2020

Transactional Scripts in Contract Stacks

Shaanan Cohny Hoffman, David A.

Follow this and additional works at: <https://scholarship.law.umn.edu/mlr>

Part of the [Law Commons](#)

Recommended Citation

Cohny, Shaanan Hoffman, David A., "Transactional Scripts in Contract Stacks" (2020). *Minnesota Law Review*. 3226.

<https://scholarship.law.umn.edu/mlr/3226>

This Article is brought to you for free and open access by the University of Minnesota Law School. It has been accepted for inclusion in Minnesota Law Review collection by an authorized administrator of the Scholarship Repository. For more information, please contact lenzx009@umn.edu.

Article

Transactional Scripts in Contract Stacks

Shaanan Cohn[†] and David A. Hoffman^{††}

Introduction	320
I. Designing Expensive, Buggy Scripts	328
A. How (Commercial) Coding Works	328
B. An Introduction to the Blockchain Platform	331
C. Ethereum and Scripting	335
D. Coding on Ethereum	341
E. Summary	346
II. Transactional Scripts in the Real World	348
A. Tokens	348
B. Exchanges	351
C. Oracles	354
III. Scripts and Stacks	358
A. The Canonical Stack	362
B. Tensions Within the Stack	368
C. Recapitulating the Canons: <i>Quoine</i> and Non-Public Scripts	385
IV. The Future of the Contract Stack	386

[†] Ph.D., Postdoctoral Research Associate, Center for Information Technology Policy, Princeton University. Copyright © 2020 by Shaanan Cohn.

^{††} Professor of Law, University of Pennsylvania Law School. We thank Alexander Altieri and Chrissy Pak for research assistance, and Yonathan Arbel, Bridget Fahey, James Grimmelmann, Greg Klass, Raina Haque, Bob Hillman, Drew Hinkes, Cathy Hwang, Gabe Kaptchuk, Max Raskin, Elizabeth Pollman, Gabe Shapiro, Jeremy Sklaroff, Tim Swanson, Andrea Tosato, Alec Webley, Kevin Werbach, Tal Zarsky, Eyal Zamir, and participants at workshops at Alabama, Penn, and Utah for comments. This work was supported in part by grants from the Ripple Research Fund at the Wharton School and at the Princeton Center for IT Policy. Copyright © 2020 by David A. Hoffman.

INTRODUCTION

In early 2019, a group of people founded Edgeware, a blockchain-based platform designed to host software development.¹ Edgeware made potential users a deal: if they agreed to temporarily sequester some cryptocurrency (essentially, placing an initial investment in escrow), they'd gain governance rights over the platform at a later date.² The mechanism for that investment was a piece of carefully-audited software, called "Lockdrop," which was deployed on a blockchain.³ Lockdrop seemed to embody and constitute a novel contracting technology, written in a programming language (Solidity) that didn't even exist a decade ago.⁴

By July of 2019, investors committed nearly \$300,000,000 to the project using Lockdrop.⁵ Then someone looked with particular care at the following piece of code:⁶

```
assert(address(lockAddr).balance == msg.value);
```

As it turns out, this line was susceptible to a software hack that would have permanently impounded investor assets.⁷ Luckily, the error was discovered before it caused harm. But, as the coder who discovered the vulnerability put it, "smart contracts are software. Even carefully audited, well tested software will (almost always) contain bugs. Therefore, and despite our best efforts . . . Smart contracts will

1. COMMONWEALTH LABS, EDGEWARE: AN ADAPTIVE SMART-CONTRACT BLOCKCHAIN 1 (2019) [hereinafter EDGEWARE WHITEPAPER], <https://arena-attachments.s3.amazonaws.com/3850782/1928e31873075de95992d4987eb14a2e.pdf?1552432633> [<https://perma.cc/J3GP-ERC4>].

2. See *id.* at 5 ("EDG entitles holders both staking and voting rights . . ."). For an introduction to cryptocurrencies, see Shaanan Cohny, David Hoffman, Jeremy Sklaroff & David Wishnick, *Coin-Operated Capitalism*, 119 COLUM. L. REV. 591, 603–05 (2019). See also DAVID FOX & SARAH GREEN, CRYPTOCURRENCIES IN PUBLIC AND PRIVATE LAW (2019).

3. *Certificate of Smart Contract Audit for Edge-Lockdrop*, QUANTSTAMP (Apr. 8, 2019), <https://arena-attachments.s3.amazonaws.com/4282493/a155dc84aa1dfba4cfd3dc6be1e1ebdc.pdf?1557965252> [<https://perma.cc/UPC5-ZUPT>].

4. *Id.* (noting that Lockdrop used Solidity programming language).

5. Neil McLaren, *Gridlock (A Smart Contract Bug)*, MEDIUM (July 1, 2019), <https://medium.com/@nmcl/gridlock-a-smart-contract-bug-73b8310608a9> [<https://perma.cc/V2G2-WT9L>] ("Edgeware's Lockdrop smart contract has processed over \$900 million of ETH and locked up over \$290 million.").

6. *Id.*

7. If the attacker designated cryptocurrency to the next lock address, the total balance would exceed the amount sent by Lockdrop, causing the check to fail. This would freeze the sequestering process permanently in place.

(almost always) contain bugs!”⁸ The question we ask in this paper is simple: can contract law make sense of intractable bugs in transactional code? The answer is likewise simple: yes.

