Grokster, Ltd.*, 545 U.S. 913 (2005) (allowing lawsuit against companies that created peer-to-peer software for inducing copyright infringement).

frustrates the ability of the government to stop copyright violations.⁵² Peer-to-peer governance, then, might be defined as a system of decision-making generally regarded as authoritative even though it lacks a centrally designated authority (or authorities) to make and enforce decisions. The United States includes centrally designated authorities (the legislative, the executive, and the judicial branch), so it is not peer-to-peer governance. A state of anarchy, moreover, is also not peer-to-peer governance—while it may lack centrally designated authorities, it does not produce authoritative rules or adjudications.⁵³ Perhaps there are fringe cases of peer-to-peer governance, such as social norms and practices that function like legal institutions,⁵⁴ but Lisa Bernstein's diamond merchants have formalized systems of arbitration,⁵⁵ and Bob Ellickson's ranchers rely substantially on unwritten rules rather than an alternative system of creating legislation.⁵⁶ For rules and adjudications to be authoritative, they generally must be in writing,⁵⁷ and familiar institutions either have centralized processes for lawmaking or function without relying on authoritative written law.

There are thus only limited precedents for peer-to-peer governance before Bitcoin, which engages in peer-to-peer governance of a limited sort. The Bitcoin protocol does produce written decisions—recording transfers of property rights and granting new property rights to Bitcoin miners who successfully solve hash problems—without designating a central authority to produce or store the decisions. But Bitcoin is a rather feeble system of peer-to-peer governance, because Bitcoin cannot produce open-ended rules (whether written in natural or computer language). Bitcoin does require important multidimensional decisions about how the protocol should evolve, and humans make those decisions based on arguments and written discussion,⁵⁸ but these decisions are made by a centralized

52. See generally Tim Wu, *When Code Isn't Law*, 89 VA. L. REV. 679, 726–45 (2003) (describing the development of peer-to-peer filesharing as a mechanism of interest-group behavior designed to minimize legal costs).

53. Anarchism, however, does not necessarily entail the rejection of all authority. See generally PAUL McLAUGHLIN, *ANARCHISM AND AUTHORITY: A PHILOSOPHICAL INTRODUCTION TO CLASSICAL ANARCHISM* 28–29 (2007).

54. Norms can have large effects on behavior, but they are contestable and can change. For an account of changes in social norms, see Richard H. McAdams, *The Origin, Development, and Regulation of Norms*, 96 MICH. L. REV. 338, 391–400 (1997).

55. Lisa Bernstein, *Opting Out of the Legal System: Extralegal Contractual Relations in the Diamond Industry*, 21 J. LEGAL STUD. 115, 124–30 (1992) (describing a highly developed system of arbitration that serves as an alternative to legal enforcement).

56. See Robert C. Ellickson, *Of Coase and Cattle: Dispute Resolution Among Neighbors in Shasta County*, 38 STAN. L. REV. 623 (1986). The ranchers do recognize the existence of central authorities, and will occasionally complain to their elected bounty supervisors. See *id.* at 680.

57. Jed Rubenfeld argues that writtenness is central to the case for judicial supremacy. See JED RUBENFELD, *FREEDOM AND TIME: A THEORY OF CONSTITUTIONAL SELF-GOVERNMENT* 163–68 (2001).

58. See, e.g., *BIP Proposal: Eliminate No-Fee Transactions in Bitcoin*, BITCOIN F. (Mar. 7, 2013), <http://www.bitcointalk.org/index.php?topic=150194.0> (providing discussion of a Bitcoin Improvement Proposal).

group with the power to modify a particular version of the Bitcoin software code repository generally regarded as authoritative.⁵⁹ The Bitcoin source code was created as an open-source project, but any particular fork of an open-source project has a central repository and is thus not itself peer-to-peer. Bitcoin, in short, uses peer-to-peer governance to approve transactions but not to approve changes to Bitcoin.

This Part describes the essential components of a robust system of peer-to-peer governance, capable of generating rules in natural or computer language and of providing incentives to reward those who enforce those rules or otherwise advance the interests of the institution. The three essential components are a decentralized ledger for recording decisions, a decentralized means of making decisions, and a decentralized fiscal power. This Part argues that these form a three-legged stool, each bolstered from a robust version of the other two. A decentralized ledger cannot work, and is of little use without decentralized decision-making (at least as to whether a purported ledger is valid) and spending. A decentralized tool for making decisions is but philosophy if those decisions cannot be recorded in an authoritative way outside anyone's control or if there is no means of enforcing those decisions with financial incentives. And the ability to spend money cannot be exercised if there is no means to decide how to spend it or to record such decisions.

A cryptocurrency is just one example of a possible peer-to-peer institution, but it is a critical example because it enables the decentralized fiscal power. This Part will elaborate on the three essential components of peer-to-peer governance by focusing on cryptocurrencies. Bitcoin has a decentralized ledger (though only for transactions), decentralized decision-making (though only for a very particular type of decision), and a decentralized fiscal power (but only to reward a specific type of activity). But Bitcoin's advances point the way to the possibility of a true peer-to-peer governance institution—built on extensions to the Bitcoin protocol, or similar cryptocurrency protocols, and capable of performing tasks more complex than keeping track of currency transactions. The core extension needed is the facility to play tacit coordination games based on normative questions.

It may seem odd to imagine building a cryptocurrency on a tacit coordination game, but, in fact, each of Bitcoin's components already depend on tacit coordination. As a recent economic analysis of Bitcoin notes, "Participants must maintain consensus (1) on the rules to determine validity of transactions, (2) on which transactions have occurred in the system, and (3) that the currency has value."⁶⁰ These three challenges correspond to the three powers to be discussed here. Bitcoin's decentralized ledger requires consensus on the rules for whether a transaction is valid; Bitcoin's decentralized decision involves tacit coordination

59. See *Bitcoin*, GITHUB, <http://www.github.com/bitcoin/bitcoin> (last visited Feb. 23, 2016) (providing official Bitcoin source code).

60. Joshua A. Kroll et al., *The Economics of Bitcoin Mining, or Bitcoin in the Presence of Adversaries*, WEIS 2 (2013), <http://www.econinfosec.org/archive/weis2013/papers/KrollDaveyFeltenWEIS2013.pdf>.

about which collection of individually valid transactions is the complete and authoritative one; and Bitcoin can exercise its decentralized fiscal power to reward miners only because there is a consensus that cryptocurrencies in general, and Bitcoin in particular, are valuable. This Part describes these forms of tacit coordination, illustrating how Bitcoin works from a new perspective.

A. *The Decentralized Ledger*

Some have argued that the writtenness of the Constitution mandates particular forms of interpretation or judicial review.⁶¹ Even critics of that claim acknowledge that writing serves important legal functions.⁶² Among those, “text serves as a focal point for legal coordination.”⁶³ Absent land records, for example, it would be difficult to coordinate concerning use of land. Similarly, although statutes and judicial decisions could exist without recording, the written record reduces the risk that legislative intent or judicial doctrine is distorted. With land, expensive title searches are needed to verify the extent of rights,⁶⁴ and the risk of fraud requires insurance.⁶⁵ While it is possible to ascertain the literal content of most legal decisions cheaply and with near certainty, disputes arise over implications and meaning.

The central technological advance of Bitcoin is the invention of the block chain, which tracks decisions of Bitcoin owners. In a typical decision, one user (identified by a Bitcoin address) transfers Bitcoins to another user (similarly identified),⁶⁶ though more complicated transactions are possible.⁶⁷ The block chain includes only transactions that are verified as legitimate.⁶⁸ After the most recent block-chain update, one can ascertain a prospective transaction’s legitimacy by examining the block chain to determine the number of Bitcoins associated with the originating address. Though designed specifically for currency transactions, the block chain is a tool of general applicability, and at least one alternative currency is designed with the purpose of enabling its block chain to store metadata in its

61. See, e.g., Randy E. Barnett, *An Originalism for Nonoriginalists*, 45 *LOY. L. REV.* 611, 635 (1999) (“[A] proper respect for the writtenness of the text means that those committed to this Constitution have no choice but to respect the original meaning of its text until it is formally amended in writing.”).

62. Andrew B. Coan, *The Irrelevance of Writtenness in Constitutional Interpretation*, 158 *U. PA. L. REV.* 1025, 1047–70 (2010).

63. *Id.* at 1048.

64. Ordinarily, it is optimal not to search too far in the past, even though this means that title will not be established with certainty. See Matthew Baker et al., *Optimal Title Search*, 31 *J. LEGAL STUD.* 139, 139–40 (2002).

65. See, e.g., David E. Woolley & Lisa D. Herzog, *MERS: The Unreported Effects of Lost Chain of Title on Real Property Owners*, 8 *HASTINGS BUS. L.J.* 365, 398–99 (2012) (noting that title insurance covers fraud concerning the chain of title).

66. A Bitcoin address is generally a randomly generated string of characters. See *Address*, BITCOIN WIKI [hereinafter *Address*], <http://en.bitcoin.it/wiki/Address> (last visited Feb. 23, 2016).

67. See, e.g., *Contract: Example 2: Escrow and Dispute Mediation*, *supra* note 10 and accompanying text.

68. *Block Chain*, *supra* note 7.

transactions.⁶⁹ Thus, in principle, a block chain can be an authoritative, chronologically ordered record of any type of legal decision.

Any database can store records chronologically. What differentiates the block chain is that it is a database without a central repository.⁷⁰ Any number of copies of the block chain may exist, and the Bitcoin protocol is designed to ensure that they are in sync, or more precisely that they are *eventually consistent*⁷¹—i.e., that temporary deviations are resolved over time. This should work even if noncooperative individuals seek to falsify the block chain to their own advantage—e.g., to allow them to spend their Bitcoins twice or more. The block chain may function even if some subset of the servers fail—e.g., because of a natural disaster or governmental interference. Thus, if an authoritative written record of all decisions is a prerequisite for effective governance, the block chain is a mechanism that satisfies this requirement in a peer-to-peer way.

The most significant prior art underlying the block chain is public key cryptography. A mathematical technique can be used to quickly generate two keys of a specific length (say, 256 bits). One of the keys can be used to scramble a communication, and the other key can then unscramble it.⁷² This method can be used to authenticate documents. For example, a “hash function” can create a short code from a document, essentially a fingerprint.⁷³ The authenticator then scrambles (encrypts) this code using the private key. The public key can be used to decrypt it, producing the original hash. Thus, anyone who knows the relevant algorithms and the public key can conclude, with near certainty, that someone who knew the private key corresponding to the public key must have performed the encryption. The only way to determine the private key from the public key is to guess, and this would take eons.⁷⁴ Producing the encrypted code is taken to signify agreement with the content of the document. A Bitcoin address is a public key, and a transaction must be signed with the private key corresponding to that public key.⁷⁵

69. See FLORINCOIN, <http://www.florincoin.org/florincoin.pdf> (last visited Feb. 23, 2016).

70. *Block Chain*, *supra* note 7.

71. See generally Werner Vogels, *Eventually Consistent*, 52 COMMS. ACM 40 (2009), <http://dl.acm.org/citation.cfm?doid=1435417.1435432> (explaining eventual consistency in database design).

72. See generally CHRISTOF PAAR & JAN PELZL, UNDERSTANDING CRYPTOGRAPHY: A TEXTBOOK FOR STUDENTS AND PRACTITIONERS 149–73 (1st ed. 2010) (providing an overview of symmetric key cryptography).

73. See, e.g., Ralph C. Losey, *HASH: The New Bates Stamp*, 12 J. TECH. L. & POL'Y 1, 12–16 (2007) (providing an accessible introduction to hashing).

74. PAAR & PELZL, *supra* note 72, at 156–57 (discussing the relationship of key lengths to the difficulty of guessing keys). With a 256-bit key, there are 2^{256} combinations, more than a 1 followed by 77 zeros.

75. *Address*, *supra* note 66.

A computer can verify a digital signature quickly,⁷⁶ and this simplifies construction of the block chain. It is trivial to confirm the legitimacy of all transactions that led to particular Bitcoins being owned by the person purporting to transfer them. No one could create a block chain with fake transfers or unauthorized transactions because any collection of transactions can be easily verified. That is, anyone can easily confirm the transaction was approved by a holder of the private key that corresponds to the public key representing the Bitcoin account. The block chain overcomes a different challenge: the danger that an authorized transaction will be *omitted* from the block chain. If one could spend Bitcoins but keep the transaction off the block chain, then one might be able to spend those Bitcoins again.⁷⁷

Bitcoin addresses the problem in part by adding transactions in ordered groups, called blocks.⁷⁸ Bitcoin provides incentives, to be discussed in the next Section, for miners to create these blocks periodically, but also makes creating blocks difficult to do—sufficiently difficult that a new block will be created on average only every ten minutes.⁷⁹ The blocks are linked by hashes. When a block is added, a hash function produces a hash based on fields including the previous block's hash and the transactions on the new block.⁸⁰ Thus, it is not possible to omit a block or a transaction on a block without changing the hash on all subsequent blocks. Accordingly, if one knows of a legitimate transaction and the hash of its block,⁸¹ one can verify the legitimacy of all transactions reported on the block chain up to that block. It is not possible with a reasonable amount of computer time to create a series of fake transactions that will result in a hash that exactly matches the real hash.⁸² If the blocks in the block chain are not properly linked, then the client software will recognize the block chain as fake. Each client will reject such a block chain because those programming and hosting client software know that others will reject such a block chain as fake.

76. For example, one public key signature system can verify 71,000 signatures per second on an ordinary quad-core processor. See Daniel J. Bernstein et al., *Ed25519: High-Speed High-Security Signatures*, 2 J. CRYPTOGRAPHIC ENGINEERING 77, 77 (2011).

77. A recipient of payments might defend against this by waiting to perform its side of the contract (such as transferring goods) until the new payment is confirmed on the block chain, preferably some blocks before the most recent one. See *infra* note 83 and accompanying text (noting the possibility that the synchronization process might lead to the removal of some blocks from the block chain).

78. *Block*, BITCOIN WIKI, <http://en.bitcoin.it/wiki/Block> (last visited Feb. 23, 2016).

79. This time period was apparently selected arbitrarily in the original Bitcoin paper. See Satoshi Nakamoto, *Bitcoin: A Peer-to-Peer Electronic Cash System*, BITCOIN 4 (2009), <http://www.bitcoin.org/bitcoin.pdf>.

80. *Block Chain Hashing Algorithm*, BITCOIN WIKI, http://en.bitcoin.it/wiki/Block_hashing_algorithm (last visited Feb. 23, 2016).

81. Client software tracks some previous transactions, designated as “checkpoints.” See *infra* notes 90–94 and accompanying text.

82. See, e.g., Losey, *supra* note 73, at 17–18 (discussing the irreversibility of hash functions).

This provides a preliminary illustration of tacit coordination's centrality to Bitcoin. Anyone could fork the block chain and make a client with some other set of rules, for example deleting some previously accepted transactions because the Bitcoins were reported stolen.⁸³ If this fork were widely accepted, then the Bitcoin protocol could change to include this ad hoc list of exceptions to the general principles. This seems highly unlikely, however, because even if there is a strong normative argument for this change, there is also a strong argument, grounded in the need to protect Bitcoin's stability, against ad hoc exceptions. When, however, the authoritative Bitcoin software repository changes the rules of Bitcoin, as it often does,⁸⁴ the changes are widely accepted—at least they have been so far. Nevertheless, revolution always remains possible, and for sufficiently strong reasons, the tacit coordination game could someday lead to many or all Bitcoin users accepting a set of rules, even rules different from those followed by the software in the generally recognized official repository. Indeed, two of Bitcoin's core developers recently forked the software (though they mitigated the danger by providing that the new software would come into effect only if adopted by a supermajority of miners).⁸⁵

B. The Decentralized Decision

The mechanism described so far does not address the risk that two miners will simultaneously add different blocks, some containing one transaction and some containing another. This will not happen often because of design decisions (to be covered in the next Section) that make it difficult to create a valid block, but it is always possible. Bitcoin needs a system for determining which block chain is *authoritative*. It resolves this with a coordination rule. The valid block chain is considered to be the block chain that required the most work to form,⁸⁶ which will generally be the *longest* block chain.⁸⁷ Once a miner adds another block to one of the two competing block chains, one chain becomes longer, and anyone aware of this block chain's existence will ignore the other. Of course, blocks could be added

83. For example, the Bitcoin protocol could have been changed to nullify a large theft. Cf. Timothy B. Lee, *Hackers Allegedly Stole \$400 Million in Bitcoins. Here's How to Catch Them*, WASH. POST: THE SWITCH (Feb. 28, 2014), <http://www.washingtonpost.com/blogs/the-switch/wp/2014/02/28/hackers-allegedly-stole-400-million-in-bitcoins-heres-how-to-catch-them/>.

84. See *Releases*, GITHUB, <http://github.com/bitcoin/bitcoin/releases> (last visited Feb. 23, 2016) (providing a list of Bitcoin software releases).

85. See Maria Bustillos, *Inside the Fight Over Bitcoin's Future*, NEW YORKER (Aug. 25, 2015), <http://www.newyorker.com/business/currency/inside-the-fight-over-bitcoins-future>.

86. See *How Does a Client Decide Which Is the Longest Block Chain if There Is a Fork?*, STACKEXCHANGE, <http://bitcoin.stackexchange.com/questions/936/how-does-a-client-decide-which-is-the-longest-block-chain-if-there-is-a-fork> (last visited Feb. 23, 2016) (explaining that block chains are compared based on their score target).

87. The longest block chain might not be the authoritative one if someone sought to create a new block chain from scratch. One could easily falsify a block chain containing a large number of blocks, but it would be clear that the level of difficulty of adding a block to this block chain was low, and it would be rejected in favor of the authentic block chain.

simultaneously again, but the coordination rule and the difficulty of adding blocks according to a set schedule ensures that eventually, one block chain will emerge as the consensus longest one. The coordination rule is akin to the intuition that the time to meet someone in New York without communication is noon. One can imagine an infinite number of functions that would determine which is the better block chain, but the longest-block-chain rule stands out as particularly salient.

Bitcoin's approach to this issue means that decision-making is not instantaneous. It could not be. Public key cryptography can provide near-instant verification that the owner of Bitcoins (or, more precisely, a holder of the private key associated with the public key address) has authorized their transfer or a particular script that ultimately may lead to their transfer.⁸⁸ However, the possibility that someone might make two such transfers simultaneously means that instant confirmation is impossible. A merchant who wishes to confirm that a transaction with Bitcoin is valid must not only wait for the transaction to be added to the block chain, but indeed wait long enough to ensure that this block chain remains the authoritative one. In theory, even after several blocks have been added, it is possible that some longer block chain could emerge, but empirically, this is highly unlikely,⁸⁹ especially if no competing block chain has yet emerged.

The fear that an inconsistent block chain might emerge at some later time, along with a desire to facilitate quick rejection of long block chains that required only small amounts of effort to create,⁹⁰ has led to the addition of another mechanism for identifying a valid block chain: checkpointing.⁹¹ A checkpoint is a record of the block-chain hash as of a certain point in time, and the Bitcoin reference software itself records checkpoints that must be included in the block chain for it to be valid.⁹² Thus, if the new software is generally accepted, then it is impossible for any transactions older than the checkpoint to be reversed. Checkpointing represents a deviation from a pure peer-to-peer system, because a checkpoint is the result of a decision by the authoritative Bitcoin software designers that it is wise, all things considered, to add this safety device. That

88. See *Script*, *supra* note 8 and accompanying text (discussing Bitcoin scripts).

89. The largest number of blocks that were added to a version of the block chain before being orphaned as a result of a longer chain emerging is four blocks, as of this writing. See *What Is the Longest Blockchain Fork that Has Been Orphaned to Date*, STACKEXCHANGE, <http://www.bitcoin.stackexchange.com/questions/3343/what-is-the-longest-blockchain-fork-that-has-been-orphaned-to-date> (last visited Feb. 23, 2016). Presumably, however, the vast majority of transactions in the orphaned blocks were still ultimately incorporated in the block chain.

90. Checkpointing is motivated by a need to combat denial-of-service attacks, in which attackers present artificially constructed block chains that are longer than the authentic block chain but required less effort to create. See *What Are Checkpoints in Bitcoin Code?*, BITCOIN F. [hereinafter *What Are Checkpoints in Bitcoin Code?*], <http://www.bitcointalk.org/index.php?topic=194078.35;wap2> (last visited Feb. 23, 2016).

91. For a discussion of checkpointing, including complaints by some that it is inconsistent with peer-to-peer decision-making, see *id.*

92. See, e.g., *Add a New Checkpoint at Block 295,000*, GITHUB, <http://www.github.com/bitcoin/bitcoin/pull/4541/commits> (last visited Feb. 23, 2016) (adding a new checkpoint).

decision may itself result from a type of focal point coordination—general agreement in the community that a checkpoint should be added—but it is still a centralized decision. One commentator has argued that checkpointing is, as a practical matter, essential, rendering a decentralized currency impossible.⁹³

Checkpointing can be seen as a reflection of the limits of the Bitcoin decision-making mechanism. The mechanism can be used to make only one type of decision, and the developers of Bitcoin do not even trust the protocol entirely to make that decision without the help of another mechanism that is the direct product of human judgment. Those human judgments are thus hard-coded into the protocol itself. There is nothing inherently wrong with adding a small, centralized component to a peer-to-peer protocol, just as there is nothing inherently wrong with running a non-peer-to-peer web service. But it shows that even in Bitcoin, there is a perception that centralized dictates can be useful to ensure continued successful coordination and will be broadly accepted by the relevant community. If Bitcoin had a system for aggregating human judgment, it might still include checkpointing, because checkpointing makes it easier to identify a valid version of the block chain, but the decisions to add checkpoints would be made peer-to-peer, instead of as a result of a centralized software update.

C. *The Decentralized Fisc*

The most celebrated and controversial aspect of Bitcoin protocol is the incentive that Bitcoin uses to ensure that blocks are generated at regular intervals. The incentive is financial. Bitcoin provides a reward for generating a block of transactions to add to the end of the block chain.⁹⁴ The “miner” who generates a block receives some quantity of Bitcoin, though not from any individual.⁹⁵ Rather, mining creates new Bitcoins that the protocol recognizes as valid.⁹⁶ The size of the reward is fixed according to a schedule, with the number of new Bitcoins offered decreasing approximately 50% every four years.⁹⁷ A miner also can receive transaction fees from transferors of Bitcoins, who voluntarily include these fees in their transactions to encourage miners to include the transactions in a block.⁹⁸

93. See Ben Laurie, *Decentralised Currencies Are Probably Impossible: But Let's At Least Make Them Efficient* (July 5, 2011) (unpublished manuscript), <http://www.links.org/files/decentralised-currencies.pdf>.

94. *Mining: Reward*, BITCOIN WIKI, <http://en.bitcoin.it/wiki/Mining#Reward> (last visited Feb. 23, 2016).

95. *Id.*

96. *Controlled Supply*, BITCOIN WIKI [hereinafter *Controlled Supply*], http://en.bitcoin.it/wiki/Controlled_supply (last visited Feb. 9, 2016).

97. See *id.* (illustrating the schedule).

98. *Transaction Fees*, BITCOIN WIKI, http://en.bitcoin.it/wiki/Transaction_fees (last visited Feb. 25, 2016). A small transaction fee is required for very small transfers of Bitcoins; this mechanism is designed to discourage Bitcoin “dust” or “spam” from filling the block chain. See *Bitcoin Transaction Fees Explained*, BITCOIN FEES, <http://www.bitcoinfees.com/> (last visited Mar. 14, 2016). Ordinarily, most clients will not include larger transactions in a block that do not include some transaction fees, though if a block does include such transactions, other clients will consider it to be a legitimate part of the block chain. *Id.*

Thus, the miners are engaged in an activity (adding blocks) that is socially useful to the Bitcoin community, and Bitcoin incentivizes miners to engage in this activity by granting new Bitcoins. Because Bitcoins are valuable, the Bitcoin protocol is able to provide financial incentives in a peer-to-peer way.

This would be straightforward if confirming transactions were an inherently expensive activity. If, for example, it took a great deal of computer power to arrange transactions in a block, confirm their digital signatures, and calculate a new hash value, then the reward for the Bitcoin miners could be explained by the difficulty of their task. In fact, however, this is trivial. The danger is not that too few miners will confirm transactions and add them to the block chain but that *too many* will and that some may intentionally omit transactions. So, the Bitcoin protocol makes it *artificially* difficult to mine blocks. A block can be added to the block chain only if the block's hash results in a number lower than a specified target.⁹⁹ Under the Bitcoin protocol, this target will fluctuate depending on the success of miners so that, on average, a block is added once every ten minutes.¹⁰⁰ If over time more miners enter Bitcoin and computer hardware improves,¹⁰¹ the target will fall.

Bitcoin miners are thus engaging in an activity that is useful to the Bitcoin community, but only an infinitesimal portion of the computing power is used to generate digital signatures. As of this writing, the target is so low that it begins with 16 zeros.¹⁰² A miner hoping to win Bitcoins collects some set of transactions and fills out the fields of the block record, including a field containing the hash value of the previous block and a field containing a *nonce*.¹⁰³ The nonce can be any 32-bit value of a certain size, so the miner's strategy is to try many nonce values, calculating the block hash for each one, hoping to produce a hash less than the target. A miner must decide what transactions to include before attempting a hash, but the miner has an affirmative reason to include all transactions with positive transaction fees, and no reason to exclude transactions.¹⁰⁴ In fact, sometimes miners include transactions without transaction

99. For the current target, in hexadecimal form, see <http://www.blockexplorer.com/api/status?q=getDifficulty> (last visited Mar. 12, 2016).

100. See *Target*, BITCOIN WIKI [hereinafter *Target*], http://en.bitcoin.it/wiki/Target#What_is_the_target_now.3F (last visited Jan. 14, 2016).

101. Bitcoin miners today generally use specialized hardware that can calculate hashes much quicker than general-purpose computers. See, e.g., Tom Simonite, *Custom Chips Could Be the Shovels in a Bitcoin Gold Rush*, MIT TECH. REV. (Dec. 5, 2012), <http://www.technologyreview.com/news/508061/custom-chips-could-be-the-shovels-in-a-bitcoin-gold-rush/>.

102. The value as of this writing is 000000000000000003AAEA2000. *Target*, *supra* note 100. Thus, the probability of any single hash being successful is less than 1 in 16^{16} , or approximately $1.8 * 10^{19}$ —i.e., 18 billion billion.

103. See *Cryptographic Nonce*, WIKIPEDIA, http://en.wikipedia.org/wiki/Cryptographic_nonce (last visited Feb. 14, 2016) (defining a “nonce” as “an arbitrary number used only once in a cryptographic communication”).

104. It might seem that a block with more transactions would slow down hashing, but the hash is of a fixed-size header to eliminate this incentive to drop transactions. See

fees; perhaps because this contributes to the general welfare of Bitcoin from which they benefit.¹⁰⁵

Bitcoin mining may largely be characterized as rent-seeking,¹⁰⁶ and just as the expected investments of ships searching for buried treasure will generally average about the value of the treasure,¹⁰⁷ one should expect the cost of mining to equal the number of Bitcoins that miners receive (in equilibrium). Although rent-seeking can be socially wasteful, it is not inherently so. Patent theorists, for example, recognize that races to invent dissipate rents and the challenge of patent policy is to ensure that the process of rent dissipation produces as much social benefit as possible.¹⁰⁸ Bitcoin's rent dissipation uses large amounts of energy, imposing negative environmental externalities.¹⁰⁹ A partial solution would be for miners to solve problems that require large amounts of memory instead of fast computation.¹¹⁰ One proposed variant creates problems whose solution would contribute to a social need, such as storage of archival information.¹¹¹

Weaknesses: Dropping Transactions, BITCOIN WIKI, http://en.bitcoin.it/wiki/Weaknesses#Dropping_transactions (last visited Feb. 25, 2016).

105. Itai Damti, Comment to *What Is the Incentive for Bitcoin Miners to Add Transactions Without Fees to a Block?*, QUORA (Apr. 15, 2014), <http://www.quora.com/What-is-the-incentive-for-Bitcoin-miners-to-add-transactions-without-fees-to-a-block>.

106. The legal literature typically focuses on rent-seeking through the political process. See, e.g., CASS R. SUNSTEIN, *AFTER THE RIGHTS REVOLUTION: RECONCEIVING THE REGULATORY STATE* 70 (1990). But it is often defined considerably more broadly. See, e.g., GORDON TULLOCK ET AL., *GOVERNMENT FAILURE: A PRIMER IN PUBLIC CHOICE* 43 (2002) (offering the following definition: "The use of resources for the purpose of obtaining rents for people where the rents themselves come from some activity that has some negative social value"); Kevin M. Murphy et al., *Why Is Rent-Seeking so Costly to Growth?*, 83 AM. ECON. REV. PAPERS & PROCEEDINGS 409, 409 (1993) (encompassing within the definition "any redistributive activity that takes up resources").

107. See, e.g., RICHARD A. POSNER, *ECONOMIC ANALYSIS OF LAW* 41 (5th ed. 1998) (explaining how rent-seeking can lead to complete dissipation of rent).

108. See John F. Duffy, *Rethinking the Prospect Theory of Patents*, 71 U. CHI. L. REV. 439, 458–75 (2004) (explaining that patent rents may be dissipated in more or less efficient ways).

109. See, e.g., Hass McCook, *Under the Microscope: Economic and Environmental Costs of Bitcoin Mining*, COINDESK (June 21, 2014, 11:02 AM), <http://www.coindesk.com/microscope-economic-environmental-costs-bitcoin-mining/>.

110. See, e.g., Colin Percival, *Stronger Key Derivation via Sequential Memory-Hard Functions*, TARSNAP, <http://www.tarsnap.com/scrypt/scrypt.pdf> (last visited Feb. 25, 2016). This proposal is the basis for Litecoin. See *What is Litecoin?*, LITECOIN, <http://www.litecoin.org> (last visited Feb. 25, 2016).

111. See Andrew Miller et al., *Permacoin: Repurposing Bitcoin Work for Data Preservation*, 35 IEEE SYMP. ON SECURITY & PRIVACY 475 (2014), <http://www.cs.umd.edu/~amiller/permacoin.pdf> (offering an alternative "scratch-off puzzle" for cryptocurrencies that is memory hard and would serve the socially beneficial function of preserving data).

Similar to Bitcoin's strategy, this strategy rewards "proof of work,"¹¹² differing only in the type of work to be rewarded. Proof of work of some type is essential because it provides a defense against Sybil attacks.¹¹³ Suppose that Bitcoin drastically reduced the number of new Bitcoins issued with each block and eliminated transaction fees. Altruism alone would likely be sufficient for some people to set up servers to verify transactions of nontrivial size.¹¹⁴ But malicious users might then take advantage of this by setting up servers to create block chains in ways that benefit themselves. For example, they might remove some number of previous blocks from the block chain and then generate many new blocks, creating a new longest chain that the nonmalicious Bitcoin servers would recognize as well. This could allow the malicious users to recover Bitcoins they have previously spent.

This type of manipulation is much more difficult with Bitcoin's demanding proof-of-work standard, because a manipulator would need to create blocks faster than everyone else combined. A manipulator could do this with ownership of more than 50% of the computing power dedicated to solving the hashing problem. Specifically, the manipulator could execute what is known as a 51% attack,¹¹⁵ producing more blocks than everyone else combined. For example, the manipulator could remove a block (containing a transaction in which it spent money) and continue hashing until it had produced at least one more block than everyone else combined. It could then add all of these blocks at once to the block chain, including a block sending the funds to a different account on one of the blocks. Legitimate Bitcoin miners would then accept the new block chain as the legitimate block chain. Such a manipulator also might be able to selectively keep new transactions off the block chain, continuing to lengthen the block chain but only with transactions that it selects. A 51% attack, however, would be extraordinarily expensive (recently estimated at over \$1 billion),¹¹⁶ so it would not be worth the luxury of a double-spend transaction. Perhaps the greater risk is that a

112. Proof of work was originally developed as an anti-spam mechanism. See Cynthia Dwork & Moni Naor, *Pricing via Processing or Combatting Junk Mail*, in 740 LECTURE NOTES IN COMPUTER SCIENCE 139 (E.F. Brickell ed., 1993).

113. A Sybil attack is an attack on a peer-to-peer system in which the attacker presents many different identities. See John R. Douceur, *The Sybil Attack*, in 2429 LECTURE NOTES IN COMPUTER SCIENCE 251 (P. Druschel et al. eds., 2002) (discussing how such attacks can be prevented).

114. Many people voluntarily devote computer resources to peer-to-peer projects. See YOCHAI BENKLER, *THE WEALTH OF NETWORKS: HOW SOCIAL PRODUCTION TRANSFORMS MARKETS AND FREEDOM* 81–89 (2006).

115. See, e.g., Daniel Cawrey, *Are 51% Attacks a Real Threat to Bitcoin?*, COINDESK (June 20, 2014, 11:42 AM), <http://www.coindesk.com/51-attacks-real-threat-bitcoin/> (discussing the possibility of 51% attacks). Someone with less than 50% of the computing power has some chance of generating two blocks before everyone else generates one block, but the chance is not as high, and it is even less likely someone could generate three blocks before everyone else generates two.

116. See *Bitcoin*, COINOMETRICS, <http://www.coinometrics.com/bitcoin/brix> (last visited Sept. 12, 2014).

government might do this with the goal of destroying Bitcoin rather than enriching itself, but even this seems far-fetched.¹¹⁷

Proof of work thus provides robust protection of the block chain. Arguably, however, it is not necessary, or at least not to the same degree, and a number of alternative cryptocurrencies either greatly reduce reliance on proof of work or eliminate it altogether. For example, Nxt uses a system that it calls “transparent forging,” in which users take turns “forging” (instead of “mining”)¹¹⁸ new blocks. The order is based on a hash function and is thus quasi-random,¹¹⁹ but each user’s opportunity to hash is proportional to that user’s ownership, so Nxt’s system is based on the principle of “proof of stake.”¹²⁰ The protocol will ignore a block that is mined when it is not one’s turn. Peercoin, meanwhile, does not explicitly use the concept of turns.¹²¹ Any coin owner may attempt to mine a block, but using coins to do so uses up those coins’ “coin age.”¹²² If there are competing block chains, the chain with the greatest “coin age” is the authoritative one.¹²³ In both systems, creating a block requires minimal computing power, and block-creators have incentive to include transactions with minimal transaction fees. Proof-of-stake systems present their own manipulation challenges,¹²⁴ but to mount a 51% attack, one would need to own more than 50% of the total currency value. Someone in that position would have no incentive to double-spend, because any benefit from double-spending would be offset by a decrease in value to the currency as a whole when such manipulation was perceived.

117. See Kroll et al., *supra* note 60, at 13–14 (modeling the possibility of a “Goldfinger” attack by the government).

118. The Nxt Wiki explains that the terminology is because “all possible coins already exist, and accounts earn coins from transaction fees alone.” See *The Nxt Wiki: FAQ*, WIKIA [hereinafter *The Nxt Wiki*], <http://nxtcoin.wikia.com/wiki/FAQ> (last visited Feb. 25, 2016).

119. See *Introduction: What is Nxt?*, NXT WIKI, http://www.wiki.nxtnet.org/wiki/Nxt_Wiki (last visited Feb. 25, 2016).

120. See *The Nxt Wiki*, *supra* note 118 (“Your ability to forge Nxt depends solely on your total account balance as a percentage of all available coins. This is what sets Nxt apart as a pure ‘Proof-of-Stake’ cryptocurrency.”).

121. See, e.g., Sunny King & Scott Nadal, *PPCoin: Peer-to-Peer Cryptocurrency with Proof-of-Stake*, PEERCOIN (Aug. 19, 2012), <http://www.peercoin.net/assets/paper/peercoin-paper.pdf>.

122. *Id.*

123. See *Peercoin Minting*, PEERCOIN, <http://www.peercoin.net/minting> (last visited Feb. 25, 2016).

124. See Iddo Bentov et al., *Cryptocurrencies Without Proof of Work*, TECHNION 4 (July 18, 2014), <http://www.cs.technion.ac.il/~iddo/CoA.pdf>. One problem they identify is that if multiple parties simultaneously create a block, it becomes rational for the next forger to sign both blocks to reduce the danger that the forger will pick the wrong one. *Id.* at 2–3. More troublesome is the possibility that a forger might seek to bribe the party that forges next, or perhaps several such parties, to enable a double-spend transaction. *Id.* at 3. Bentov et al., however, offer a number of solutions to these problems. *Id.* at 2–9. For a thorough discussion of proof-of-stake, see NARAYANAN ET AL., *supra* note 1, at 231–36.