But to their promoters, even buggy “smart contracts” like Lockdrop are the vanguard of a revolution, heralding an age in which code will depose both contract theory and practice.⁹ Enthusiasts argue that when contracts are embodied and performed by code, contracting costs (like negotiation, monitoring, and performance) will fall. Better still, parties won’t need to trust each other, or courts, to be assured that they’ll get what they bargained for. The result: “smart contracts” are offered as a potential solution to an astounding variety of business and social problems. They may transform insurance,¹⁰ financial derivatives,¹¹ consumer protection,¹² corporate governance,¹³ tax filing,¹⁴ voting,¹⁵ supply chain management,¹⁶ bankruptcy,¹⁷ property

8. McLaren, *supra* note 5.

9. Cf. Primavera De Filippi & Samer Hassan, *Blockchain Technology as a Regulatory Technology: From Code Is Law to Law Is Code*, FIRST MONDAY (Dec. 5, 2016), <https://firstmonday.org/ojs/index.php/fm/article/view/7113/5657> [<https://perma.cc/2PX9-8C4K>] (“What makes the blockchain different from other technologies is that smart contracts are actually meant to replace legal contracts.”); Mark Verstraete, *The Stakes of Smart Contracts*, 50 LOY. U. CHI. L.J. 743, 743 (2019) (“The most ardent supporters of smart contracts . . . claim that smart contracts might replace large swaths of the traditional contract system.”).

10. *Smart Contracts: 10 Use Cases for Business*, AMBISAFE, <https://ambisafe.com/blog/smart-contracts-10-use-cases-business> [<https://perma.cc/S9WA-FHMM>].

11. See generally INT’L SWAPS & DERIVATIVES ASS’N, LEGAL GUIDELINES FOR SMART DERIVATIVES CONTRACTS: INTRODUCTION (2019), <https://www.isda.org/a/MhgME/Legal-Guidelines-for-Smart-Derivatives-Contracts-Introduction.pdf> [<https://perma.cc/N5XN-J5KQ>].

12. See generally Joshua Fairfield, *Smart Contracts, Bitcoin Bots, and Consumer Protection*, 71 WASH. & LEE L. REV. ONLINE 35 (2014).

13. Fiammetta S. Piazza, *Bitcoin and the Blockchain as Possible Corporate Governance Tools: Strengths and Weaknesses*, 5 PA. ST. J.L. & INT’L AFFS. 262, 282–86 (2017).

14. Valentine P. Vishnevsky & Viktoriia D. Chekina, *Robot vs. Tax Inspector or How the Fourth Industrial Revolution Will Change the Tax System: A Review of Problems and Solutions*, 4 J. TAX REFORM 6, 20 (2018).

15. See Tsui S. Ng, *Blockchain and Beyond: Smart Contracts*, A.B.A. BUS. L. TODAY (Sept. 28, 2017), https://www.americanbar.org/groups/business_law/publications/blt/2017/09/09_ng [<https://perma.cc/MD66-A7V4>] (“Governments may use smart contracts to manage . . . e-voting.”).

16. See generally Horst Treiblmaier, *The Impact of the Blockchain on the Supply Chain: A Theory-Based Research Framework and a Call for Action*, 23 SUPPLY CHAIN MGMT. INT’L J. 545 (2018).

17. See generally Alan Rosenberg, *Automatic Contracts and the Automatic Stay: A Primer on “Smart Contracts” in Bankruptcy*, 38 AM. BANKR. INST. J. 18 (2019).

rights,¹⁸ and repossession through the internet of things.¹⁹ But there's more. Jurisprudence—in the sense of the fundamental utility of contract doctrine—is on the chopping block.²⁰ For many, smart contracts are the first transformative legal innovation of the millennium.²¹

Perhaps inevitably, “smart contracts,” a term that connotes money, computers, and modernity, has invited a stampede of commentators to speculate about a wide variety of possible contracting technologies.²² Many marvel at the innovation as some kind of utopian simulacra:²³ an immutable,²⁴ transparent exchange,²⁵ occurring

18. See, e.g., Michael Casey, *Could Blockchain Technology Help the World's Poor?*, WORLD ECON. F. AGENDA (Mar. 9, 2016), <https://www.weforum.org/agenda/2016/03/could-blockchain-technology-help-the-worlds-poor> [<https://perma.cc/4U9N-4S2F>].

19. For the canonical description, see Jeremy M. Sklaroff, Comment, *Smart Contracts and the Cost of Inflexibility*, 166 U. PA. L. REV. 263, 271–78 (2017).

20. For a lucid review, see generally Marco Dell'Erba, *Demystifying Technology: Do Smart Contracts Require a New Legal Framework? Regulatory Fragmentation, Self-Regulation, Public Regulation*, U. PA. J.L. & PUB. AFFS. (forthcoming 2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3228445 [<https://perma.cc/48NP-CWX7>].

21. See, e.g., Alexander Savelyev, *Contract Law 2.0: “Smart” Contracts as the Beginning of the End of Classic Contract Law*, 26 INFO. & COMM'C'N TECH. L. 116 (2017); see also *What is Ethereum?*, ETHERSCRIPTER, https://etherscripter.com/what_is_ethereum.html [<https://perma.cc/LV8P-6BJD>] (describing popular smart contract Ethereum as “a new kind of law” that can be “perfectly observed and enforced”).

22. The fish rots from the head. See NICK SZABO, SMART CONTRACTS: BUILDING BLOCKS FOR DIGITAL MARKETS 1 (1996), <http://www.truevaluemetrics.org/DBpdfs/Blockchain/Nick-Szabo-Smart-Contracts-Building-Blocks-for-Digital-Markets-1996-14591.pdf> [<https://perma.cc/6BAR-V45A>] (defining smart contracts as “a set of promises, specified in digital form, including protocols within which the parties perform on these promises”). In recent years, Szabo has vociferously defended the term. See, e.g., Nick Szabo (@NickSzabo4), TWITTER (Oct. 14, 2018, 6:38 PM), <https://twitter.com/NickSzabo4/status/1051603179526270976> [<https://perma.cc/5HC7-D95Z>] (“‘Smart contract’ is a very useful concept & phrase. ‘Smart’ as in ‘smart phone’ (shorthand for computerized phone), ‘contract’ meaning it does some important things we previously relied on contracts to do for our deals, especially controlling assets & incentivizing performance.”).