Though there remains the possibility that any cryptocurrency is vulnerable to attacks not yet conceived, the continuing viability of proof-of-stake systems suggests that proof of work is not essential to a cryptocurrency. Bitcoin mining harms Bitcoin holders by diluting their share in Bitcoin. Issuance of new Bitcoins is a form of seignorage revenue for the Bitcoin institution but that revenue is currently spent entirely on mining. The proof-of-stake currencies show that the seignorage revenue could have a neutral effect on currency owners, but it is also possible that a peer-to-peer cryptocurrency could use its decentralized fiscal power for other purposes. If a cryptocurrency had some means of engaging in nonmechanical decision-making about what interests to support, it could assign new coins to individuals who advance those interests. Moreover, a robust decision-making mechanism could allow other peer-to-peer institutions to piggyback on a cryptocurrency, accepting cryptocurrency from private parties and then spending it.

The case against complete reliance on proof of work is a normative argument based on economic efficiency, but existing proof-of-stake systems confront a normative argument based on conceptions of economic equality. The objection is that in a pure proof-of-stake system, all coins are allocated to the initial creators of the system. As one online commentator objects, “This scheme is completely unacceptable because it’s not ‘compatible’ to decentralized nature of cryptocurrencies.”¹²⁵ A counterargument is that the entrepreneurs and programmers who create and promote a currency are providing value.¹²⁶ Arguably, this is better than a system like Bitcoin, which still gives great value to its founders (because they can mine coins when the hashing is easy) and then subsidizes wasteful activity. But some may not find this argument persuasive. Perhaps a system that allocates cryptocurrency over time to those who contribute to the project would satisfy both efficiency and equity concerns.

A full assessment of these normative arguments is beyond the scope of this Article, but their superficial appeal has relevance. The Bitcoin protocol is able to perform the function of a decentralized fisc only because people believe that Bitcoin is valuable—they believe that others will believe that Bitcoin is valuable. This is the highest level tacit coordination game that already exists in Bitcoin, and it can be broken down into separate types of tacit coordination. For Bitcoin to maintain value, people must continue to: (1) believe that cryptocurrencies are valuable; (2) believe that Bitcoin in particular has value; and (3) agree on exactly what the Bitcoin protocol is. The case for (1) is presented above,¹²⁷ but normative arguments may also be relevant to (2). The relative appeal of cryptocurrencies depends on tacit coordination, which may depend in part on saliency, financial features, and normative appeal. Meanwhile, anyone can produce a “hard fork” of

125. Come-from-Beyond, Comment to *What Are Some Counter Arguments for NXT (nxtcoin) “Premine” or Initial Distribution Setup?*, STACKEXCHANGE (July 8, 2014, 7:03 PM), <http://www.bitcoin.stackexchange.com/questions/28366/what-are-some-counter-arguments-for-nxt-nxtcoin-premine-or-initial-distribut/28367#28367>.

126. See *The Nxt Wiki*, *supra* note 118.

127. See *supra* notes 32–39 and accompanying text.

Bitcoin, changing the protocol but accepting the existing block chain,¹²⁸ and normative arguments would then be relevant to the question of which resulting block chain should be viewed as authoritative.

Perhaps the greatest existential threat to Bitcoin is the possibility that there will be a tipping point that leads to some other cryptocurrency dominating it. This could also destabilize cryptocurrency markets more generally, for relative value instability makes the broader project unstable. This presents challenging design questions for Bitcoin developers. Arguably, they should seek to incorporate features of leading alternative currencies, much as the leader in a yachting race should tilt its sails in the same direction as the follower, to prevent even a chance of losing the lead.¹²⁹ But one could also argue that Bitcoin should be conservative, reinforcing the perception of its stability and reducing the risk associated with experimentation. The current structure of Bitcoin decision-making promotes conservative decision-making. The centralized developers incorporate suggested changes only when there is consensus.¹³⁰ This is, in part, because of concerns that lack of consensus would lead to not only a new competitor currency but also—more problematically—a fork of the Bitcoin block chain itself.¹³¹

This may be the smart course. Peer-to-peer decision-making could, however, be useful as a bulwark against a hard fork. The most plausible scenario in which a hard fork could occur is if Bitcoin miners collude to change the rules, presumably to give themselves more Bitcoin—e.g., by increasing transaction fees¹³² or changing the schedule at which new Bitcoins would be created.¹³³

128. See David Kirk, *Cryptocurrency: What is a Fork?*, TECH-RECIPES (Mar. 11, 2014), <http://www.tech-recipes.com/rx/48517/cryptocurrency-what-is-a-fork/>. A hard fork in the rules concerning a valid block occurs only when the new rules would result in acceptance of blocks that the old rules would reject. See *Bitcoin Developer Guide: Consensus Rule Changes*, BITCOIN, <http://www.bitcoin.org/en/developer-guide#consensus-rule-changes> (last visited Feb. 26, 2016). With a soft fork, all new blocks continue to meet the requirements of the old rules, so the old clients will accept new blocks as valid additions to the block chain. *Id.* Any change in the rules governing what constitutes the authoritative block chain will necessarily be a hard fork.

129. See Ian Ayres, *Supply-Side Inefficiencies in Corporate Charter Competition: Lessons from Patents, Yachting and Bluebooks*, 43 U. KAN. L. REV. 541, 550, 553–56 (1995) (using the yachting example to explain why Delaware may have incentives to imitate other states in the race for corporate charter revenue).

130. The Bitcoin Wiki states that when a Bitcoin improvement proposal “is contentious and cannot be agreed upon . . . the conservative option will always be preferred.” *Bitcoin Improvement Proposals*, BITCOIN WIKI, http://wiki.bitcoin.com/w/Bitcoin_Improvement_Proposals (last visited Feb. 26, 2016).

131. See, e.g., Comments to *How Is a Hard Fork Resolved?*, STACKEXCHANGE, <http://www.bitcoin.stackexchange.com/questions/9986/how-is-a-hard-fork-resolved> (last visited Feb. 26, 2016) (discussing the incentives that Bitcoin developers and miners have to come to a consensus resolution in the event of a hard fork).

132. There may be good reasons to increase transaction fees. See, e.g., Kerem Kaskaloglu, *Near Zero Bitcoin Transaction Fees Cannot Last Forever*, in THE INTERNATIONAL CONFERENCE ON DIGITAL SECURITY AND FORENSICS (DIGITALSEC2014) 91

Miners already join together in mining pools, and commentators have noted the possibility that this collusion could facilitate agreements to change the Bitcoin protocol.¹³⁴ This seems especially plausible because the rate at which new Bitcoins are issued is planned to reduce exponentially, and there is no current plan for mandating minimum transaction fees.¹³⁵ Miners have large fixed investments in computers custom-built for mining¹³⁶ and might have an incentive to accept the risks associated with a hard fork, especially if individual enterprises faced bankruptcy as a result of decreased revenue.¹³⁷ In principle, Bitcoin could survive with a hard fork, and Bitcoins that exist on one block chain but not the other would have a value based on the relative perceived legitimacy of the block chains. But this would add complexity, as Bitcoin payment processors would need to offer different prices for block chain *A*, block chain *B*, and block chain *A & B* Bitcoins. Additional forks would complicate things even further.

A peer-to-peer decision-making mechanism could reduce the risk of a successful hard fork in two ways. First, if such a mechanism were established and used to make normative decisions about the evolution of Bitcoin, the resulting decisions might have greater perceived legitimacy. Miners colluding to execute a hard fork of Bitcoin would recognize that the success of their project would depend on the outcome of the tacit coordination game in which people assess the relative authoritativeness of the two forks, and this might depend on normative considerations. One argument made on behalf of a recent attempt at collusion is that centralized decision-making by the core Bitcoin developers is less legitimate than decisions made by the mining community.¹³⁸ A peer-to-peer decision-making

(2014), <http://www.sdiwc.net/digital-library/near-zero-bitcoin-transaction-fees-cannot-last-forever.html>.

133. The Bitcoin Wiki assumes that the supply schedule will never change. *See Controlled Supply*, *supra* note 96 (“The Bitcoin generation algorithm defines, in advance, how currency will be created and at what rate.”). But the algorithm can be changed with sufficient consensus.

134. *See, e.g.*, Ed Felten, *Bitcoin Mining Now Dominated by One Pool*, FREEDOM TO TINKER (June 16, 2014), <http://www.freedom-to-tinker.com/blog/felten/bitcoin-mining-now-dominated-by-one-pool/>. Miners might collude to change the rules even absent existing mining pools. Miners can easily communicate via internet forums, and if it were clear that at some point, some miners would begin running a new version of client software not approved by the core developers of Bitcoin, miners would need to choose sides.

135. Each miner has an incentive to include any transaction that includes a transaction fee sufficient to cover the marginal cost of processing the transaction, because there is virtually no transaction cost associated with a fee. *See* Kroll et al., *supra* note 60, at 12–13. A miner who ignores such a transaction will simply be yielding its voluntary transaction fee to another miner.

136. *See How to Set Up a Bitcoin Miner*, COINDESK, <http://www.coindesk.com/information/how-to-set-up-a-miner/> (last visited Feb. 26, 2016).

137. In general, the prospect of bankruptcy can lead to risky business decisions because it is preferable for the existing owners of the business to have a small chance of survival than allow creditors to take over the business. *See, e.g.*, Barry E. Adler, *Bankruptcy and Risk Allocation*, 77 CORNELL L. REV. 439, 461–63 (1992) (discussing the incentives for equity holders to take risks on the eve of bankruptcy).

138. *See* Bustillos, *supra* note 85.

mechanism could help neutralize that argument by privileging the interests of the Bitcoin community as a whole over those of miners. Second, peer-to-peer decision-making could allow more rapid evolution of Bitcoin. To avoid the perception that Bitcoin is an oligarchy, the centralized developers make changes only when they perceive strong consensus in favor of them. But high supermajority requirements can block useful improvements,¹³⁹ including decisions necessary to either appease miners or protect against their assumption of greater power.

II. PEER-TO-PEER GOVERNANCE FOR CRYPTOCURRENCIES

As Part I showed, Bitcoin offers an ingenious scheme for maintaining a consistent ledger without using a central server. The protocol uses a simple coordination rule to decide which block chain required the most work to create and is, therefore, authoritative. It incentivizes third parties to perform the artificial tasks that make up this work by promising them new Bitcoins and transaction fees. Because simple examination of the block chain makes it possible to determine how much work was performed to create it, this coordination arrangement makes falsification of the block chain virtually impossible. Other cryptocurrencies rely on other simple coordination rules to determine the true block chain. None of the cryptocurrencies require human judgment. Mechanical rules that reduce the chance that disagreement about which block chain is correct will lead to a hard fork of the currency, where some users own Bitcoins valid on one block chain but not the other.

The need for human judgment cannot be avoided when the questions at issue become more complex. Part III will address issues that would arise in using peer-to-peer governance beyond cryptocurrencies, such as determining whether to authorize a payment from an insurance fund. The purpose of this Part is to argue that even for an institution with goals as simple as those of a cryptocurrency—essentially, maintaining a reliable ledger of transactions—incorporating human judgment may strengthen the institution rather than harm it. The success of Bitcoin unavoidably depends on tacit coordination around which version of the protocol should count as authoritative. Creating a formal coordination game with Bitcoin payments could focus the results of the informal tacit coordination game, stabilizing Bitcoin and reducing the possibility that it will be administered for the benefit of particular groups (such as miners) rather than users as a whole. Recent disputes about whether the core developers of Bitcoin have conflicts of interest highlight the need for Bitcoin to adopt peer-to-peer governance.¹⁴⁰

139. Scholars have suggested that supermajority rules may sometimes be useful. See John O. McGinnis & Michael Rappaport, *The Condorcet Case for Supermajority Rules*, 16 SUP. CT. ECON. REV. 67, 68–71 (2008) (describing situations in which supermajority requirements may be efficient). But extreme versions of supermajority rules, such as unanimity requirements, will block changes that even the vast majority of observers believe are beneficial.

140. See Bustillos, *supra* note 85.

Section II.A will begin the task of illustrating how a cryptocurrency could make decisions peer-to-peer with a simple decision currently conducted centrally: approval of a proposed checkpoint. This can be analogized to an administrative adjudication resolving a yes-or-no issue. Individual binary decisions can be aggregated into more complex decisions, including how to improve a text or code. Accordingly, Section II.B will explore how Bitcoin could make peer-to-peer decisions about how to evolve the Bitcoin protocol itself. Because the protocol is expressed as code, this is a general decision-making task analogous to rulemaking.¹⁴¹ Then, Section II.C will show how a cryptocurrency could be designed to award new coins to those who promote the currency, which would require decisions about how much money various parties should receive. This demonstrates how to create a *discretionary*, decentralized fisc. In combination, these capabilities cover the essential building blocks of any decision-making system.

A. Checkpointing

Checkpointing is a useful starting point because it constitutes a centralized element to the peer-to-peer system. Moreover, it can be thought of as a simple binary decision, where someone proposes the addition of a checkpoint, and a decision must be made whether it should be created. A binary decision, such as an “aye” or “nay” vote, is a fundamental building block in any system of decision-making. A checkpoint is a hash of a block that is in the block chain.¹⁴² Software honoring a checkpoint will reject a proposed block chain that does not contain the checkpoint hash without even calculating the total effort used to produce the presented block chain.¹⁴³ In fact, once a Bitcoin mining client accepts a checkpoint it will reject a hypothetical *longer* block chain, reducing the damage that a hypothetical 51% attack could accomplish.¹⁴⁴ Currently, checkpoints are sometimes included when the Bitcoin reference code is updated for other

141. See *Protocol Rules*, BITCOIN WIKI, http://en.bitcoin.it/wiki/Protocol_rules (last visited Feb. 26, 2016) (discussing rules enforced by the protocol code).

142. David Gilson, *Feathercoin Secures Its Block Chain with Advanced Checkpointing*, COINDESK (Aug. 28, 2013, 4:00 PM), <http://www.coindesk.com/feathercoin-secures-block-chain-advanced-check-pointing/> (describing checkpointing and a new cryptocurrency that includes a centralized checkpointing feed).

143. This does not take long, but if this process takes even a second or two, it may facilitate denial-of-service attacks on Bitcoin miners. See *What Are Checkpoints in Bitcoin Code?*, *supra* note 90.

144. See Cawrey, *supra* note 115 and accompanying text. A 51% attacker would not be able to remove blocks past the checkpoint.

reasons.¹⁴⁵ This means that Bitcoin includes few checkpoints,¹⁴⁶ but some other cryptocurrencies include a much larger number of centralized checkpoints.¹⁴⁷

Before describing how a tacit coordination game can produce checkpoints, it may be useful to consider other decentralized options. This will show that peer-to-peer governance is capable of producing normative decisions even if one cannot rely on a system of tacit coordination or prefers not to do so. The main purpose of this Article is to defend the proposition that peer-to-peer governance is possible. The subsidiary purpose is to argue that peer-to-peer governance should be based on tacit coordination games because of the weaknesses in other approaches. Thus, Section II.A.1 will consider peer-to-peer governance through voting, vote buying, and jury-like mechanisms. Section II.A.2 will discuss how an explicit tacit coordination game can determine checkpoints. The purpose of both sections is to highlight how peer-to-peer governance may be used in general. Checkpointing is selected as an example not for its importance, but for its simplicity.

1. Resolution Without Tacit Coordination

The most obvious mechanism for accomplishing peer-to-peer governance is voting. For example, if someone proposes a new checkpoint, we could allow anyone to vote on the new checkpoint and count up all the votes. But the general principle of one-person, one-vote¹⁴⁸ will not work with peer-to-peer governance based on Bitcoin. In theory, a peer-to-peer governance system might maintain a list of people (or just people authorized to vote) and provide some mechanism for them to authenticate themselves.¹⁴⁹ But Bitcoin does not do that, and it is not possible to know whether hundreds of different Bitcoin addresses correspond to the same person. Thus, in the absence of reliance on some external non-peer-to-peer people-tracking mechanism, Bitcoin cannot allow voting based on one-person, one-vote, even if that were desirable.

145. Gilson, *supra* note 142.

146. See *Checkpoints*, GITHUB [hereinafter *Checkpoints*], <http://www.github.com/bitcoin/bitcoin/blob/master/src/checkpoints.cpp> (last visited Sept. 18, 2014) (listing only 13 checkpoints as of September 18, 2014).

147. See King & Nadal, *supra* note 121, at 4 (proposing to include several checkpoints per day in a new cryptocurrency).

148. See *Gray v. Sanders*, 372 U.S. 368, 381 (1963) (discussing importance of this principle in U.S. government). Other forms of peer governance do not involve voting, but emphasize different versions of equality. See *Peer Governance*, P2P FOUND., http://www.p2pfoundation.net/Peer_Governance (last visited Feb. 26, 2016) (listing “equipotentiality” first among peer-to-peer governance’s “main characteristics,” and explaining that “in a peer project all the participants have an equal ability to contribute, although that not all the participants have the same skills and abilities”).

149. Some have argued that a mechanism like the block chain might be used to produce a more reliable mechanism for counting votes in democratic elections, though a critical first step would be to distribute to each authorized voter the ability to vote exactly once. See *VoteCoin*, START JOIN, <http://www.startjoin.com/VoteCoin> (last visited Feb. 26, 2016) (discussing the idea for block chain based voting).

As a result, the most obvious mechanism for implementing voting is to allow voting proportionate with ownership interests. This is the general system for voting in corporate law.¹⁵⁰ At least one cryptocurrency, Nxt, allows for voting on certain types of issues based on ownership interest.¹⁵¹ Presumably, voters will share an interest in a cryptocurrency's success, and voting by interest may work for other types of peer-to-peer institutions as well. But voting in proportion to interest has two problems, both familiar from corporate law. First is the problem of oppression—namely, that those with a majority of interests may make decisions to benefit themselves at the minority's expense.¹⁵² Second is the problem of apathy. Many voters may not take time to fully study the issues, causing uninformed voting.¹⁵³ Further, many voters may choose not to vote, leading to voting that is not broadly representative of the population. These are, of course, familiar problems in both republican government¹⁵⁴ and direct democracy.¹⁵⁵

These problems may not be large in the context of checkpointing, because relatively little is at stake, but there are dangers. Suppose that a sufficiently large coalition of miners seeks to exclude other miners. For example, established miners might try to block new entrants into the mining market. They might accomplish this with checkpoints that validate versions of the block chain with only their own recent blocks, implicitly rejecting the blocks of new entrants. Such a coalition might simply change the rules of Bitcoin and add their own checkpoints, but it may be easier to exploit collusion by adding checkpoints within the protocol rules. If the rules allowed for peer-to-peer decision-making about checkpointing, then even new entrants following these rules would be forced to accept the superiority of the establishment block chain. Although the establishment miners would not own all Bitcoins, they would be a large interest group who would vote. It might be

150. See Lyman Johnson, *Sovereignty over Corporate Stock*, 16 DEL. J. CORP. L. 485, 496–97 (1991) (describing the development of the principle of “one share, one vote” in Delaware).