23. Frank Pasquale, *A Rule of Persons, Not Machines: The Limits of Legal Automation*, 87 GEO. WASH. L. REV. 1, 24–25 (2019).

24. Jeffrey M. Lipshaw, *The Persistence of “Dumb” Contracts*, 2 STAN. J. BLOCKCHAIN L. & POL'Y 1, 24 (2019) (stating that the key characteristics of smart contracts, “in addition to consensus, include immutability and finality from the time they are created and going forward”).

25. Adam J. Kolber, *Not-So-Smart Blockchain Contracts and Artificial Responsibility*, 21 STAN. TECH. L. REV. 198, 208 (2018) (“Since Ethereum smart contracts consist of particular computer code on a decentralized blockchain, it is easy to verify program execution.”).

through “consensus,” that is “self-executing”²⁶ or “automated.”²⁷ With that capacious definition in mind, is the chip in your Visa Card part of a smart contract network? The code comprising the Venmo App? Your Metrocard?²⁸ If we can’t understand the scope of this phenomenon, how can we fairly evaluate the risk it poses, or the benefits it promises, to our commercial and social life?

Smart contracts’ flabby meaning is not this Article’s precise target.²⁹ Rather we offer a focused description of the most celebrated aspect of the underlying technology and its relationship to legal order. We’ll start by introducing a new term that we think captures what most people in this field think of when they consider “smart contracts” as they are currently deployed. We name and describe the *transactional script*.³⁰ Here is a parsimonious definition (that we’ll unpack later):³¹

A transactional script is a persistent piece of software residing on a public blockchain. When executed as a part of an exchange, the code effectuates a consensus change to the state of a ledger.

26. Professors Kevin Werbach and Nicolas Cornell argue that smart contracts are distinctive because “juridical forums are powerless to stop the execution of smart contracts—there is no room to bring an action for breach when breach is impossible.” Kevin Werbach & Nicolas Cornell, *Contracts Ex Machina*, 67 DUKE L.J. 313, 332 (2017); see also Amy J. Schmitz & Colin Rule, *Online Dispute Resolution for Smart Contracts*, 2019 J. DISP. RESOL. 103, 106, 107, 113 (describing smart contracts as “self-enforcing,” “self-governing,” and with “no ambiguity around the parties’ obligations”).

27. Christopher D. Clack, Vikram A. Bakshi & Lee Braine, *Smart Contract Templates: Foundations, Design Landscape and Research Directions 2* (Mar. 15, 2017) (unpublished manuscript), ARXIV: 1608.00771; see also Max Raskin, *The Law and Legality of Smart Contracts*, 1 GEO. L. TECH. REV. 305, 309 (2017) (“A smart contract is an agreement whose execution is automated.”).

28. See Kolber, *supra* note 25.

29. Cf. Ed Felten, *Smart Contracts: Neither Smart nor Contracts?*, FREEDOM TO TINKER (Feb. 20, 2017), <https://freedom-to-tinker.com/2017/02/20/smart-contracts-neither-smart-not-contracts> [<https://perma.cc/NZ2Q-GFHE>].

30. For a definition of smart contracts that tracks with our transactional script, see Carla L. Reyes, *If Rockefeller Were a Coder*, 87 GEO. WASH. L. REV. 373, 383–84 (2019), which states that:

The term “smart contract” refers to decentralized computer code that runs on a DLT protocol and manifests some combination of the following characteristics: exerts some control over assets digitally recorded on a DLT protocol, takes some action upon receipt of specified data, is often, but not always, part of a DLT-based application, guarantees execution, and writes the resulting state change from the operation of the smart contract into the DLT’s ledger.

31. Cf. Peter G.L. Hunn, *Smart Contracts as Techno-Legal Regulation*, 7 J. ICT STANDARDIZATION 269, 275 (2019) (focusing on “a deterministic state machine” and a “consensus protocol” that provides agreement “on the same sequence of operations”).

We stress that our focus—for the moment—is narrower than all digitized exchanges,³² or even all deals accomplished through block-chain-style ledgers.³³ That is, transactional scripts sit at the core of the rapidly expanding group of things called “smart contracts,” but do not encompass the whole field. Crucially, transactional scripts are exchanges that operate in *public*—indeed, they are valuable in large part because their resolution must be agreed to by multiple different entities, jointly operating on an operating system with fixed and translucent rules.

Transactional scripts are a striking legal innovation and much of the current interest in smart contracts in fact regards them.³⁴ But law (and lawyers’) role in the creation and execution of scripts is unclear. Even apart from the performative techno-libertarian claims about law’s abnegation, many scripting projects are missing many of the accoutrements of transactional law. More plainly put, currently deployed transactional scripts—those that are exposing real people to financial and personal risks—often rely on the code alone as their primary risk-allocation mechanism. And the code errs.

In the Edgeware project, none of the governance promises were expressed in a natural language “contract” as we understand that term, even in the digital sphere, as in a click-wrap agreement. They existed rather in a “white paper”—a self-published document with no

32. For a tremendous survey of digitized exchanges, see Harry Surden, *Computable Contracts*, 46 U.C. DAVIS L. REV. 629 (2012). Surden discusses “autonomous computable contracting” only briefly. *Id.* at 694–95.