151. *Voting*, Nxt, <http://nxt.org/what-is-nxt/voting> (last visited Feb. 26, 2016).

152. This is particularly a concern in close corporations, where shareholders' interests are more likely to vary than in public corporations. See generally Robert C. Art, *Shareholder Rights and Remedies in Close Corporations: Oppression, Fiduciary Duties, and Reasonable Expectations*, 28 J. CORP. L. 371, 376–402 (2003) (reviewing states' approaches to oppression).

153. Frank H. Easterbrook & Daniel R. Fischel, *Voting in Corporate Law*, 26 J.L. & ECON. 395, 396, 420 (1983); Michael S. Kang, *Shareholder Voting as Veto*, 88 IND. L.J. 1299, 1305–15 (2013) (explaining how shareholder ignorance affects corporate governance, including leading to primacy of the board).

154. See, e.g., BRYAN CAPLAN, *THE MYTH OF THE RATIONAL VOTER: WHY DEMOCRACIES CHOOSE BAD POLICIES* (2008) (discussing voter ignorance); ILYA SOMIN, *DEMOCRACY AND POLITICAL IGNORANCE: WHY SMALLER GOVERNMENT IS SMARTER* (2013) (arguing that voter ignorance justifies small government); Philip K. Hastings, *The Voter and the Non-voter*, 62 AM. J. SOC. 302 (1956) (discussing selection bias effects resulting from nonvoting).

155. See, e.g., Michael S. Kang, *Democratizing Direct Democracy: Restoring Voter Competence Through Heuristic Cues and “Disclosure Plus,”* 50 UCLA L. REV. 1141 (2003) (discussing how referenda could be improved by providing better information to voters).

irrational for most other Bitcoin owners to take the time to learn about checkpointing issues and vote their own shares. Thus, self-interested miners—even just a small minority of Bitcoin owners—might be able to make decisions to benefit their own interests.

A vote-buying mechanism faces even greater problems along these lines.¹⁵⁶ A vote-buying scheme is effectively an auction, and the outcome that receives the most financial support is chosen as policy. Bitcoin owners could send Bitcoins to one address to register support for a checkpoint and a different address to register opposition. These would be public keys created without corresponding private keys, so sending the Bitcoins would destroy them. Holders of small stakes would have little incentive to try to buy their preferred outcomes. This would be true even if there were a policy providing for refund of the Bitcoins spent by the losing side or the winning side; as long as there is any probability one will lose one's Bitcoins, the optimal individual strategy is to free-ride. Moreover, small holders of Bitcoins would have little incentive to become informed in the first place.

Vote buying has long been viewed as undemocratic,¹⁵⁷ but recent research has suggested that a variation on vote buying could work in the corporate context. E. Glen Weyl describes quadratic vote buying, in which the cost of votes purchased is a quadratic function of the number of votes.¹⁵⁸ For example, someone buying two votes would pay four times as much as someone buying one vote. The cost of a *marginal* vote is thus linear in the amount of votes purchased,¹⁵⁹ counterbalancing the increasing marginal benefit of votes. Weyl and Eric Posner argue that this could be especially useful for corporate law because it addresses the concern that existing shareholder voting relies on shareholders who may not have sufficient information to vote.¹⁶⁰ An approximation they call square-root voting would simply allow shareholders to vote the square root of the number of shares they own.¹⁶¹

Weyl recognizes, however, the danger of “de-merging,” in which a single individual pretends to be multiple individuals.¹⁶² While he argues that quadratic voting reduces the danger of this, he also shows that the relative inefficiency

156. For an analysis of vote buying in corporate law, see Thomas J. Andre, Jr., *A Preliminary Inquiry into the Utility of Vote Buying in the Market for Corporate Control*, 63 S. CAL. L. REV. 533 (1990).

157. See Richard L. Hasen, *Vote Buying*, 88 CALIF. L. REV. 1323, 1348 (2000).

158. E. Glen Weyl, *Quadratic Vote Buying* (Apr. 2013) (unpublished manuscript), <http://economics.uchicago.edu/workshops/Weyl%20Glen%20-%20Quadratic%20Vote%20Buying.pdf>.

159. *Id.* at 1.

160. See Eric A. Posner & E. Glen Weyl, *Quadratic Voting as Efficient Corporate Governance*, 81 U. CHI. L. REV. 251, 262–63 (2014) [hereinafter Posner & Weyl, *Efficient Corporate Governance*]; see also Eric A. Posner & E. Glen Weyl, *Voting Squared: Quadratic Voting in Democratic Politics*, 68 VAND. L. REV. 441, 499–500 (2015) (arguing that quadratic voting also could be useful for democratic institutions).

161. Posner & Weyl, *Efficient Corporate Governance*, *supra* note 160, at 266–67.

162. Weyl, *supra* note 158, at 21–22.

caused by de-merging will be on the order of the number of separate identities created by a de-merge.¹⁶³ With Bitcoin, it would be trivially cheap for a Bitcoin owner to separate its interests into any arbitrary number of interests, perhaps using a “mixing” service to make it impossible to prove a common origin.¹⁶⁴ This defeats quadratic voting. The scheme could be used with Bitcoin only if voting were restricted to verified identified individual owners of Bitcoins. Perhaps a peer-to-peer system for such verification could be developed, but at least for now, it does not seem a viable mechanism for achieving peer-to-peer governance.

The jury system ensures that some decisions are made by individuals that have scrutinized the evidence with some care. Could we adapt the mechanism to Bitcoin by designating a random sample of Bitcoin owners to make decisions, such as whether to approve a checkpoint? It would certainly be possible to select random Bitcoin owners, with the probability of being selected proportional to their interests. One might even imagine a system for punishing selected users who refused to cast a vote.¹⁶⁵ But forcing them to engage in reasoned decision-making is likely to be much more difficult. We could require evidence of reasoned decision-making, such as a written opinion.¹⁶⁶ But there are two problems with this approach. First, some Bitcoin owners might offer to create such evidence for others, but that again would shift power to those with greater stakes in decisions. Perhaps we could police such activity, but that would require normative judgment. Second, assessing whether someone has engaged in reasoned decision-making requires normative judgment. Thus, the problem is recursive. Insisting on reasoned decision-making requires more reasoned decision-making.

This does not mean that it would be impossible to build peer-to-peer governance on a voting mechanism or on a random selection jury-like mechanism. The problems that we have identified exist in our own familiar democratic and corporate institutions, yet they endure. The adaptations are complex. In corporations, for example, voters elect board members and entrust those board members to make decisions.¹⁶⁷ And of course, voting for representation is the critical feature of republican government. Perhaps similar adaptations could be imagined for Bitcoin, with identifiable Bitcoin owners electing representatives who have the limited role of supervising voting by anonymous Bitcoin owners. There may, however, be an alternative, a decision-making process that is peer-to-

163. *Id.* at 21.

164. Simon Barber et al., *Bitter to Better — How to Make Bitcoin a Better Currency*, in FINANCIAL CRYPTOGRAPHY AND DATA SECURITY 399, 411 (Angelos D. Keromytis ed., 2012) (discussing the use of mixers by Bitcoin users to provide anonymity).

165. *Cf.* Note, *The Case for Compulsory Voting in the United States*, 121 HARV. L. REV. 591, 600 (2007) (discussing the case for compulsory jury service).

166. The norm of written opinions for judicial decision-making can be justified in part on the ground that it forces judges to be careful in their reasoning. *See generally* Gerald Lebovits et al., *Ethical Judicial Opinion Writing*, 21 GEO. J. LEGAL ETHICS 237, 294 (2008) (discussing the assumption that judges deliberate each issue carefully).

167. *See generally* Stephen M. Bainbridge, *Director Primacy: The Means and Ends of Corporate Governance*, 97 NW. U. L. REV. 547 (2003) (highlighting the importance of the board of directors in corporate governance).

peer to the core. All peer-to-peer mechanisms—including Bitcoin, file-sharing,¹⁶⁸ and other projects¹⁶⁹—are built to achieve cooperation, even in the face of hostile adversaries who would destroy the system or manipulate it to their own benefit. Group decision-making can be seen as just another such problem, and the peer-to-peer challenge is to create a coordination rule that produces clear and generally justifiable decisions in such conditions.

2. Resolution with Tacit Coordination

How would a peer-to-peer checkpointing system work? The goal of this Article is not to describe the full range of peer-to-peer normative decision-making systems or even to identify the best system, but rather to illustrate proof of concept. Accordingly, we will begin with a simple peer-to-peer mechanism for constructing a tacit coordination game to make a binary decision.

Someone could support a new checkpoint by paying a proposal fee to a pre-established address not under any individual's control and including in the transaction metadata a reference to the hash representing the block that would serve as the checkpoint being proposed. Once the proposal transactions were added to the block chain, all could recognize the normative decision-making process's initiation. Individuals would then have a fixed period of time within which to dedicate Bitcoins to supporting or opposing the checkpoint. One would demonstrate support or opposition by transferring Bitcoins to designated addresses, such as public keys generated from hashes of the proposal transaction followed by strings such as "Yes" and "No." The chance that anyone would own these addresses is infinitesimally low because the system for generating keys prevents anyone from purposefully generating a private key corresponding to any particular public key.

The purpose of all these transactions is simply to create a convention for announcing support or opposition to a particular proposal, and the proposal fee would count as an initial announcement of support for the proposal. The fixed period could be measured in number of blocks to be added to the block chain from the time of the initial proposal, but if there were a sufficient amount of activity at the end, the time period would be automatically extended.¹⁷⁰ The winning position would be the position with the most support, and the money dedicated by the losers would be allocated to those supporting the winning position. Earlier supporters would receive money before later ones, so there would be no incentive

168. The file-sharing example highlights that cooperation may exist not merely as a result of rational self-interested calculation, but also as a result of social norms. See Lior Jacob Strahilevitz, *Charismatic Code, Social Norms, and the Emergence of Cooperation on the File-Swapping Networks*, 89 VA. L. REV. 505, 537–47 (2003) (discussing the role of social norms in file-sharing communities).

169. See Moshe Babaioff et al., *Incentives in Peer-to-Peer Systems*, in ALGORITHMIC GAME THEORY 593–94 (Noam Nisan et al. eds., 2007).

170. The extension criterion could be that the resolution will be extended (perhaps for two more blocks) if either of the most recent two blocks would change the outcome. Thus, an attempt to engineer a surprise, large allocation at the last minute would fail to surprise, and others would then have an opportunity to make opposing allocations.

to add money to the winning side at the last minute. This reallocation would be accomplished by the generation of new Bitcoins for the winners.

For example, suppose that a proposal is made to add a checkpoint based on block *X*. The proposal is initiated by *A*, who is required to pay a proposal fee of at least (let us suppose) 1 Bitcoin, and *A* meets that minimum obligation. Suppose that *B* places 2 Bitcoins against the proposal, and then *C* places 1 Bitcoin against the proposal. (We will count *B*'s transaction as occurring first if it is listed earlier in the block chain.) If there are no other transactions, then the proposal fails, 3 to 1. This would entitle *B* to a generation transaction of 3 Bitcoins altogether (including the 2 Bitcoins to be refunded) and *C* to a generation transaction of 1 Bitcoin. Because all the original Bitcoins that were voted are effectively destroyed, the total number of spendable Bitcoins in existence remains constant. Of course, if *D* had buttressed *A*'s position by spending another 10 Bitcoins, and no other transactions occurred, then the checkpoint would be approved. *A* would receive 2 Bitcoins (equal to its investment plus 1 of *B*'s Bitcoins) and *D* would receive 12 Bitcoins (its investment plus 1 of *B*'s and 1 of *C*'s). It does not matter whether *D* and *A* are, in reality, accounts owned by the same person. If another party, *E*, also supported *A*'s position, then *E* would simply have the funds invested refunded, without receiving anyone else's Bitcoins.

This is a tacit coordination game in which potential participants must anticipate whether more funds will be distributed in favor of one position than in another. After *A* devotes its 1 Bitcoin in favor of the checkpoint, *B* must consider whether to match *A*'s Bitcoin. If *B* matches, *B* will want to at least marginally exceed the amount offered by *A* so that *B* will win if there are no further transactions.¹⁷¹ *B* will have an incentive to match and at least marginally exceed *A* if *B* believes that no one else will participate, or if there is participation, more participants eventually will place money against the proposal than in its favor.

Of course, *A* will have an incentive to fight back to win its initial bet against *B*. One side might exceed the other's contribution by a sizable amount as a way of signaling its fortitude. Taking a large position, however, has two effects. On the one hand, it shows the resoluteness of the party putting that amount of money in support of a position, perhaps implying that the party is willing to put even more money down in favor of the same position. On the other hand, it also increases the chance that third parties will be drawn into the game. There is at least some fixed cost associated with initial entry into the game, including consideration of the issue to be resolved, but a contribution that exceeds the prior one functions as an offer to enter into a bet, and it will be worth taking the time to consider this if the offer is large. A party supporting a position must ask the ultimate question: What would third parties decide to do if they focused on the issue to be resolved.

171. *B* could be required to exceed *A*'s payment by at least some amount, such as the original proposal fee. This has the advantage of providing a subsidy to encourage additional participation after *B*'s turn, but may reduce the incentives *B* has to participate. If *B* only barely exceeds *A*, there will be at least one party with an incentive to challenge *B*: *A*.

The game is thus one of tacit coordination in which any participant must anticipate what hypothetical other participants will do in the future, recognizing that those hypothetical other participants will be looking prospectively at still other hypothetical participants. The dynamics of this particular tacit coordination game are similar to those of an all-pay auction,¹⁷² where existing investments are not sunk, so participants have incentives to bolster their investments. But rational participants in such an auction would recognize that others have the same incentives, and risk aversion would counsel against throwing good money after bad.

Everyone's incentive is to do what people in the future will do, with no authoritative answer disciplining the participants. So the incentive is to look for focal points that serve as tacit coordination devices, and the ultimate question—whether a particular checkpoint should be added to the block chain—serves as a natural focal principle. Collusion is difficult, because even if existing participants collude, someone could try to combat the collusion and create interest that would draw more third parties—from the virtually unlimited pool of Bitcoin owners—into the game. The initial participants are likely to have relatively high knowledge because of the need to anticipate others' decisions, and later participants are likely to have low knowledge. Later participants, however, might be drawn by the high stakes and conduct research to gain knowledge.

Of course, it is possible that there could be alternative focal points, but there will be so many of them that they will tend to cancel out.¹⁷³ For example, one could argue that the original proposal is focal or the first position that someone takes is focal or the most recent position that someone takes is focal, but it is hard to see why any of these differentiates itself from any other. Similarly, one could look to see who is making the most noise in favor of a particular position, but if that could change a focal point, then everyone would scream.

This argument is, admittedly, somewhat informal, and it may seem inconsistent with game theory. As Hykel Hosni points out,¹⁷⁴ coordination games generally involve multiple Nash equilibria.¹⁷⁵ If one expects others to follow a particular focal point, one should follow that as well, so the Nash equilibrium concept does not predict a particular focal point. However, when there are multiple Nash equilibria, coordination will often be around the solution that produces the highest payoffs to the players.¹⁷⁶ The question of whether to use normative focal

172. An all-pay auction is one in which the losers pay the amount of their bids. See, e.g., Michael R. Baye et al., *The All-Pay Auction with Complete Information*, 8 ECON. THEORY 291 (1996) (providing an economic model).

173. See Abramowicz, *supra* note 22, at 548–56.

174. See Hykel Hosni, *Interpretation, Coordination and Conformity*, in GAMES: UNIFYING LOGIC, LANGUAGE, AND PHILOSOPHY 37 (Ondrej Majer et al. eds., 2009).

175. *Id.* at 46–47.

176. See *id.* at 47 (discussing a version of the Battle of the Sexes game in which both members of a couple, who must make independent decisions without communication, prefer attending a Bach concert to a Stravinsky concert, but the Stravinsky concert is also a Nash equilibrium because if one attends that concert, the other would prefer attending together than attending separately); see also Anna Gunnthorsdottir & Palmar Thorsteinnsson,

points in general is a tacit coordination problem that will apply across many checkpoint disputes, and those who participate in these disputes will likely be better off if the system works than if it fails. Moreover, Hosni argues that in coordination games “agents should apply reasons to discard those possible strategies that will prevent them from conforming on their mutual expectations”¹⁷⁷—for example, because multiple strategies that point in different directions are indistinguishable—and “a *perfect reason* will be a choice function which always returns a singleton, a unique strategy.”¹⁷⁸ The normative argument for using normative focal points is, in this framework, a reason that returns a single strategy using normative focal points.

What does it mean for the normative question—in this case, whether a checkpoint should be added—to serve as a focal point? Perhaps a more precise statement of the principle would be that the focal point is the best answer to the question. The Bitcoin software code itself includes a comment indicating that “a good checkpoint block”¹⁷⁹ should be “surrounded by blocks with reasonable timestamps”¹⁸⁰ and “[c]ontains no strange transactions.”¹⁸¹ The existing checkpoints seem to have round block numbers,¹⁸² perhaps to emphasize that the checkpoint is arbitrary, rather than designed to achieve some advantage. A checkpoint should be sufficiently recent that it is useful but sufficiently old that the probability of it being dropped from the block chain is extremely small. Thus, the checkpoint functions solely to speed up and solidify block chain analysis rather than to change the outcome. One could, of course, debate the relative importance of all of these considerations. Some of these considerations may not even matter at all. Participants in the peer-to-peer checkpoint process may engage in some form of online debate. But someone considering all of these factors will likely come to some conclusion about how strong the case is for a new checkpoint. Different people may come to different conclusions, and people may change their views once others credibly signal their own views. But it is this familiar process of trying to figure out the best answer to a problem that seems likely to constitute the search for the focal point.

Tacit Coordination and Equilibrium Selection in a Merit-Based Grouping Mechanism: A Cross-Cultural Validation Study (July 11, 2011) (unpublished manuscript), <http://www.ssrn.com/abstract=1883465> (demonstrating tacit coordination on the highest payoff option in a laboratory experiment).