33. The reader would benefit from reading the following excellent works on scripts (broadly defined): Sklaroff, *supra* note 19; J.G. Allen, *Wrapped and Stacked: “Smart Contracts” and the Interaction of Natural and Formal Language*, 14 EUR. REV. CONT. L. 307 (2018); Werbach & Cornell, *supra* note 26; Karen E.C. Levy, *Book-Smart, Not Street-Smart: Blockchain-Based Smart Contracts and the Social Workings of Law*, 3 ENGAGING SCI. TECH. & SOC’Y 1 (2017); James Grimmelman, *All Smart Contracts Are Ambiguous*, 2 J.L. & INNOVATION 1 (2019); PRIMAVERA DE FILIPPI & AARON WRIGHT, *BLOCKCHAIN AND THE LAW: THE RULE OF CODE* (2018); Jonathan Rohr, *Smart Contracts in Traditional Contract Law, or: The Law of the Vending Machine*, 67 CLEV. ST. L. REV. 67 (2019); Edmund Schuster, *Cloud Crypto Land*, MOD. L. REV. (forthcoming 2020) (manuscript at 23–25), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3476678 [<http://perma.cc/YAY3-TA63>].

34. See *Global Smart Contracts Market*, MKT. RSCH. FUTURE, <https://www.market-researchfuture.com/reports/smart-contracts-market-4588> [<https://perma.cc/F3DN-L98D>] (“The global smart contracts market is expected to reach approximately 300 USD Million by the end of 2023 . . .”); cf. *Venture Capital Firms Go Deep and Wide with Blockchain Investments*, DIAR (Oct. 1, 2018), <https://diar.co/volume-2-issue-39> [<https://perma.cc/V6Q2-ZLVR>] (“[I]n just . . . three quarters of 2018, blockchain and crypto companies have raised nearly 3.9Bn through traditional VC—280% more [than] last year . . .”).

standard form and untested legal effect.³⁵ That white paper states that users will receive “[d]ownside protection” in the case of a “malicious attack or exploit,” and describes the lockdrop contract as a failsafe technique.³⁶ Moreover, the promoters suggest that it would be impossible to falsely claim ownership for an address.³⁷ Needless to say, these representations do not match the reality of what the buggy code delivered. But, equally obviously, since the code was public, its latent errors were also theoretically knowable, at least to sophisticated counterparties.

Other examples of coding errors and oversights abound—including the now infamous DAO hack, which we will discuss later.³⁸ In many cases, the losers of coding errors have paid off the winners to undo transactions.³⁹ Such settlements occur in the very indistinct shadow of law: there isn’t a body of cases that directly address the question of what happens when transactional scripts go wrong.⁴⁰ The academic literature too has largely downplayed the role and capabilities of law in resolving transactional scripts gone wrong.⁴¹

Even the most sophisticated treatments in the literature largely focus on what happens when scripts accomplish their promised aims.⁴² Some argue that the technology is simply incapable of the sorts

35. EDGEWARE WHITEPAPER, *supra* note 1, at 9–15.

36. *Id.* at 6–7.

37. *Id.* at 11–12 (“Upon inspection [of the transaction hash], a verifier should have enough proof that if the owner of the . . . account did not also own the recipient Edgeware account (represented as the target address . . .), then they would not issue such a transaction.”).

38. *See infra* text accompanying notes 276–79.

39. The literature has noted that often firms price bug bounties too low when compared with their after-market (and illegitimate) use. Lorenz Breidenbach, Philip Daian, Florian Tramer & Ari Juels, *Enter the Hydra: Towards Principled Bug Bounties and Exploit-Resistant Smart Contracts*, 27 PROC. USENIX SEC. SYMP. 1335, 1335 (2018). It is difficult to know how common self-help is in the world of transactional scripts. Victims have little incentive to publicize their failure to write better code, while successful attackers may wish to avoid public renown (if not the tax authorities).

40. *See, e.g.*, Lauren Henry Scholz, *Algorithmic Contracts*, 20 STAN. TECH. L. REV. 128, 141 (2017) (discussing the application of contract and agency law to algorithmic contracts and noting lack of caselaw).

41. *Cf.* Dell’Erba, *supra* note 20 (manuscript at 21) (“It could be that there may be a bug in the code or that the parties may reconsider what they want.”).

42. *Compare* Werbach & Cornell, *supra* note 26, at 350–64 (things going well), *with id.* at 365–67 (not as well).

of error that law cares about: the “[c]ode typically entails no ambiguity, and no variant interpretation is possible.”⁴³ Others lament that because the blockchain code is self-contained, it contains no place within it for “default law” to exist.⁴⁴ Worse, even if there are places for law’s tools to have purchase, jurists “may not be able to hypothesize a reasonable human’s interpretation of a given smart contract.”⁴⁵

Such handwaving is an unwarranted, and ultimately unworkable, surrender of the law’s role in adjudicating disputes. Code that embodies commercially-significant scripts will inevitably contain ambiguities, and disappointed parties will ask judges to adjudicate their rights. At the basic level, as James Grimmelmänn has recently observed, the “meaning of any specific program rests on a foundation of some prior agreement about how to interpret some larger class of programs.”⁴⁶ He concludes that while there may be fewer superficial examples of interpretative gaps in formal code, “when the bottom drops out, it can really drop out.”⁴⁷ It’s then that law will be asked to step in and provide rational and predictable solutions, which will almost certainly be developed by analogy to off-chain exchanges.