177. Hosni, *supra* note 174, at 49.

178. *Id.*

179. *Checkpoints, supra* note 146.

180. *Id.* Sometimes, a Bitcoin block will have a timestamp before a block that is nominally earlier in the block chain. See CodesInChaos, Comment to *What Is the Standard Deviation of Block Generation Times?*, STACKEXCHANGE (Sept. 14, 2012, 9:13 AM), <http://www.bitcoin.stackexchange.com/questions/4690/what-is-the-standard-deviation-of-block-generation-times> (explaining that timestamps may not be accurate and that the difference between blocks may even be negative). The Bitcoin software does not seek to provide a peer-to-peer mechanism for ensuring that timestamps are accurate.

181. See *Checkpoints, supra* note 146.

182. See *id.* (including blocks such as 168,000 and 295,000 as checkpoints).

Does this focal point decision-making mechanism share the flaws of other peer-to-peer mechanisms? Superficially, it might appear to be quite similar to the vote-buying mechanism. That approach and the focal-point approach both ultimately resolve a question based on which of two positions attracts a larger number of Bitcoins. But the incentives are critically different. In the focal-point mechanism, those who contributed to the winning side receive the contributions in support of the losing side. We must still address, however, whether apathy might lead to poor decision-making and whether the process is likely to be biased in favor of concentrated interests. The financial incentive to be on the winning side is central to addressing both questions. Regarding the concern about apathy, the mechanism requires only a few individuals to participate, and it gives those individuals incentives to inform themselves sufficiently to enable predictions of what the final resolution might be. The larger the amount at stake, the greater the incentives will be to acquire information and generate arguments.

The danger that some Bitcoin owners' interests might be given a high degree of weight, however, is more serious. An initial concern might be that anyone with self-interest would be able to bias the process, even if that owner has only a small number of Bitcoins relative to the broader community of potential decision-making participants. There is some danger of this because it will be rational for participants to change their assessment of the focal point given signals from others. An investment in a particular position might reflect a genuine view of the focal point *or* an attempt at manipulation, but the former possibility will receive at least some weight. The more common attempts at manipulation are, however, the less weight they are likely to receive in focal-point analysis. Moreover, such attempts will generally encourage others to participate in the process because making an investment inconsistent with the focal point provides a financial opportunity for those on the opposite side.

Overall, the effect is similar to that of "noise traders" in the stock market, who make their decisions for reasons other than market fundamentals. These noise traders can influence prices, but they also attract more participation from sophisticated parties.¹⁸³ This can make stock prices more precise overall,¹⁸⁴ though less precise in cases in which there is more self-interested participation than expected. Whether self-interested participation, on net, improves or decreases accuracy is an empirical question, though empirical evidence from analogous contexts is encouraging.¹⁸⁵ Even isolated effects of self-interested advocacy in individual cases may be troubling, but self-interest affects many types of decision-making, including lobbying in the legislative arena and hiring local counsel who

183. See, e.g., JOHN L. TEALL, *FINANCIAL TRADING AND INVESTING* 118 (2012) (discussing how investment and trading decisions made by noise traders create opportunities for more sophisticated traders and investors).

184. See, e.g., Matthew Spiegel & Avandhar Subrahmanyam, *Informed Speculation and Hedging in a Noncompetitive Securities Market*, 5 *REV. FIN. STUD.* 307 (1992).

185. See Robin Hanson et al., *Information Aggregation and Manipulation in an Experimental Market*, 60 *J. ECON. BEHAV. & ORG.* 449 (2006) (showing that attempts to manipulate prediction markets generally increase market accuracy by improving liquidity).

knows the judge in the judicial arena. The mechanism described here at least gives third parties incentives to identify manipulation and challenge it.

A concern of potentially greater magnitude is that the existence of self-interest may change the focal point. This seems unlikely when the self-interest is contained to individuals with a relatively small number of Bitcoins. The problem is more severe, however, if the self-interest affects a large proportion of the Bitcoin community—especially the community that participates in adjudicative decision-making. If participants in the tacit coordination game expect that there is a high probability that the later participants will be Bitcoin miners, for example, they might try to identify the normative focal point from the *perspective* of the Bitcoin miners. This provides a strong argument for including individuals other than Bitcoin miners in the normative decision-making process.

It seems unlikely, however, that the interests of Bitcoin miners will receive disproportionate weight in any normative analysis. Each participant will have more at stake from the money being placed at risk than from collateral consequences, so even if someone would like to collude with other Bitcoin miners, a participant will have an incentive to defect if the ultimate focal point is expected to be some distance from what the miners would prefer. The focal point is unlikely to depend much on who participates in the tacit coordination game, because there is a strong normative argument for considering the welfare of the entire Bitcoin community. Even if it did depend on the identity of the participants, there will be at least some participants who are not Bitcoin miners, and indeed there is little reason to think that those who mine will be especially likely to participate in decision-making. And even if most active participants are miners (and this seems unlikely), in the pool of *potential* participants, the proportion of Bitcoin miners seeking to obtain some collateral advantage will be small. In short, it seems doubtful that the interests of large concentrated groups, like miners, will receive greater weight than the interests of the broader public. They might receive slightly more weight, but large concentrated groups receive *much* more weight in democratic processes,¹⁸⁶ so at least this seems likely to be an improvement.

We cannot predict the result of tacit coordination games based on theory alone. There are multiple equilibria in any tacit coordination game, so game theory alone cannot determine which equilibrium will result.¹⁸⁷ We have seen, though, that Bitcoin's success already is dependent on multiple forms of tacit coordination, and this is true for other institutions. The "ultimate rule of recognition"¹⁸⁸ that

186. See generally MANCUR OLSON JR., *THE LOGIC OF COLLECTIVE ACTION: PUBLIC GOODS AND THE THEORY OF GROUPS* (1971) (providing the seminal analysis of this phenomenon).

187. Experiments, however, have suggested that efficiency concepts can narrow down the range of potential equilibria to those that are "payoff-dominant," i.e., "not strictly Pareto dominated by any other equilibrium point." Van Huyck et al., *supra* note 19, at 236.

188. H.L.A. HART, *THE CONCEPT OF LAW* 107–08 (Penelope A. Bulloch & Joseph Raz eds., 2d ed. 1994) (defining this as a rule not validated by any superior norm or rule); see also Andrei Marmor, *Legal Conventionalism*, 4 *LEGAL THEORY* 509 (1998) (discussing whether the "rule of recognition" can be viewed as a coordination mechanism).

results in the acceptance of legitimacy of governments can be viewed as the outcome of a tacit coordination game. If that game produced a general perception tomorrow that Bozo the Clown is dictator, then the Era of Bozo would begin. We do not worry about that in the United States because the tacit understanding making the government legitimate is strongly entrenched. History teaches that tacit coordination can, but does not always, produce great stability.

Peer-to-peer governance could be introduced gradually, allowing for testing. The decision-making apparatus initially might serve as a tool for recommending decisions to the Bitcoin software repository administrators. Initially, they might ignore it altogether, choosing their own checkpoints instead. But if it produced reasonable recommendations, they might establish a weak presumption in favor of following the checkpoint recommendations of the peer-to-peer decision-making system, and perhaps later, a strong presumption. A decision-making apparatus may potentially become the exclusive mechanism by which administrators decide whether to add checkpoints. Eventually, the client software might be modified so that it automatically incorporates a checkpoint whenever the analysis of the block chain reveals a completed decision-making process recommending one. If this proves problematic, the administrators could remove this feature from the client software. But gradually increased reliance on peer-to-peer decision-making will build legitimacy over time.¹⁸⁹

B. Evolution of the Reference Code

Checkpoints are a relatively trivial aspect of Bitcoin operations. Currently, a checkpoint is added only as part of a regular client-software update. Some competing cryptocurrencies checkpoint much more often, but this may be because they think checkpointing is a necessary security precaution in the absence of a proof-of-work system.¹⁹⁰ So, decentralizing checkpointing would make only an incremental difference in the degree to which Bitcoin decision-making is peer-to-peer. A more fundamental innovation would be to use peer-to-peer decision-making to resolve whether to change the client software in the official repository. Fully implemented, this innovation would allow decisions recorded on the block chain to determine whether changes should be made to the source code. In principle, this could be used for any open-source project, and to generate or amend documents of any kind, including public or private rules and regulations.

189. See Emanuela Carbonara et al., *Legal Innovation and the Compliance Paradox*, 9 MINN. J.L. SCI. & TECH. 837, 854–56 (2008) (discussing how phased implementation can be useful in building support for reforms).

190. See Vitalik Buterin, *Feathercoin: Interview with Peter Bushnell*, BITCOIN MAG. (Aug. 12, 2013), <http://www.bitcoinmagazine.com/6263/feathercoin-interview-with-peter-bushnell/> (arguing that Feathercoin has an advanced checkpointing system that makes it more resilient to attacks than Bitcoin).

Open-source projects are generally managed with the assistance of versioning software (the current most popular versioning protocol is Git¹⁹¹), which amounts to a more powerful version of the “track changes” feature in popular word processors. This software allows users to make a version of the software code, change it, and then propose that it be integrated into the official version. For example, a recent proposed change to Bitcoin involved adding a feature enabling the software host to limit the total bandwidth it uses.¹⁹² The user who proposed this had created a remote fork of the master branch of the project, copying all of the code files to a repository under that user’s control. After making changes, the user later filed a “pull request”¹⁹³ for the remote fork to be merged into the master branch, which would involve changes to 14 different files. Ordinarily, a user filing a pull request will have incorporated changes made in the master branch because of the original creation of the remote fork.¹⁹⁴ The centralized software developer or developers can choose whether to accept a pull request.¹⁹⁵ Periodically, the centralized developer will create a new branch within the repository, designating a new version of the open source software by forking from the master branch. This new branch may include several sets of new features and other changes, such as documentation improvements and bug fixes. Anyone can then compile the software, and some websites host the compiled versions,¹⁹⁶ including installers for multiple operating systems.

The critical determinations necessary to control the development of a software repository are whether to accept a proposal to pull changes from a remote fork into the master fork and whether to create a new branch of the software based on the current master branch. Ideally, peer-to-peer decision-making also would control whether to create experimental branches and whether to approve pulling

191. See Comparing Trends of Git, Subversion, and Svn, GOOGLE TRENDS, <http://www.google.com/trends/explore#q=git,subversion,svn> (from Google Trends, search “git, subversion, svn”) (showing that Git overtook subversion in popularity around 2009).

192. Jmcorgan, Comment to *Adds Publishing Blocks and Transactions Over ZMQ*, GITHUB (July 27, 2014, 5:05 PM), <http://www.github.com/bitcoin/bitcoin/pull/4594> (offering a set of changes that purport to facilitate broadcasting of information on newly generated blocks and new transactions among Bitcoin nodes).

193. A pull request is a request for software changes to be incorporated in the master branch. *Using Pull Requests*, GITHUB HELP, <http://help.github.com/articles/using-pull-requests/> (last visited Feb. 26, 2016). The software hosting the repository creates discussion forums built around each pull request. See Jmcorgan, *supra* note 192. This allows users to discuss the changes. The user who creates the pull request can then make further changes in response to feedback. *Id.*

194. A significant function of versioning software is to facilitate integration of different sets of changes, which may conflict with one another. See, e.g., *Resolving Conflicts*, GITHOWTO, http://www.githowto.com/resolving_conflicts (last visited Feb. 26, 2016) (explaining how to resolve conflicts in Git).

195. *Merging a Pull Request*, GITHUB HELP, <http://help.github.com/articles/merging-a-pull-request/> (last visited Feb. 29, 2016) (“Anyone with push access to the repository can complete the merge.”).

196. See, e.g., BITCOINCORE, <http://www.bitcoin.org/en/download> (last visited Mar. 29, 2016).

changes into these branches. This would enable peer-to-peer decision-making about the development of features, rather than only about whether a final proposed version of a feature should be accepted in the master branch. For any particular decision, the process could mirror the normative decision-making process for determining whether to accept a proposed checkpoint described above.¹⁹⁷ A user would pay a proposal fee by making a Bitcoin (or other cryptocurrency) payment and using metadata to indicate in some concise way the nature of the proposal.¹⁹⁸ Owners of Bitcoins (or other cryptocurrency) could then send money in support of or opposition to a proposal. Eventually, the losing side's contribution would be distributed to the winning side.

Those who have control over a repository could observe when decisions were final and update it accordingly. Of course, they might choose to disregard changes, but anyone else could create a version of the repository that included the relevant changes. Peer-to-peer decision proposals could specify the hash of the repository that would exist if those proposals were implemented, and those downloading repository code could confirm a hash match. So, if a norm of peer-to-peer decision-making were clearly established, there might be many repositories that were mirrors of one another, and operators of client software simply would reject the repositories that were not up-to-date. A peer-to-peer decision-making process removes the need for any one software repository to be designated or even thought of as the official one. We have seen that the Bitcoin protocol establishes a mechanism for determining which of the competing block chains should be accepted as the correct one, and the peer-to-peer decision-making protocol would ensure that the authoritative block chain could determine which of the competing software repositories should be considered authoritative.

With these decision-making elements in place, the peer-to-peer decision-making process would resemble legislative processes for proposing legislation, offering amendments, and amending amendments. In contrast to the process followed in *Robert's Rules of Order*,¹⁹⁹ however, more than one set of issues could be debated at any particular time.²⁰⁰ Of course, the normative evaluation to approve a proposal to merge a set of changes into a master branch involves, in

197. See *supra* Section II.A.2.

198. For example, the user might report a hash of the proposed changes. Others could then search the Internet using services such as Google or file-sharing services to find the file with the reported hash. Presumably, users would reject a proposal with a hash that could not be identified. It would also be possible to place full proposals directly on the block chain, though a proposal that contained a significant amount of data might contribute to the problem of "block chain bloat." See generally Daniel Cawrey, *Why New Forms of Spam Could Bloat Bitcoin's Block Chain*, COINDESK (Sept. 3, 2014), <http://www.coindesk.com/new-forms-spam-bloat-bitcoins-block-chain/> (discussing the bloat problem).

199. See HENRY MARTYN ROBERT III, *ROBERT'S RULES OF ORDER NEWLY REVISED* 371–75 (Sarah Corbin Robert et al. eds., 11th ed. 1970).

200. For an analysis of how Robert's Rules could be adapted to an online setting, see Phil Reiman, *In Congress Electric: The Need for On-Line Parliamentary Procedure*, 18 J. MARSHALL J. COMPUTER & INFO. L. 963 (2000).

part, an assessment of whether this is the appropriate time to do so. It might not make sense to accept Change *X* yet because it makes sense for Change *Y* to be considered, either because Change *Y* is more important, or because Change *Y* was developed earlier and might affect the desirability of Change *X*. Often, it will make sense to achieve group consensus on a general principle for proceeding before full development of that principle into code. So, a proposal to allow a particular change might fail at one time but succeed later.

Similarly, the normative case for creating a new version of a repository might be weak at one time, but stronger a few weeks later when more time has passed and more testing has taken place. Peer-to-peer decision-making may be more chaotic than a structured meeting with a recognized chair, but it should be able to resolve issues in a reasonable order. A peer-to-peer decision-making mechanism also might support group decisions on whether to change the *time* at which a particular decision is to be resolved. This would reduce the risk associated with supporting or opposing a proposal, because decision-makers on the primary question could focus on the overall merits, while others could focus on questions of timing.

Software tools could be developed that would automatically update repositories based on determinations in the block chain. But this is not essential. What *is* essential is the general acceptance of the principle that the block chain, pursuant to the decision-making mechanism described above or some other peer-to-peer mechanism, determines the software. This completes a circle: the Bitcoin protocol determines the block chain, and the block chain determines the Bitcoin protocol. The existence of this circle would enable evolution with respect to both the rules determining the authoritative block chain—e.g., if decision-makers incorporated a proof-of-stake component into Bitcoin, and the rules governing the determination of what counts as an authoritative decision—e.g., the mechanics of the formal tacit coordination game.

The possibility of changing the decision-making process may decrease the chance of total rejection of the decision-making system. But such rejection will always be possible. Anyone can make a normative argument that other participants in an open-source software project should use some version of the software other than the officially sanctioned one, or that one set of agreed-upon rules should be disregarded in favor of another set, regardless of their respective pedigrees. Establishing a peer-to-peer system for making decisions, however, can provide perceived legitimacy to the corresponding software repositories, at least if peer-to-peer decision-making came to be accepted over time. It would seem strange for someone to advocate immediate change to some alternative software repository not recognized by the official process, simultaneously repudiating *both* the decision-making rules and the decisions made pursuant to those rules. Constitutional law analogously often successfully channels demand for change into either calls for

changes consistent with the Constitution or for changes to the Constitution, pursuant to the provisions set forth by the Constitution itself.²⁰¹

Decisions might be made to entrench some rules of decision or some aspects of the Bitcoin protocol by establishing a higher threshold of decision,²⁰² creating a form of higher-order law analogous to constitutional law. This provides a possible response to a plausible objection to using peer-to-peer decision-making for Bitcoin in particular. One feature of Bitcoin trumped by some advocates is that there is no central bank deciding on monetary policy, because the schedule of Bitcoins to be produced was fixed at the outset.²⁰³ People will be more comfortable holding Bitcoin, the argument goes, knowing that this serves as a check against inflation. The argument is similar to that offered by those, most notably Milton Friedman,²⁰⁴ who argue that monetary policy should be conducted according to rules set forth in advance. There are familiar rebuttals that can be adapted to Bitcoin. For example, the ability to determine the currency's course of growth would make it possible to adapt to unexpected needs.²⁰⁵ If, however, flexibility in decision-making produced too much inflation, the Bitcoin mining schedule could be made unchangeable or difficult to change.

Creation of higher-order principles is not a foolproof safeguard against change, because there could be a decision to change a higher-order decision.²⁰⁶

201. In the United States, legal change is manifested in statutes or in constitutional amendments under Article V of the Constitution. It remains possible, however, that the people could reject the Constitution. Akhil Amar has argued that the Constitution specifically recognizes the right of the people to do so. *See* Akhil Reed Amar, *The Consent of the Governed: Constitutional Amendment Outside Article V*, 94 COLUM. L. REV. 457, 458–59 (1994).

202. For example, a proposal might require that two-thirds of money placed down be in favor of changes, if those changes affected certain documents or code sections in the repository. This would decrease the potential gains from supporting such a change. One could also imagine a provision simply requiring some high standard, such as “very high confidence,” for certain types of changes. The peer-to-peer decision-makers would then decide whether the particular proposal met that high standard.