But analogies can deceive. Whereas scholars and jurists who confront questions about ordinary written contracts have a deep working knowledge of how such exchanges function, transactional scripts are functionally innovative and the legal community has yet to coalesce around even a basic understanding of what they are, let alone what they do. What’s needed, then, is a working knowledge of the ways in which the vocabulary and functioning of the code itself can create disconnects between the intent of the humans coding transactional scripts and the code’s function.⁴⁸ Describing the plumbing of these phenomena is the first contribution of this Article, and occupies Part I. We make two fundamental observations.

43. Wulf A. Kaal & Craig Calcaterra, *Crypto Transaction Dispute Resolution*, 73 BUS. LAW. 109, 136 n.94 (2017–2018).

44. Usha R. Rodrigues, *Law and the Blockchain*, 104 IOWA L. REV. 679, 682 (2019) (“Because the smart ‘contract’ is code alone, there is no gap, in the sense of an entry point, for the law to step in to fill.”).

45. Kaal & Calcaterra, *supra* note 43, at 136; Schmitz & Rule, *supra* note 26, at 103 (“Those with no coding background cannot easily interpret a smart contract in its rawest form.”).

46. Grimmelmänn, *supra* note 33, at 12.

47. *Id.* at 20 (“The relevant community can redefine the programming language in a way that radically alters the meaning of programs written in it.”).

48. On the problems caused by blockchain jargon, see Angela Walch, *Blockchain’s Treacherous Vocabulary: One More Challenge for Regulators*, J. INTERNET L., Aug. 2017, at 1.

First, the software language development environment of almost all extant transactional scripts has features (and entails uncommon practices) which make coding errors, and gaps between coders' intent and expression, both likely to occur and difficult to resolve.⁴⁹ *Second*, the dominant platform for scripts—Ethereum—assesses a tax on complex programs. This fee makes transactional scripts unsuited for many knotty contracting problems, unless their drafters make parts of the code non-public, thus stripping scripts of much of what makes them an elegant solution to problems of trust in exchange.

With a better grasp of how scripts operate, Part II uses case examples to argue—contrary to the dominant account—that these scripts fail in ways that are legible to traditional contractual frameworks. Put simply, they fail to accomplish their parties' intent. Those examples start with "tokens" issued in "initial coin offerings," continue with decentralized exchanges, and finally consider so-called "oracles." We describe both the coded and natural language promises accompanying these projects. Each such category of transactional script has been touted as a revolutionary financial innovation. Each, as we'll show, have already produced errors with real legal consequences.

In Part III, we ask how law ought to respond to the sorts of problems occasioned by such systematic failures of transactional scripts. We argue that the layering of transactional script and natural language promise is best understood as producing a *contract stack*. Contract stacks are inevitable when parties come together to accomplish commercial ends, even if they try their hardest to make the code the final answer to all their problems. Collecting the meager extant caselaw, and deploying old fashioned rules of interpretation, we offer a novel framework through which courts ought to compile such stacks and thus make sense of these new forms of commercial exchange.

The framework we offer, though focused on the scripted ecosystem of the moment, has general application. Future iterations of digitized transactions, whether using blockchain or other forms of software, will require courts to make sense of gaps between natural language promises and coded executions. Our approach brings old principles of common law to bear on the problems that the next generation of contract practitioners and academics will face.

49. See Elaine Ou, *Blockchain Is Littered with 'Smart' Contracts Gone Bad: Opinion*, INS. J. (Nov. 16, 2017), <https://www.insurancejournal.com/news/international/2017/11/16/471387.htm> [<https://perma.cc/P5MF-67N3>].

I. DESIGNING EXPENSIVE, BUGGY SCRIPTS

We begin with a precise account of the creation and function of transactional scripts. To the extent that the discussion requires struggling with new jargon, our petard has been well-hoisted.⁵⁰ Our goal is not merely to demystify this technology, but also the world and social practices of coders whose work increasingly matters to law and transactional practice. We make two primary contributions. *First*, we try to explain why code is intrinsically buggy. *Second*, we offer some reasons to be skeptical about scripts' innate ability to simultaneously solve problems of trust, complexity and automation of deals.

We start with the functional question of how coders go about their work.

A. HOW (COMMERCIAL) CODING WORKS

The process of taking a programmer's intent from code to execution involves multiple steps, each of which can and does introduce error. Just as when law firms draft contracts, programming progresses in pieces and through teams.⁵¹ Programming teams typically start with a human-driven goal. They then choose a language in which to code. Typically, coders will choose a high-level language that abstracts away the details of the machine's hardware and allows programming in syntax closer to natural language.

Transactional scripts are commonly coded in Solidity, the language conceived alongside of, and for use on, Ethereum, the platform on which most scripts operate.⁵² It is syntactically similar to the popular, web-development language, Javascript and thus looks familiar to many non-smart-contract coders. To create a transactional script in Solidity, a coder creates a text file with contents that conform to the publicly available specification of the Solidity language, and that capture in programmatic form the design goal.⁵³

50. Cf. ELIZABETH MERTZ, *THE LANGUAGE OF LAW SCHOOL: LEARNING TO "THINK LIKE A LAWYER"* (2007) (providing the classic text on how law students are taught to think and argue in distinctive ways).

51. On "flexible standardization" and the production of contracts in law firms, see Matthew Jennejohn, *The Architecture of Contract Innovation*, 59 B.C. L. REV. 71 (2018).

52. For a good overview of the coding ecosystem, see Raina S. Haque, Rodrigo Seira Silva-Herzog, Brent A. Plummer & Nelson M. Rosario, *Blockchain Development and Fiduciary Duty*, 2 STAN. J. BLOCKCHAIN L. & POL'Y 1, 16-17 (2019).

53. For more on the process of open source development in blockchain generally, see Angela Walch, *Open-Source Operational Risk: Should Public Blockchains Serve as Financial Market Infrastructures?*, in 2 HANDBOOK OF BLOCKCHAIN, DIGITAL FINANCE, AND INCLUSION 252-54 (David Lee Kuo Chuen & Robert Deng eds., 2017).

Like traditional contract drafters, transactional script developers freely use boilerplate. A significant fraction of code in open source software projects is copied.⁵⁴ This practice is also prominent within the transactional scripts world, with one study finding in 95% of implementations that performed a common function, coders had used identical syntax.⁵⁵ Reusing or retrofitting code for new ends serves multiple purposes for developers. Reuse decreases development time and may provide ready-made, secure, and compatible solutions for difficult coding problems. Code reuse can also cause problems, making developers reliant on code they may not understand, propagating bugs, and increasing development time where reused code is difficult to understand or adapt.⁵⁶

Knitting together these pieces of boilerplate with bespoke code is an iterative process, and programs are created across multiple sessions consisting of coding and testing. To keep track of changes, and the contributions of many individuals,⁵⁷ software projects use version control systems that themselves capture a log of each change.⁵⁸ To add a change to a version control system, coders are typically required to explicitly identify the change/s they wish to add, describe it in a “commit message” and send it to a server that tracks the full set of changes across users.⁵⁹

54. Developers report a mean thirty percent of functionality is derived from re-used code. See Manuel Sojer & Joachim Henkel, *Code Reuse in Open Source Software Development: Quantitative Evidence, Drivers, and Impediments*, 11 J. ASS'N FOR INFO. SYS. 868, 869 (2010).

55. See Yi Zhou, Deepak Kumar, Surya Bakshi, Joshua Mason, Andrew Miller & Michael Bailey, *Erays: Reverse Engineering Ethereum's Opaque Smart Contracts*, 27 PROC. USENIX SEC. SYMP. 1371, 1371 (2018), <https://www.usenix.org/system/files/conference/usenixsecurity18/sec18-zhou.pdf> [<https://perma.cc/GJ94-S6F8>].

56. *Ethereum Smart Contract Best Practices*, GITHUB, https://consensys.github.io/smart-contract-best-practices/general_philosophy [<https://perma.cc/NS6T-MQ34>] (“A smart contract system from a software engineering perspective wishes to maximize reuse where reasonable.”).

57. Programmers will sometimes pair program: a reviewer assesses each line of code as the primary programmer is typing. The reviewer suggests changes, spots errors and often drives the strategic direction of the code. See, e.g., Laurie Williams, Robert R. Kessler, Ward Cunningham & Ron Jeffries, *Strengthening the Case for Pair Programming*, IEEE SOFTWARE, July/August 2000, at 19.

58. Distributed version control systems have much in common with blockchains. They record sequences of changes to a common, replicated record using a chained hash structure that ties together each new change with the complete past history. They do not solve consensus problems, relying on individuals to determine the outcome of conflicts.

59. While decentralized version control systems (DVCS) such as git can in theory operate without a centralized server, the convenience of a host that is always online

```

commit eb0b56b19017ab5c16c745e6da39c53126924ed6
Author: Pieter Wuille <pieter.wuille@gmail.com>
Date:   Fri Aug 1 22:57:55 2014 +0200

    Simplify serialize.h's exception handling

    Remove the 'state' and 'exceptmask' from serialize.h's stream
    implementations, as well as related methods.

    As exceptmask always included 'failbit', and setstate was always
    called with bits = failbit, all it did was immediately raise an
    exception. Get rid of those variables, and replace the setstate
    with direct exception throwing (which also removes some dead
    code).

    As a result, good() is never reached after a failure (there are
    only 2 calls, one of which is in tests), and can just be replaced
    by !eof().

    fail(), clear(n) and exceptions() are just never called. Delete
    them.

```

Figure 1: A well written commit message from the Bitcoin core repository explaining the changes made and the reasoning behind them.⁶⁰

When added to the system, the set of changes is termed a “commit.” The commit log generated by a version control system thus contains evidence of the drafting process, both in code and human readable form. The log and code are available to all developers on a project and may sometimes be stored on a publicly accessible server.

Large firms enforce discipline over this commit log system: commit messages must capture the intent and effect of a change. In smaller development outfits, programmers often fail to include meaningful commit messages and the log may thus poorly capture intent.⁶¹

Another prominent practice common to well-disciplined development teams is *code review*. Before a server or team will accept a commit, it passes through human review. The review is performed by other team members, who comment on a platform integrated with the

and authoritative ensures that most deployments use such a server. Common commercial service providers of these servers include Github, Gitlab, and Bitbucket.

60. <https://github.com/bitcoin/bitcoin/commit/eb0b56b19017ab5c16c745e6da39c53126924ed6.patch> [<https://perma.cc/UL38-WEK3>].

61. Shortform commit messages can range from the insightful, “Simplify serialize h’s exception handling,” to the banal, “fuck fuck holy shit fuck I think I finally fixed my shitty git fuck.” Chris Beams, *How to Write a Git Commit Message*, INCEPTION INNOVATION, <https://inceptioninnovation.com/blog/tutorial-5/post/how-to-write-a-git-commit-message-14> [<https://perma.cc/3Z9G-CPFV>]; Ramiro Gómez, *Exploring Expressions of Emotions in GitHub Commit Messages*, GEEKSTA (May 10, 2012), <https://geeksta.net/geeklog/exploring-expressions-emotions-github-commit-messages/> [<https://perma.cc/NT3D-UPGT>].

version control system. The reviewers assess all elements of the commit—the message, the quality of the code implementing the change and the intent of the change—and either accept, reject, or send the commit back for further review.