203. *See, e.g., How Does Bitcoin Work?*, ECONOMIST (Apr. 11, 2013, 11:50 PM), <http://www.economist.com/blogs/economist-explains/2013/04/economist-explains-how-does-bitcoin-work>; *Controlled Supply*, *supra* note 96.

204. *See* Milton Friedman, *Monetary Policy: Theory and Practice*, 14 J. MONEY, CREDIT & BANKING 98, 100–01 (1982) (discussing the rules that should be adopted to manage monetary policy).

205. *See, e.g.,* Scott Sumner, *In Defense of a Flexible Monetary Policy*, CATO UNBOUND (Nov. 8, 2013), <http://www.cato-unbound.org/2013/11/08/scott-sumner/defense-flexible-monetary-policy>.

206. An analogy in American constitutional law might be a change to the representation of states in the Senate. The Constitution guards against this change even by constitutional amendment. *See* U.S. CONST. art. V (“[N]o state, without its consent, shall be deprived of its equal suffrage in the Senate.”). Some observers, however, have argued that it would be possible first to amend the Constitution to remove the obstacle to this type of amendment to the Constitution, and then to amend the Constitution either to change the representation of states in the Senate or to allow subsequent legislation to do so. *See, e.g.,*

Arguably, however, such a higher-order decision could provide better protection against inflation than the existing decision-making system. After all, there is nothing to stop the software developers of the official Bitcoin repository from changing the schedule at which Bitcoins can be released, other than the possibility that others will reject the authoritativeness of their repository. The most plausible, foreseeable clash of competing interests is the possibility that miners will demand greater rents, whether in the form of additional Bitcoin mining allowances or transaction fees. This would provide some benefit to the public by increasing the costs of mounting a 51% attack.²⁰⁷ But miners' interest would be in higher rents than the public would favor. The miners could claim to be the authentic representatives of the Bitcoin community and reject the official software repository. Faced with a credible threat of a hard fork, those who control the central repository seem likely to give into the miners' position, at least partly. Current developers may insist that they act solely on the basis of consensus, but the drastic future reduction in the issuance of new Bitcoins means that transaction fees will have to increase at least somewhat. Given differing interests, consensus as to how great the increase is seems unlikely. The most probable outcome will be quite close to the interests of the miners.

Creating a peer-to-peer decision-making system could help avoid this outcome. If formal tacit coordination games became the accepted mechanism for determining change to the Bitcoin protocol, there would remain the possibility of tacitly coordinating around some other result. For example, peer-to-peer decision-making could decide against minimum transaction fees, but if most miners acted as if blocks with transactions below some hypothetical minimum were invalid, the Bitcoin protocol effectively would insist on transaction fees. But such coordination might be more difficult than it would be *absent* a formal mechanism for making decisions about the protocol. Miners seeking to coordinate amongst themselves to create minimum transaction fees would not be merely advocating that policy, but advocating rejection of the entire peer-to-peer decision-making system and replacement with some other system. This makes the change more radical and, thus, more difficult to tacitly coordinate upon.

The greatest challenge for peer-to-peer decision-making may be the difficulty of initiating it, even with gradual introduction. Miners, of course, would be able to see that such a system might lead, in the long term, to the reduction of their power. They might, therefore, resist peer-to-peer decision-making, likely focusing on legitimate concerns that it has not sufficiently been tested. As long as the Bitcoin software developers proceed by consensus, it is unlikely to be adopted. But if the developers did gradually institute such a system, the miners today might not yet have sufficient incentive to risk a hard fork of the currency. After all, any long-term adverse consequences for the miners might be years away, and the rents for miners so far in the future will likely be dissipated by hardware investments in the interim. Today, there is probably sufficient tacit coordination around a

Robert W. Bennett, *Democracy as Meaningful Conversation*, 14 CONST. COMMENT. 481, 486 n.14 (1997).

207. See Cawrey, *supra* note 115 and accompanying text.

particular software repository that a decision by developers to move gradually to a system that eliminates the need for such a repository would not create so much resistance as to destabilize the currency. Perhaps the greater obstacle, then, might be among the software developers themselves, who might prefer to control the main repository than to have peer-to-peer decision-making.

C. Rewarding Institution-Promoting Activities

The two examples detailed above illustrate that peer-to-peer decision-making can be used both for binary decisions and for decisions about whether to accept and change a particular text or code. Another type of decision is a quantity decision. Legal systems frequently make quantity decisions, for example when juries decide how much damages to award a plaintiff in a case in which the defendant has been found liable. Bitcoin or some other peer-to-peer institution, meanwhile, might wish to implement a more robust decentralized fisc. Bitcoin's mining mechanism, we have seen, provides rewards only for a particular activity that provides benefits to the Bitcoin community: mining.²⁰⁸ Most institutions, however, choose to spend money on a variety of purposes, so a mechanism for committing to spend money or reward activities undertaken on behalf of an institution could be central to some peer-to-peer decision-making institutions.

Bitcoin itself might benefit if rewards were available for other activities benefiting Bitcoin. For example, one might argue that Bitcoin should reward those who make significant contributions to the code base. Some claim that there are not enough volunteers interested in working on low-level aspects of the code.²⁰⁹ Monetary payment might be counterproductive by making individuals less likely to make altruistic contributions,²¹⁰ but for certain contributions with less inherent interest, monetary payment might be useful. Or, perhaps it makes sense to reward commitments to help stabilize the currency by buying at least a certain amount of the currency should its value on exchanges fall below a certain level. Or, perhaps businesses that enable Bitcoin payment or developers of services complementary to Bitcoin should receive some grant subsidy as a means of further extending Bitcoin. Bitcoin could establish policies allowing or prohibiting payments for different classes of contributions, and then, where permitted, use quantity decisions to determine the size of contributions.

The existing approaches to decision-making could easily be used to make quantity decisions. Numbers can be represented in binary form, so a group of binary decisions could be used to make a quantity decision. One binary decision could represent a single unit; another, two units; another, four units; and so forth, with decisions on the largest conceivable numbers of units being resolved first. The text decision-making approach could work as well. A text, after all, could

208. See *supra* Section I.C.

209. See Danny Bradbury, *Bitcoin Core Development Falling Behind, Warns Bitcoinj's Mike Hearn*, COINDESK (Feb. 24, 2014, 5:57 AM), <http://www.coindesk.com/bitcoin-core-development-falling-behind-warns-mike-hearn/>.

210. See, e.g., BENKLER, *supra* note 114, at 378 (explaining that monetary rewards may reduce contributions in peer production).

consist simply of a number, and user proposals to change the number would be assessed with the same decision-making approach described above for changes to the reference code.²¹¹ A number likely would become more refined over time, as preliminary decisions on an approximate level would not be revisited in fixing on a final value.

It would also be possible to design peer-to-peer decision-making processes geared specifically to decision-making about quantities. For example, given a need to reach agreement on a quantity such as a reward, one participant could propose a particular number by sending a proposal fee to an address in a transaction indicating in metadata the purpose of the payment. The metadata would also contain the participant's proposal of a certain number. Another participant might then propose a different number by paying the same amount or a higher amount to the same address, with metadata specifying the new number proposal, and subsequent participants could do the same. Each new proposal amounts to a bet with the prior proposing participant that the new participant's proposal will be closer to the final number than the prior participant's proposal. The size of this bet is the amount of the prior participant's bet, after deducting the amount that the prior participant had bet with the participant before that. The tacit coordination game is thus much like the earlier ones, with each participant considering what participants will decide in the future. The decision can be deemed final once a sufficient period elapses either with no proposals or with volatility in the most recent proposal below some predetermined threshold.²¹²

For example, suppose *A* proposes 15 with 1 Bitcoin, *B* proposes 30 with 2 Bitcoin, *C* proposes 20 with 3 Bitcoin, and *D* proposes 40 with 3 Bitcoin. If *D*'s transaction is the last one, then *B* would have won its 1 Bitcoin bet with *A*, *C* would have lost its 1 Bitcoin bet with *B*—i.e., the 2 Bitcoin that *B* invested minus the 1 Bitcoin of which corresponded to *B*'s investment with *A*—and *D* would have won its 2 Bitcoin bet with *C*—i.e., 3 Bitcoin minus the 1 Bitcoin that *C* bet *B*. The remaining investment, *D*'s extra 1 Bitcoin, would be refunded, as would the winning bets. There is, of course, some risk of simultaneous transactions (suppose *B*₁ and *B*₂ simultaneously challenge *A*), but the protocol could resolve which block should be considered authoritative in this case. For example, if a block contained multiple challenges to a particular transaction, then the block with the greatest challenge amount could control, and the other transaction would be void;²¹³ if the amounts were equal, then the block that appears earliest in the block chain would be authoritative.

211. See *supra* Section II.B.

212. Volatility might be measured, for example, as the standard deviation of the most recent proposal at the time each block is added to the block chain for the most recent 100 blocks. The threshold could be defined by the protocol.

213. A challenge to a void block would also be void. So, if *C*₁ challenged *B*₁, and *C*₂ challenged *B*₂, and *B*₁'s transaction were void pursuant to this rule, then the *C*₁ transaction would also be void. A void transaction would have no effect on the user's Bitcoin balance, though it would be included in the block chain.

It might seem that there is a flaw in this scheme, and indeed the flaw may exist to some extent in the earlier proposals as well: there may be no incentive for the first participant, *A* in the above example, to pay the proposal fee. Once *A* pays the proposal fee, a subsequent participant will challenge whenever it expects to be able to improve on the estimate more than some hypothetical subsequent participant could improve on its own estimate. But *A* will have no incentive, unless *A* has some intrinsic interest in the question at hand. If resolution of the quantity decision is important for peer-to-peer governance, then, it may make sense for the peer-to-peer institution to cover some portion of the proposal fee as a means of subsidizing the decision-making. Similarly, it may make sense to cover some portion of any increase in the amount at stake. It is not inherently obvious, however, how large any such reward should be. So, one could use peer-to-peer decision-making to determine the size of the reward on some other question being resolved by peer-to-peer decision-making.

The recursion inherent in this can be resolved by providing for some default reward proportion to be paid unless someone pays a proposal fee to initiate decision-making on some other reward proportion. For example, we might imagine a default reward proportion of 0. Suppose that *A* initiates a decision-making urging that a reward be paid to some Bitcoin owner on account of that Bitcoin owner's work promoting Bitcoin. *A* thus pays a proposal fee of, say, a mandated 1 Bitcoin. The Bitcoin owner might be *A* itself, or might not be. Either during or after the process of determining the reward to be paid, someone might propose some reward to *A* to offset some or all of the expense of the proposal fee. The proposer of this reward might be *A* as well, or might be someone else willing to pay the proposal fee, with a lower mandate of, say, 0.1 Bitcoin, because the stakes will be lower. In principle, someone could initiate yet another decision-making process to offset a portion of the 0.1 Bitcoin proposal fee by paying a proposal fee at some pre-established minimum level, though at some point, the level of recursion will be such that participants are likely to reject the proposal.

III. THE POSSIBILITIES AND PERILS OF PEER-TO-PEER GOVERNANCE

A cryptocurrency such as Bitcoin, Part I showed, can perform the central tasks of traditional institutions—maintaining a ledger, spending money, and making decisions—peer-to-peer, though in a constrained way. Incorporating formal tacit coordination games into a cryptocurrency, we saw in Part II, can allow for a more flexible decision-making apparatus. Nonetheless, if peer-to-peer decision-making were limited to cryptocurrencies, it would hold relatively little interest for legal theorists, other than perhaps those specialized in specific types of financial institutions or transactions. This Article has focused on cryptocurrencies, however, only because they are a central building block for any peer-to-peer institution. Naturally, we should not expect or want peer-to-peer decision-making to take over our central democratic institutions. But it is possible that peer-to-peer decision-making could assume niche responsibilities, most obviously in private law contexts, but perhaps in public law as well.

A. Peer-to-Peer Arbitration

Perhaps the most obvious application of peer-to-peer decision-making would be as a form of arbitration. Under the Federal Arbitration Act, parties can voluntarily by contract use private arbitrators to resolve their disputes, and the federal courts will honor those private resolutions.²¹⁴ The courts have interpreted the Act broadly, allowing, for example, arbitration provisions in contracts of adhesion to preclude class-action litigation.²¹⁵ At least one commentator considering the possibility of online arbitration has argued that online arbitration would be permissible.²¹⁶ The vision for such arbitration, however, is not of a peer-to-peer institution, but simply of arbitrators, chosen either by the parties or by the arbitration agency, using technology such as chat rooms or videoconference to lower some of the transaction costs associated with arbitration.

Peer-to-peer arbitration could represent a far greater departure from existing litigation and arbitration. First, peer-to-peer arbitration by definition would not require the selection of particular arbitrators. Second, such arbitration could avoid the need for legal enforcement of judgments (and the danger that the courts might refuse to honor peer-to-peer arbitration decisions; for example, on the theory that they violate due process rights)²¹⁷ if the arbitration is used simply to resolve disputes over funds placed in escrow. If possession is nine-tenths of the law,²¹⁸ then courts are unlikely to interfere with the outcome of a self-executing peer-to-peer arbitration. Third, peer-to-peer arbitration would not require formalized rules governing the presentation or consideration of evidence or arguments. Once a party initiates the decision-making process, participants would consider whatever evidence they considered relevant. The litigants would have some incentives to release information that benefited their respective cases and potentially even information that hurt their cases if that information was actually less harmful than decision-makers would think if release were refused.²¹⁹

214. See Federal Arbitration Act, 9 U.S.C. §§ 1–16 (2012); see also *supra* note 24 and accompanying text.

215. See *AT&T Mobility v. Concepcion*, 563 U.S. 333, 352 (2011) (voiding inconsistent state statutes).

216. See Frank A. Cona, *Application of Online Systems in Alternative Dispute Resolution*, 45 *BUFF. L. REV.* 975 (1997).

217. There is, of course, no clear original intent on this issue, and the courts' resolution of any questions about whether peer-to-peer decision-making offends due process is likely to be pragmatic. *Matthews v. Eldridge* recognizes that due process is highly context-specific and considers factors including the risk of error and the costs and burdens of procedures. Resolving a due process inquiry would thus likely depend in part on an empirical assessment, either rigorous or anecdotal, of the peer-to-peer decision-making. 424 U.S. 319, 343–47 (1976).

218. See Joseph William Singer, *Nine-Tenths of the Law: Title, Possession & Sacred Obligations*, 38 *CONN. L. REV.* 605, 605 (2006).

219. Economists have recognized that incentives to release information can be powerful when inferences will be drawn from refusal to release the information. See Paul Milgrom & John Roberts, *Relying on the Information of Interested Parties*, 17 *RAND J. ECON.* 18, 30–31 (1986) (“[R]ational skepticism by a decisionmaker can lead to a full-

It might seem that an absence of procedural rules would be a serious disadvantage of peer-to-peer arbitration. Some rules may be unnecessary or less necessary with peer-to-peer decision-making. Rules of jurisdiction²²⁰ and associated doctrines, such as venue²²¹ and forum non conveniens,²²² determine the court in which a lawsuit should be filed. In a peer-to-peer arbitration, there is no need to select a particular arbitrator or arbitration forum for peer-to-peer decision-making, because anyone may participate. Other rules, such as provisions allowing for hearings, help ensure that judges cannot shirk from the tasks of hearings and evidence. Peer-to-peer arbitration, by contrast, provides financial incentives for careful consideration.²²³ Still other rules, especially those that allow appeal, help prevent idiosyncratic decision-making by a single individual and ensure that the law is followed. Peer-to-peer decision-making involves multiple decision-makers, thus reducing the risk of idiosyncratic judgment, though perhaps exacerbating the risk that decision-makers might consider factors not strictly relevant from a legal standpoint.²²⁴

The ultimate question is the empirical and subjective one of whether peer-to-peer arbitration, whether procedure free or with a well-developed set of procedural rules (perhaps created by peer-to-peer decision-making itself, resolving issues such as time limits for a defendant to answer a plaintiff's complaint²²⁵), would be superior to a more traditional system of adjudication or arbitration. It is impossible to take a firm position on this. The experiment seems a worthwhile one, if the worth of an experiment is measured by the degree of uncertainty as to its outcome. Perhaps peer-to-peer arbitration would be cheaper than traditional arbitration, both because of saved transport costs and because of the expense

information decision by inducing one party to reveal information that is damaging to its interests. The party reveals this information for fear that withholding it will lead to an *even more unfavorable supposition* by the skeptical decisionmaker.") (emphasis in original).

220. See, e.g., U.S. CONST. art. III (limiting federal courts' subject matter jurisdiction); *Int'l Shoe Co. v. Washington*, 326 U.S. 310 (1945) (introducing the modern framework for personal jurisdiction).

221. See, e.g., 28 U.S.C. § 1391 (2012) (setting forth the federal venue rules).

222. See, e.g., *Piper Aircraft Co. v. Reyno*, 454 U.S. 235 (1981) (developing federal version of the doctrine).

223. Hearings and trials may also serve a psychological function, helping satisfy litigants' desire that someone consider their perspective. See, e.g., E. ALLAN LIND & TOM R. TYLER, *THE SOCIAL PSYCHOLOGY OF PROCEDURAL JUSTICE* (1988). How peer-to-peer decision-making would compare in imparting a sense of procedural justice is an empirical question. One might assume that peer-to-peer decision-making would be inferior because of a lack of in-person contact, but trial is so rare in civil adjudication that the benefits of in-person contact cannot be a primary benefit of the system. See generally Marc Galanter, *The Vanishing Trial: An Examination of Trials and Related Matters in Federal and State Courts*, 1 J. EMPIRICAL LEGAL STUD. 459 (2004) (discussing the increasing rarity of trials).

224. Even if peer-to-peer decision-makers believed that some evidence *should* be disregarded, they might, nonetheless, have trouble ignoring it. See Andrew J. Wistrich et al., *Can Judges Ignore Inadmissible Information? The Difficulty of Deliberately Disregarding*, 153 U. PA. L. REV. 1251, 1323 (2005) (showing that judges similarly have difficulty ignoring inadmissible evidence).

225. Cf. FED. R. CIV. P. 12(a)(1) (providing similar time limits).

associated with formal proceedings, but this is not guaranteed. Maybe peer-to-peer arbitration will lead to more predictable decisions, because no single person will control the outcome. It is also possible, though, that freedom from legal constraints will add randomness and arbitrariness.