Once a program is ready to be tested, it must be converted from the high-level language to machine instructions. The conversion is done through the aid of a compiler, a secondary program developed for this express purpose.⁶² Compilers, as software themselves, are imperfect. Rarely, they contain bugs that cause a mistranslation from the high-level language to the target language, further obscuring the link between programmer intent and program. They might also be designed maliciously or negligently.⁶³

However, if all goes well, the output from the compiling process is bytecode, a low-level representation of the high-level program that the computer can execute more directly. To provide some concrete detail, we now turn to the blockchain platforms of interest.

B. AN INTRODUCTION TO THE BLOCKCHAIN PLATFORM

Blockchains are already the subject of a large legal literature. Here, we sketch only the broadest strokes. Essentially, they are platforms for distributed data processing that create incentives for users to agree on, and store, outcomes of computation.⁶⁴ A blockchain generally consists of two components: storage structures that track

62. There are two primary ways programs are executed: directly or through an interpreter. Directly executed programs undergo multiple translation (compilation) steps taking the program from human-readable source code, to less-human readable assembly code, to machine readable instructions. See *Introduction to Programming Languages/Compiled Programs*, WIKIBOOKS (Sept. 29, 2019, 12:45 AM), https://en.wikibooks.org/w/index.php?title=Introduction_to_Programming_Languages/Compiled_Programs&oldid=3581047 [https://perma.cc/RLC4-9QYN]. Interpreted programs generally undergo a preliminary form of compilation but are not themselves converted into machine instructions. Rather, an interpreter executes the preliminary form by a set of rules acting on a virtual machine. See *Introduction to Programming Languages/Interpreted Programs*, WIKIBOOKS (Sept. 27, 2017, 6:49 PM), https://en.wikibooks.org/w/index.php?title=Introduction_to_Programming_Languages/Interpreted_Programs&oldid=3304944 [https://perma.cc/4UBZ-FZRZ].

63. See Ken Thompson, *Reflections on Trusting Trust*, Turing Award Lecture, 27 COMM'NS ACM 761 (1984). Solidity compilers, with their relative closeness to blockchain based assets, make a similarly attractive target for attacks. The Solidity foundation maintains a list of known vulnerabilities. *List of Known Bugs*, SOLIDITY, <https://solidity.readthedocs.io/en/v0.5.12/bugs.html> [https://perma.cc/XM2S-DAY4].

64. See generally DE FILIPPI & WRIGHT, *supra* note 33, at 33–49; Theophanis C. Stratopoulos & Jesús Calderón, *Introduction to Blockchain for Accounting Students* (Aug. 20, 2020) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3395619 [https://perma.cc/6YP7-NZ8B]; DYLAN YAGA, PETER MELL,

changes in the system and algorithms that ensure consistency of data across storage locations.

A blockchain identifies data records by hashes: the output of an algorithm called a cryptographic hash function.⁶⁵ This easy-to-compute equation takes potentially voluminous data as input and produces a short, fixed-length, output—the hash. Critically, there are no known methods to easily perform the reverse computation.⁶⁶ It is similarly hard to find two different sets of data that when fed into the function both produce the same hash.⁶⁷ This creates a tie between the input data and the hash. The hash acts as the “name” of the block, uniquely identifying it by its contents.

Data is organized into blocks, with each block linking to a set of stored data. Each block contains a hash (below in bold) corresponding to the storage records linked to the block and the hash of the previous block (shown as “prev”).

A block’s hash is computed by feeding the contents of the block into the hash function. The hash therefore is strongly tied to the data within the block, and also ties the block to those that come before. This property ensures an ordering to the blocks, creating the chain aspect of a blockchain. Updates (proposed or accepted) to a blockchain are generally contained within “transactions,” small sets of machine instructions or transactional records that comprise the data within blocks. A block is associated with transaction data through an additional hash stored within the block, in our figure this hash is depicted as *H*(data).⁶⁸

NIK ROBY & KAREN SCARFONE, NAT’L INST. OF STANDARDS & TECH., INTERNAL REPORT 8202, BLOCKCHAIN TECHNOLOGY OVERVIEW (2018), <https://doi.org/10.6028/NIST.IR.8202> [<https://perma.cc/8E82-6GQR>].

65. Yaga et al., *supra* note 64, at 7–13; Stratopoulos & Calderón, *supra* note 64, at 34–40.

66. Note: the cryptographic hash functions used for systems such as blockchains bear additional security properties. See JONATHAN KATZ & YEHUDA LINDELL, INTRODUCTION TO MODERN CRYPTOGRAPHY 128–30 (2007).

67. It would take around 720 times the age of the universe to find a collision for its hash function. See Ashif Shereef, *A Physicist’s Journey into Cracking the Bitcoin*, HACKERNOON (Mar. 27, 2018), <https://hackernoon.com/a-physicists-journey-into-cracking-bitcoin-4631e57158cc> [<https://perma.cc/LTS6-M4BC>].

68. The hash over the data is computed with the aid of an additional structure known as a Merkle Tree. The properties of a Merkle Tree allow one to easily notice if the contents of a particular transaction have changed without checking the entirety of

Using a chain as a public and trusted record requires a mechanism to achieve consensus on the contents of the record. Participants follow a well-defined set of rules that allow them to agree both on the validity of a particular chain, and the ordering (and acceptability) of any proposed updates to the chain.⁶⁹ Participants connect to one another over the internet, forming a subnetwork within the larger network. Participants send and receive messages in a set format to indicate their responses to proposed changes to the blockchain.