The care that peer-to-peer decision-making participants take in their evaluation of evidence would depend partly on the protocol rules. The larger the peer-to-peer proposal fee,²²⁶ the greater the incentive that peer-to-peer decision-makers will have to educate themselves. Who should pay the fee and how large it should be is a question alien to public adjudication, where taxpayers subsidize the courts,²²⁷ but familiar in the arbitration context, because arbitrators must be paid.²²⁸ One might use peer-to-peer decision-making to set the size of the fee to be paid by the plaintiff. The size should depend on the marginal benefit of increased adjudication accuracy. Our litigation system contains only relatively crude mechanisms for adjusting the amount spent to judge cases based on the amount at stake, such as the existence of separate courts for small claims.²²⁹ Judges are likely to use their discretion to spend more time on more important matters,²³⁰ but no financial incentives drive this result. With peer-to-peer arbitration, one's incentives to invest in researching a case will be proportional to the probability that one will conclude that prior participants have not fully taken factors into consideration and to the proposal fee.

Peer-to-peer decision-making also could be used to affect litigants' investment incentives. For example, peer-to-peer decision-making similarly might be applied on a case-by-case basis to determine whether one side must reimburse the other for their legal fees or other expenses, based on factors such as whether the case was a close one.²³¹ Perhaps to avoid idiosyncratic decision-making on such issues, legal systems generally do not allow case-specific inquiries about fee shifting.²³² Moreover, our litigation system generally makes no attempt to limit

226. See *supra* Section II.A.2.

227. See Stephen J. Ware, *Is Adjudication a Public Good? "Overcrowded Courts" and the Private Sector Alternative of Arbitration*, 14 *CARDOZO J. CONFLICT RESOL.* 899, 919–20 (2013) (arguing that some parties should have to pay market rates for adjudication provision).

228. See Christopher R. Drahozal, *Arbitration Costs and Contingent Fee Contracts*, 59 *VAND. L. REV.* 729, 736–43 (2006) (discussing arbitration fees).

229. See generally Austin Sarat, *Alternatives in Dispute Processing: Litigation in a Small Claims Court*, 10 *LAW & SOC'Y REV.* 339 (1976) (describing the general differences between litigation in civil courts and litigation in small claims courts).

230. Marin K. Levy, *The Mechanics of Federal Appeals: Uniformity and Case Management in Circuit Courts*, 61 *DUKE L.J.* 315, 345 (2011) (describing how circuit court judges may hold fewer oral arguments to spend more time on difficult and complex cases).

231. See Lucian Arye Bebchuk & Howard F. Chang, *An Analysis of Fee Shifting Based on the Margin of Victory: On Frivolous Suits, Meritorious Suits, and the Role of Rule 11*, 25 *J. LEGAL STUD.* 371, 373, 382–85 (1996) (arguing for fee-shifting in nonclose cases).

232. An exception is Israel, where judges have discretion to assign costs. See Theodore Eisenberg et al., *When Courts Determine Fees in a System with a Loser Pays Norm: Fee Award Denials to Winning Plaintiffs and Defendants*, 60 *UCLA L. REV.* 1452 (2013).

parties' spending on developing reasonable legal arguments.²³³ Because each party will not take into account the effect of its spending on the welfare of its opponent, the result is likely to be excessive spending, relative to the amount that the parties ideally would spend *ex post* to make the contract efficient *ex ante*. Arbitrators or judges could be empowered to levy fines for excessive legal investment, but we would be hesitant to place such discretion in individuals, especially if there is a danger that they might use this power to shirk their own work.

Whether this would increase predictability of decisions and better control spending on arbitrations are empirical questions. The point here is simply that peer-to-peer arbitration is not just arbitration moved onto the Internet, but a different form of decision-making with strengths different from those of conventional arbitration and litigation.

B. A Peer-to-Peer Trust

If peer-to-peer arbitration can serve as a means of producing relatively predictable decisions relatively cheaply, then it could in turn serve as the foundation for a peer-to-peer trust.²³⁴ A settlor would establish the trust by a transaction that would send Bitcoin to an address from which it could not be spent by ordinary means.²³⁵ Metadata for the transaction would indicate the purpose of the trust and circumstances in which the trust could be disbursed.²³⁶ At any later point, someone could pay at least a minimum proposal fee, which could be established by the trust, to initiate a funding request for the trust. Requests for discretionary funds might require higher proposal fees than requests for nondiscretionary payments. This would initiate a peer-to-peer arbitration to resolve whether a payout should occur and, if so, the size of the payout that should be granted. The peer-to-peer arbitration could also consider whether any part of the proposal fee or other payments made during the arbitration process by participants should be refunded from the trust. New currency could then be issued in the

233. Mechanisms do exist for penalizing *frivolous* arguments. See, e.g., FED. R. CIV. P. 11. But these are used sparingly. See Neal H. Klausner, *The Dynamics of Rule 11: Preventing Frivolous Litigation by Demanding Professional Responsibility*, 61 N.Y.U. L. REV. 300, 311 (1986) ("It is a rare occasion, however, when the court invokes its inherent equitable power. This sanction has been reserved for cases in which a claim was made in subjective bad faith.").

234. Stephan Tual of the Ethereum project has used a trust as an example of peer-to-peer decentralization built on the block chain, but without tacit coordination decision-making. See London Futurists, *The Upcoming Decentralization Singularity*, YOUTUBE at 44:52 (Nov. 11, 2014), <http://www.youtube.com/watch?v=TNDHjmbC-t8>; see also *White Paper*, *supra* note 12 and accompanying text.

235. A preset address, perhaps resulting from the hash of a phrase such as "Trust Account for John Smith," could be used. As before it would be virtually impossible for someone to find a private key to unlock the funds in such an account. See *supra* Section II.A.2.

236. The metadata might, for example, be a hash of a document with further instructions. The settlor would have incentives to make this document available through conventional online means so that individuals would know the rules governing payout of the trust.

amount specified and awarded to the public address of the party applying for a payment.

The Bitcoin protocol would need to recognize that when an adjudication concluded with a decision to make a payout, the payout should result in the creation of new currency in the specified amount. In theory, one could bake into the protocol itself a rule that total payouts from a trust cannot exceed the amount paid into the trust, but it is also possible for the protocol simply to allow peer-to-peer decisions to create new currency. This would thus delegate to the peer-to-peer decision-makers the task of ensuring that excessive payments are not made. If peer-to-peer decision-making can be used to pay out arbitrary rewards for those who help promote a currency,²³⁷ then it should also be possible to use such decision-making to pay out arbitrary amounts for other purposes. If the general mechanism of formal tacit coordination games works for subjective decision-making, it should work as well for questions with objective answers (such as whether a trust has sufficient funds remaining to support a request), particularly because participants could program computer-based agents to participate in decision-making to support objective rules. This should be profitable if one anticipates that the focal point solution will be to enforce such rules. In principle, peer-to-peer decision-making also could be used to undo mistakes of peer-to-peer decision-making, but objective errors seem highly unlikely in any event.

The potential efficiency benefit from a peer-to-peer trust is that it might lower transaction costs. In 1984, John Langbein argued that the high transaction costs associated with the probate system had led to an increased reliance on techniques for transferring assets without resort to probate.²³⁸ For example, life insurance proceeds and pension accounts name specific beneficiaries who can receive the relevant funds without direct legal intervention.²³⁹ Probate continues to perform a critical function in clearing title for real property, but there are means sometimes to evade even this, and personal property is often distributed without judicial intervention.²⁴⁰ Meanwhile, secured lending allows creditors to resolve most loans without probate.²⁴¹ The reason that all of these mechanisms are preferred to probate is that “[t]he probate system,” Langbein explained, “has earned a lamentable reputation for expense, delay, clumsiness, makework, and worse.”²⁴²

Conventional nonprobate transfers, however, are also not without transaction costs. Daniel Kelly notes that the combination of a will and a revocable trust will generally involve greater *ex ante* transaction costs than creation of a will

237. See *supra* Section II.C.

238. John H. Langbein, *The Nonprobate Revolution and the Future of the Law of Succession*, 97 HARV. L. REV. 1108, 1116–17 (1984).

239. *Id.* at 1110–11.

240. *Id.* at 1117–19.

241. *Id.* at 1123.

242. *Id.* at 1116.

alone.²⁴³ “Moreover,” he argues, “a settlor who creates a trust may have to perform additional tasks like transferring assets into the trust or changing beneficiary designations.”²⁴⁴ Transaction costs are likely to be especially large when the grantor wishes to impose subjective conditions on distribution of trust funds. For example, a common trust provision allows for funds to be used only for purposes such as education or health.²⁴⁵ A trustee will then be needed to determine whether particular claims for payment should be honored. In the case of bad faith or serious abuse, the courts can remove a trustee.²⁴⁶ Even short of that, refusals to make payouts can lead to lawsuits.²⁴⁷ A lawsuit may demand correction of an overpayment,²⁴⁸ and so trustees historically have been conservative in authorizing payments.²⁴⁹ A trust can grant a trustee “absolute” or “uncontrolled” discretion, but the *Restatement (Third) of Trusts* provides that these words “are not interpreted literally.”²⁵⁰ “It is a matter of interpretation,” the *Restatement* explains, “to determine the degree to which a settlor’s use of language of extended discretion (e.g., ‘absolute discretion’) manifests an intention to modify the normal duties of the trustee and the normal grounds of judicial intervention in the exercise of a discretionary power.”²⁵¹ A settlor thus cannot definitively avoid judicial interference and its attendant costs.

There are thus at least two possible benefits to a peer-to-peer trust that could lead settlors to prefer such a mechanism to either probate or a conventional nonprobate trust. First, creation costs could be quite low, because few or no formalities would be required. The only requirement would be making a cryptocurrency payment with sufficient metadata so that the purpose of the payment could be ascertained. Second, the peer-to-peer trust would rely on peer-to-peer decision-making, which might be cheaper than a conventional trustee. A conventional trustee will have to charge enough money to cover the risks associated with being a trustee, including the possibility that the trustee will be found to have acted in bad faith and required to replenish the trust.²⁵² The peer-to-

243. Daniel B. Kelly, *Toward Economic Analysis of the Uniform Probate Code*, 45 U. MICH. J.L. REFORM. 855, 875–76 (2012).

244. *Id.*

245. Peter B. Tiernan, *Drafting Trusts That Include Broad Invasion Powers*, 77 FLA. B.J. 74, 74 (2003).

246. *See Gould v. Starr*, 558 S.W.2d 755 (Mo. Ct. App. 1977).

247. *See In re NWFEX, Inc.*, 267 B.R. 118, 155 (Bankr. W.D. Ark. 2001) (holding that trustee had to return trustee fees to estate where trustee did not provide shareholders with payouts).

248. *See In re Murray*, 45 A.2d 636 (Me. 1946) (holding that trustees may have to repay trust money that was not properly paid out).

249. *See Edward C. Halbach, Jr., Problems of Discretion in Discretionary Trusts*, 61 COLUM. L. REV. 1425, 1427 (1961).

250. RESTATEMENT (THIRD) OF TRUSTS § 87 cmt. d (AM. LAW INST. 2003); *see also* RESTATEMENT (SECOND) OF TRUSTS § 187 cmt. j (AM. LAW INST. 1959).

251. RESTATEMENT (THIRD) OF TRUSTS § 87 cmt. d (AM. LAW INST. 2003).

252. *See generally* Stewart E. Sterk, *Rethinking Trust Law Reform: How Prudent Is Modern Prudent Investor Doctrine?*, 95 CORNELL L. REV. 851 (2010) (describing overall risks to trustee investments).

peer trust could allow a settlor to prevent judicial interference with the trust's decision-making. A settlor might wish to do this if the settlor is sufficiently confident in the peer-to-peer decision-making process. This is not for everyone. Judicial trust supervision provides benefits. But just as nonprobate transfers have allowed an end run around perceived inefficiencies of probate law, so too could cryptocurrency trusts allow an end run around perceived inefficiencies of conventional trusts. A cryptocurrency trust thus serves a niche for those who believe that they face high transaction costs with conventional trusts.

C. A Peer-to-Peer Bank

The peer-to-peer trust, as described so far, lacks one common feature: the ability to invest trust funds. The trust money is set aside until the money is needed, so the investment is ultimately in the cryptocurrency itself, rather than in a diversified form. Ideally, it would be beneficial for the trustee to be able to invest deposited Bitcoins pending trust withdrawals to grow the trust corpus. This is, of course, possible with conventional trust relationships. The trustee simply relies on a financial institution such as a bank or mutual fund, deposits the trust moneys and then withdraws them as needed. Peer-to-peer decision-making could support mechanisms for *deciding* when cryptocurrency should be exchanged for other assets controlled by a bank. The challenge for a cryptocurrency is how to execute that exchange. The problem, however, is that there is no mechanism allowing cryptocurrency accounts to own virtual assets. For a peer-to-peer institution to own assets besides virtual currency, some interface is needed between the virtual and real worlds.

A cryptocurrency bank can establish this connection. The bank would serve the role of a trusted intermediary. Potential depositors would need to decide whether to trust any bank based on its track record and any assurances it might provide with regard to its security and financial practices. Early experiments with Bitcoin banks have not inspired confidence, with at least two major bank failures from apparent failures to safeguard Bitcoins.²⁵³ But it seems plausible that a bank might establish a reputation over time. Even a wholly anonymous bank might inspire trust so long as the present discounted value of expected bank profits is greater than the benefit to the bank of stealing deposits. It is more plausible, though, that trust could be achieved through transparency and identification of the

253. The largest of these was the failure of Mt. Gox. Robert McMillan, *The Inside Story of Mt. Gox, Bitcoin's \$460 Million Disaster*, WIRED (Mar. 3, 2014, 6:30 AM), <http://www.wired.com/2014/03/bitcoin-exchange/> (describing how poor management practices led to the eventual closure of the Bitcoin exchange at Mt. Gox in Tokyo, Japan, which resulted in a loss of over 800,000 Bitcoins with an estimated worth of about \$460 million). A smaller bank failure was that of Flexcoin. See Alex Hern, *Bitcoin Bank Flexcoin Closes After Hack Attack*, GUARDIAN (Mar. 4, 2014, 7:33 AM), <http://www.theguardian.com/technology/2014/mar/04/bitcoin-bank-flexcoin-closes-after-hack-attack>. In each of these cases, the security flaw was not in the central cryptographic mechanism, but in the wallet services themselves. Presumably, someone will eventually develop wallet software that does not have vulnerabilities that allow people without the requisite private keys to withdraw other people's money.

bank owners, so that they might face criminal liability should they steal money and at least reputational sanctions should they fail to safeguard it.

A cryptocurrency bank operating in this way is not a peer-to-peer bank. A virtual currency operated by a trusted intermediary is not a peer-to-peer institution; indeed, the purpose of Bitcoin was to offer a peer-to-peer alternative to the trusted intermediary approach. A bank that serves as a trusted intermediary will maintain its own centralized records and management. Deposits into and withdrawals from the bank might be conducted entirely by Bitcoin and thus appear on the block chain, but if the bank itself is chartered in some jurisdiction, then it is not peer-to-peer. Nonetheless, it is worthwhile to consider the role that cryptocurrency banks might have in supporting peer-to-peer institutions, along with the danger that such banks might support criminal activity, before considering the possibility of a true peer-to-peer bank.

A cryptocurrency bank, in principle, could hold accounts in the name of cryptocurrency public keys. For example, a peer-to-peer decision might be to place a trust corpus into a particular cryptocurrency bank. The bank would have released a public key corresponding to an account it controls, and the peer-to-peer decision-making process could result in new currency then being assigned to this public key to offset the funds placed into trust. The owner of the private key (the bank) could do what it wished with this currency, including swapping the cryptocurrency for ordinary currency via an exchange.²⁵⁴ It thus would be able to place funds into traditional investments. Peer-to-peer decision-making could result in a withdrawal decision, and the bank would then be expected to send cryptocurrency back to the trust. Presumably, a failure to do so would mean that peer-to-peer decision-makers would not use that bank in the future.

Just as peer-to-peer arbitration or a peer-to-peer trust could offer lower transaction fees than traditional equivalents, so too might a cryptocurrency bank reduce transaction fees. But the principal reason for this is that a cryptocurrency bank might more easily escape regulation. If it becomes easy for individuals or organizations to move their funds to cryptocurrency, and they can anonymously move cryptocurrency to bank accounts, they may be able to opt out of bank regulation. One motivation for this is that bank regulation is expensive. Theorists justify the expense on the grounds that it benefits depositors²⁵⁵ and contributes to macroeconomic stability,²⁵⁶ though some individual depositors might prefer banks

254. Many Bitcoin exchanges already exist. *See, e.g.*, BITSTAMP, <http://www.bitstamp.net> (last visited Feb. 29, 2016).

255. *See, e.g.*, James R. Barth et al., *Bank Regulation and Supervision: What Works Best?*, 13 J. FIN. INTERMEDIATION 205, 210–11 (2004) (describing how capital adequacy requirements may align the interests of bank owners and depositors).

256. *See generally* Olivier Blanchard et al., *Rethinking Macroeconomic Policy*, 42 J. MONEY CREDIT & BANKING 199 (2010) (describing the general relationship between bank regulation and macroeconomic policy).

with less regulation. Another motivation for cryptocurrency banking would be to facilitate crimes such as tax evasion or money laundering.²⁵⁷

An anonymous cryptocurrency bank likely could not easily be regulated, assuming the cryptocurrency has sufficient privacy protections, because there would be no way of identifying the owners of the bank. But an anonymous bank will have a harder time drawing in cryptocurrency customers. Banks that seek to bolster their credibility by being public can be regulated in the countries in which they are located. It only takes one country to create an offshore banking haven. Such a jurisdiction would likely want to regulate these banks somewhat (to assure depositors), but might offer minimal regulation and maximal privacy protection. Cryptocurrency makes it more difficult to pressure a jurisdiction into cooperating with international transparency laws designed to deter money laundering. Existing financial regulation can target offshore banking by the indirect means of regulating transfers between offshore banks and ordinary banks.²⁵⁸ One could, however, transfer cryptocurrency directly to a cryptocurrency bank in such a haven.

Countries such as the United States could attack cryptocurrency banks in one of two ways. First, they might put pressure on the haven jurisdiction. Second, they might seek to regulate transactions in which individuals purchase cryptocurrencies, demanding disclosure of their identities, and then seeking to regulate those individuals. Authorities might, for example, regulate cryptocurrency ATMs, which in principle can make it easy to exchange cash and cryptocurrency anonymously. But it may be more difficult to regulate black markets.²⁵⁹ As long as individuals can buy and sell cryptocurrency with fiat currency, cryptocurrency banks will be difficult to regulate.

It may seem that our focus so far on cryptocurrency banks that are trusted intermediaries rather than truly peer-to-peer undermines the argument that it is possible to imagine robust peer-to-peer institutions. If one must rely on a trusted intermediary model to create a cryptocurrency bank, then perhaps true peer-to-peer institutions are impossible. The obstacle, however, is solely a legal one: a fully functional bank must be able to own real assets because a primary function of a bank is to invest funds. A peer-to-peer institution could own assets only if the legal system recognized the peer-to-peer institution as legitimately existing and having a form of personhood sufficient for the ownership of property. Real property purchased by a trust, for example, might be held in the name of the public key or in the name of the cryptocurrency as a whole.

Recognition of such ownership may seem unlikely because of concerns that cryptocurrency ownership might facilitate illegal activity by providing

257. See generally Omri Marian, *Are Cryptocurrencies Super Tax Havens?*, 112 MICH. L. REV. FIRST IMPRESSIONS 38 (2013).

258. See William F. Wechsler, *Follow the Money*, 80 FOREIGN AFF. 40, 52 (2001) (discussing difficulties in detecting wire transfers between domestic and off-shore bank accounts).

259. See Jon Matonis, *Government Ban on Bitcoin Would Fail Miserably*, FORBES (Jan. 28, 2013, 9:39 AM), <http://www.forbes.com/sites/jonmatonis/2013/01/28/government-ban-on-bitcoin-would-fail-miserably>.

anonymity. But a refusal to allow cryptocurrency ownership because of discomfort with cryptocurrencies would be self-defeating. Trusted intermediaries would still exist in other jurisdictions, and so cryptocurrencies would remain helpful for money laundering. Meanwhile, a refusal to allow ownership would impede legitimate cryptocurrency transactions and reduce the government's ability to regulate peer-to-peer banks or other peer-to-peer institutions. If a peer-to-peer bank or customer fails to follow applicable regulations, then the legal system could seize assets owned by the peer-to-peer institution. The legal system would need to develop principles for regulating such property seizure. For example, the legal system would need to assess when property ownership made a particular cryptocurrency account or cryptocurrency amenable to jurisdiction, addressing such timeless questions as whether the exercise of jurisdiction could be predicated solely on the basis of property ownership.²⁶⁰ However those questions are resolved, the legal system can regulate peer-to-peer institutions to the extent it permits them to own assets within its jurisdiction, but can do little about offshore trusted intermediaries that promote money laundering.

D. A Peer-to-Peer Business Association

A peer-to-peer bank is a specific realization of the more general concept of a peer-to-peer business association. The peer-to-peer bank accepts funds, makes investment decisions, and approves expenditures, which are the general functions of any business association. We can thus imagine peer-to-peer decision-making being used to operate a peer-to-peer business association. The business might raise funds by soliciting contributions in Bitcoin or another cryptocurrency, make investment decisions, and ultimately pay dividends or liquidation funds to the investors. The business association might sue and be sued. A peer-to-peer business association would not be a sole proprietorship, partnership, limited liability company, or corporation, at least as traditionally conceived. The traditional forms of business association differ in how they allocate ownership interests and decision-making authority, but the peer-to-peer business association allocates decision-making authority in a new way—not to a specific owner, to partners, to a board, or even to shareholders, but to the peer-to-peer decision-makers as a whole.

Whether peer-to-peer business associations fill a niche depends on whether there are situations in which such associations minimize the sum of agency costs.²⁶¹ The agents of the business association would be the peer-to-peer decision-makers who voluntarily participate in the tacit coordination game for profit motive. Such decision-makers might have less self-interest than managers or

260. See, e.g., *Pennoyer v. Neff*, 95 U.S. 714, 725, 728 (1878) (finding that quasi in rem jurisdiction would be property if Oregon property belonging to the defendant had been attached prior to the lawsuit); *Shaffer v. Heitner*, 433 U.S. 186, 186–87 (1977) (holding that the exercise of quasi in rem jurisdiction is permissible only if it meets the requirements of the minimum contacts test of *International Shoe Co. v. Washington*, 326 U.S. 310, 317 (1945)).

261. See generally Michael C. Jensen & William H. Meckling, *Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure*, 3 J. FIN. ECON. 305 (1976) (providing the seminal analysis of agency costs).

directors of a corporation, because the peer-to-peer decision-makers would not receive a salary from the entity. This would help reduce agency costs. Of course, some individual decision-makers might have some interest, for example in a contract that the peer-to-peer business association might undertake, but to avoid losing money in the peer-to-peer decision-making process, they would need to persuade others about the relevant corporate decision. Meanwhile, such decision-makers might have more information than shareholders, who often have little incentive to become informed about corporate affairs.²⁶² How much incentive they have to acquire information—and whether they would have as much information as managers or directors—depends on the size of proposal fees and thus the subsidy for decision-making.²⁶³

A limitation of peer-to-peer business associations is that their decisions would be inherently transparent. But, there may be some industries in which transparent decision-making would furnish an advantage. For example, such decision-making might reassure potential contractual partners that they are not being taken advantage of. To the extent that secrecy in business affairs is needed, however, peer-to-peer decision-makers could decide to hire employees, including executive managers, and allocate decision-making power to these managers, including the power to maintain information in confidence. The only decisions that *must* remain secret are the decisions by the peer-to-peer decision-makers themselves. Of course, to the extent that peer-to-peer decision-making controls only who the managers or directors are, the benefits and costs of peer-to-peer decision-making would be reduced.

If the legal system were to recognize peer-to-peer business associations, a doctrinal question would be whether such business associations are entitled to limited liability.²⁶⁴ If granted limited liability, a peer-to-peer business association would still face seizure of property that it owned, but the owners of the business association would not face additional liability as a result of the business association's actions. Given the ease with which business associations today can obtain limited liability, there appears to be little reason to resist limited liability for peer-to-peer business associations other than resistance to the inherent idea of peer-to-peer business associations. As a practical matter, defeating limited liability might be quite difficult because governmental authorities would not be able to identify investors.

Once again, it might seem that peer-to-peer business associations are fanciful. But it would only take one jurisdiction to recognize such business associations for them to be able to contract business in multiple jurisdictions. Just

262. Jennifer Arlen & Eric Talley, *Unregulable Defenses and the Perils of Shareholder Choice*, 152 U. PA. L. REV. 577, 580 (2003).

263. See *supra* notes 227–30 and accompanying text.

264. Issues about the extent of limited liability similarly arose with the rise of limited liability companies. See, e.g., Steven C. Bahls, *Application of Corporate Common Law Doctrines to Limited Liability Companies*, 55 MONT. L. REV. 43 (1994).

as Delaware seeks to attract corporations to receive franchise tax revenue,²⁶⁵ so too could Delaware allow for the registration of peer-to-peer business associations in exchange for payment of specified fees. If Delaware were uninterested in this business, another jurisdiction (such as Nevada, which recently tried to compete with Delaware for corporate charter business²⁶⁶) might do so. The jurisdiction might even call the peer-to-peer business association a “corporation.” Under current law, states may not discriminate against businesses incorporated in other states.²⁶⁷ Once registered or incorporated, such a business might be able to operate in other states in much the same way as other businesses.

E. Peer-to-Peer Public Law

Our examples of peer-to-peer decision-making have focused on private law for good reason. There are significant obstacles to private law peer-to-peer institutions, even placing aside the need for extension of Bitcoin or other cryptocurrency. One is the possibility that governmental hostility could prevent peer-to-peer institutions from owning real assets or that governments might directly regulate or prohibit individuals from using vehicles such as peer-to-peer trusts. Perhaps one or more governments can be persuaded to tolerate such peer-to-peer institutions. But it would be quite another large step before a government would embrace peer-to-peer decision-making for public purposes. The strongest reason to do so would be if existing governmental institutions were broken, for example because of corruption. The most likely area in which a peer-to-peer public institution might be created is central banking, though this too currently seems speculative.

1. A Peer-to-Peer Central Bank

A peer-to-peer central bank is the most obvious public institution that might be built on a cryptocurrency because a cryptocurrency essentially performs the function of a central bank. A country could adopt Bitcoin, or some new cryptocurrency, as fiat currency. Bitcoin transactions are electronic, and so Bitcoin is an imperfect substitute for cash.²⁶⁸ But mobile phones are becoming ubiquitous even in the developing world.²⁶⁹ So someday, a country *could* adopt a cryptocurrency as its fiat currency. The most obvious impetus to doing so would

265. See, e.g., Larry E. Ribstein, *Delaware, Lawyers, and Contractual Choice of Law*, 19 DEL. J. CORP. L. 999, 1012 (1994).

266. See Jill E. Fisch, *The Peculiar Role of the Delaware Courts in the Competition for Corporate Charters*, 68 U. CIN. L. REV. 1061, 1067–68 (2000).

267. See *Paul v. Virginia*, 75 U.S. (8 Wall.) 168, 180–81 (1869).

268. But see Nermin Hajdarbegovic, *10 Physical Bitcoins: The Good, the Bad and the Ugly*, COINDESK (Sept. 14, 2014, 7:15 PM), <http://www.coindesk.com/10-physical-bitcoins-good-bad-ugly/> (reporting on attempts to creating physical Bitcoins, for example with embedded private keys revelation of which will destroy the coin).

269. See generally *Emerging Nations Embrace Internet, Mobile Technology: Cell Phones Nearly Ubiquitous in Many Countries*, PEW RES. CTR. (Feb. 13, 2014), <http://www.pewglobal.org/2014/02/13/emerging-nations-embrace-internet-mobile-technology/>.

be a perception that the existing fiat currency has failed. This could occur if counterfeiting becomes widespread,²⁷⁰ but based on history, the more likely scenario is that the government has been unable to control inflation.

The macroeconomics literature has highlighted that it often will make sense for a central bank to seek to “tie its hands” to prevent it from engineering inflation surprises in the future.²⁷¹ The insight of this literature is that inflation can be a self-fulfilling prophecy, with future inflation depending not only on future central bank actions but also on current (and future) *expectations* of inflation. And so, if a central bank is in the habit of helping the government meet its bills and inflate away its debts by printing currency (or other mechanisms of expansive monetary policy), the public will anticipate that the central bank will continue to do so. The government may thus respond by hiring a new central banker with a reputation for inflation intolerance, who is more conservative about inflation than the government itself.²⁷² A more radical step is for a government to abandon its own currency and simply use a foreign currency, such as the U.S. dollar.²⁷³ This can fix the inflation-expectation problems, but it comes at a price: monetary policy can no longer be used to address business cycle fluctuations.

A country that merely adopted a fork of the current Bitcoin project would be solving its problem in a similar way. The growth rate of the new coins could be specified in advance, creating stable inflation expectations. This could be better than adopting a foreign currency, particularly if the country’s macroeconomic conditions are not likely to be correlated with those of the country whose currency otherwise would be adopted. Milton Friedman argued that a constant growth rate rule may be superior to an activist central bank with a country’s own interests in mind;²⁷⁴ a cryptocurrency, like Bitcoin, can insist on constant growth of the monetary supply. But for those who believe that a central bank can exercise discretion responsibly,²⁷⁵ adopting a cryptocurrency with a mechanical mining schedule would be harmful.

One economist, George Selgin, has considered the possibility that a central bank could adopt a cryptocurrency as a fiat currency.²⁷⁶ Recognizing the limitations of a constant growth rate rule, Selgin suggests that the currency might be “based upon a production ‘protocol’ such as might replicate the outcome of

270. See COMM. ON TECHS. TO DETER CURRENCY COUNTERFEITING ET AL., A PATH TO THE NEXT GENERATION OF U.S. BANKNOTES: KEEPING THEM REAL (2007) (describing anti-counterfeiting efforts).

271. See Finn E. Kydland & Edward C. Prescott, *Rules Rather than Discretion: The Inconsistency of Optimal Plans*, 85 J. POL. ECON. 473, 477–80 (1977).

272. See, e.g., Kenneth Rogoff, *The Optimal Degree of Commitment to an Intermediate Monetary Target*, 100 Q.J. ECON. 1169 (1985) (explaining that appointing conservative central bankers reduces expected inflation).

273. See Alberto Alesina & Robert J. Barro, *Dollarization*, 91 AM. ECON. REV. 381, 384 (2011) (discussing the adoption of the U.S. dollar by El Salvador).

274. See Friedman, *supra* note 204.

275. See, e.g., Sumner, *supra* note 205.

276. George Selgin, *Synthetic Commodity Money*, CATO INST. (Apr. 10, 2013), <http://object.cato.org/sites/cato.org/files/articles/synthetic-commodity-money.pdf>.

almost any conceivable monetary rule.”²⁷⁷ For example, he refers to Scott Sumner’s proposal for central banks to target nominal GDP, growing the currency just enough to keep nominal GDP growth rates constant.²⁷⁸ But Selgin does not explain precisely how this would work. The client software would need to be programmed with nominal GDP levels as an input. But there could be disputes about just what the nominal GDP levels are, and a government desiring to engineer inflation might prefer for the official nominal GDP levels to be artificially low so that more currency would be produced. If the government controls a central repository for the client software, it would be able to do this easily. The government could also at any point change the rule, abandoning nominal GDP targeting, by changing the client software. Anticipation of this would harm inflation expectations.

A fiat cryptocurrency with a built-in peer-to-peer decision-making apparatus can allow for monetary policy tailored to a particular country’s needs. The cryptocurrency could be targeted at a variable like nominal GDP growth, with the cryptocurrency itself used to change the reference software to incorporate nominal GDP growth figures.²⁷⁹ Or, the currency production schedule could be specified numerically, with peer-to-peer decision-making used to make changes and thus to accomplish either expansionary or contractionary monetary policy. To avoid large rents for miners, a proof-of-stake approach might be used in lieu of proof of work.²⁸⁰ There is always the danger that the government will abandon the currency for some other fiat currency.²⁸¹ But changing currencies is more destabilizing than interfering with an existing currency, and if the existing currency has proven relatively successful, the short-term economic costs from changing currencies are likely to exceed the short-term benefits of being able to create inflation.

2. Other Public Institutions

A government might be willing to replace a public institution with a peer-to-peer decision-making alternative only if several conditions are met. First, the institution must be one that seems clearly to be failing in achieving its core goals. Second, the peer-to-peer alternative must be seen as able to achieve the core goals of the institution. Third, the lack of direct governmental control over the peer-to-peer institution must be viewed as beneficial. Fourth, it must be difficult for the government to interfere with the peer-to-peer institution, once it is established. Central banking plausibly could meet all of these conditions in a country with a history of failed monetary policy, particularly because the central function of a cryptocurrency is so close to that of a central bank. Though it seems far less likely for other public institutions, perhaps it could become more plausible if peer-to-peer

277. *Id.* at 23.

278. *Id.* (citing Scott Sumner, *Re-Targeting the Fed*, NAT’L AFF., Fall 2001, at 79).

279. *See supra* Section I.B.

280. *See King & Nadal, supra* note 121 and accompanying text.

281. Selgin, *supra* note 276, at 24 (noting that a government retains “its ability to introduce and to confer legal tender status upon some new fiat currency”).

decision-making became familiar in private-law contexts and successful for central banking.

The obstacle to public institutions using peer-to-peer governance is not merely a practical one. Rather, it is a philosophical concern about the need for legitimacy of governmental authority. What creates the conditions for legitimacy is contested in political philosophy literature. A tradition traceable to John Locke emphasizes the significance of consent.²⁸² Such approaches might tolerate peer-to-peer decision-making, so long as the governed can be seen as consenting to it in particular contexts. Others emphasize the significance of representation,²⁸³ and it is difficult to make a case that peer-to-peer decision-making supports democratic participation or representation. A more recent theory, advanced by David Estlund, argues that democratic authority is based on epistemic proceduralism, which recognizes the tendency of democratic procedures to make correct decisions.²⁸⁴ This might seem to have greater potential to serve as a philosophical justification for peer-to-peer public institutions. Yet Estlund seeks to justify democratic institutions, even conceding voter ignorance and other weaknesses, as well as the possibility that there may be alternative approaches more likely to produce right answers.²⁸⁵ If peer-to-peer institutions turn out to produce right answers more effectively than alternatives, an answer to Estlund will still be necessary before peer-to-peer institutions can be considered legitimate.²⁸⁶

CONCLUSION

Although there is a long history of debate about the degree to which government should be centralized,²⁸⁷ the legal literature has not previously questioned the premise that every legally authoritative action must come from some institution that is centralized. Even advocates of direct democracy imagine some centralized system for counting votes in such elections. Peer-to-peer systems lack a centralized server for recording information, but, as Bitcoin has shown, peer-to-peer systems can still produce decisions about which there will be a high degree of consensus. The very limited form of decision-making inherent in Bitcoin could serve as a foundation for more sophisticated types of decision-making, allowing legal institutions to be created without voting or the designation of a central authority. The strongest case for application of such decision-making is for governance of Bitcoin itself because the current governance arrangement means

282. See JOHN LOCKE, *THE SECOND TREATISE OF GOVERNMENT* 54–55 (Thomas P. Peardon ed., 12th ed. 1952); see also John Dunn, *Consent in the Political Theory of John Locke*, 10 *HIST. J.* 153 (1967).

283. See M. Stephen Weatherford, *Mapping the Ties That Bind: Legitimacy, Representation and Alienation*, 44 *W. POL. Q.* 251, 259–63 (1991).

284. DAVID M. ESTLUND, *DEMOCRATIC AUTHORITY: A PHILOSOPHICAL FRAMEWORK* 8 (2009).

285. *Id.* at 258–70.

286. For a response to Estlund's argument, see Kristoffer Ahlstrom-Vij, *Democracy Without Voting* (2014) (unpublished manuscript) (on file with author).

287. The most famous work in the genre is *THE FEDERALIST PAPERS* (Clinton Rossiter ed., 1981).

that Bitcoin is, in important respects, not peer-to-peer. Bitcoin could experiment with such governance by allowing decisions to be used merely as advice about whether software features should be implemented. Peer-to-peer law is likely to emerge slowly and in unpredictable ways, but it has the potential to create authoritative decisions without authoritative decision-makers. There may be decisive arguments against particular peer-to-peer institutions, but legal theorists should at least allow peer-to-peer institutions to join the menu of possible regulatory arrangements.