The blockchain records each update that has ever happened to the data it stores. By viewing the record of changes, a viewer can access the historical state of the chain. The contents of a blockchain are fixed only so long as the participants agree about what constitutes the set of previous transactions.⁷⁰ Moreover, while the record of past transactions may be unchanged, a future transaction may modify the ledger to make the most recent contents identical in substance to the contents at a prior time (differing records of the past notwithstanding).

Protocols commonly feature *validators* who vie for the scarce opportunity to submit the next block to the network.⁷¹ Those that succeed at adding a block claim a reward.⁷² A cap on the amount of data that can be stored in a single block, in combination with the restricted opportunities to add blocks, limits the number of transactions that can be processed in a unit of time.⁷³

the data. See Shaan Ray, *Merkle Trees*, HACKERNOON (Dec. 14, 2017), <https://hackernoon.com/merkle-trees-181cb4bc30b4> [<https://perma.cc/M9CV-GDJW>].

69. These rules are the protocol governing the system. See Grimmelmann, *supra* note 33, at 8.

70. Forks serve as an “ever-present escape valve” from majority consensus with which a minority of blockchain participants disagree. Haque et al., *supra* note 52, at 34.

71. In such protocols, a subspecies of validators, miners, are also given the opportunity to mint a new unit of the corresponding cryptocurrency, updating the ledger to grant them ownership over the new coin. *Id.* at 149.

72. Opportunities to add blocks are allocated according to the consensus scheme, the two most popular of these being Proof-of-Work, which probabilistically allocates opportunities based on amount of computational effort spent, and Proof-of-Stake which probabilistically allocates opportunities in proportion to a miner’s staked cryptoasset.

73. Dubbed the scaling problem, maximizing transaction throughput is a highly active area of research. Proposed solutions include moving certain transactions off-chain and other novel protocol designs. See, e.g., Connor Blenkinsop, *Blockchain’s Scaling Problem, Explained*, COINTELEGRAPH (Aug. 22, 2018), <https://cointelegraph.com/explained/blockchains-scaling-problem-explained> [<https://perma.cc/PC48-3H75>] (giving an overall explanation of the scaling problem and presenting off-chain transactions as a possible solution).

An individual wishing to transact via the network submits their proposed transaction to the peer-to-peer network and indicates how much they are willing to pay (in cryptocurrency) to have their transaction processed.⁷⁴ Validators will process such user-submitted transactions, claiming associated processing fees.⁷⁵ The size of the fee is inversely correlated with how long the network will take to process the transaction, as a larger fee more strongly incentivizes validators to include that transaction within the next block.⁷⁶

Blockchain ledger entries are associated with identifiers known as addresses, which determine who controls particular ledger entries.⁷⁷ Digital keys that grant control over assets (or data) stored in a given set of entries are stored in “wallets,” which are files or programs containing these keys. Colloquially, some equate control with ownership, which is why you’ll hear about a wallet’s “owners.” Of course, access and the ability to modify doesn’t necessarily mean ownership, unless that apartment key you gave your dog walker was more significant than you thought. At its core, a cryptoasset is nothing more than an entry in the ledger, secured with a set of fancy tools that mean only those people who know the access digits can change its characteristics.

Finally, public blockchains are typically designed so that validators or other users can unilaterally join and leave the network. Permissioned blockchains take a fundamentally different approach, admitting only preapproved validators to the consensus forming process. While potentially useful for a consortium of known parties, or for use within an individual business, permissioned blockchains

74. In some systems, portions of the transaction fee may be optional, but this may result in the network failing to ever process a transaction.

75. Cf. Rebecca Bratspies, *Cryptocurrencies and the Myth of the Trustless Transaction*, 25 MICH. TECH. L. REV. 1, 20–21 (2018) (indicating that individual transaction fees increase as the network load increases, leading to transaction fees over \$55 dollars per transaction at one point for Bitcoin transactions).

76. *Id.*; see also FILIP LUNDIN & FREDRIK RAHM, EVALUATING RISK AND REWARD FOR VALIDATORS IN A CRYPTOCURRENCY PROOF-OF-STAKE NETWORK 11–12 (2018) (explaining the inverse relationship between validators incentivization and transaction process risk).

77. These are normally derived from users’ cryptographic keys or the result of a hash function applied to relevant data.

rely on users' trust in the list validators, preapproved by the entity operating the blockchain.⁷⁸ This limits their use in the low-trust marketplace for which transactional scripts are most often touted.⁷⁹ We therefore focus our analysis on the "permissionless" ecosystem.⁸⁰

C. ETHEREUM AND SCRIPTING

Bitcoin is an awkward commercial platform, mostly useful for recording and facilitating the flow of bitcoin transactions.⁸¹ Realizing the utility of performing more complex operations on a blockchain, a

78. See Sklaroff, *supra* note 19, at 276–77 n.50 (“[T]he advantages of [permissioned] blockchains exist in tandem with reliance on offline identity. . . . [P]articipants on permissioned blockchains are typically bound by off-chain, real-world agreements[.]” (internal quotations omitted)).

79. Most commenters agree with Ian Kane, COO of TERNIO, in thinking that “permission[ed] blockchains have their place specifically in an enterprise environment.” Shehryar Hasan, *Private Blockchains Are Bullshit, Expert Says*, BLOCKPUBLISHER (Jan. 25, 2019), <https://blockpublisher.com/private-blockchains-are-bullshit-expert-says> [<https://perma.cc/9JDU-DDMC>]; see also Justin O’Connell, *What Are the Use Cases for Private Blockchains? The Experts Weigh in*, BITCOIN MAG. (June 20, 2016), <https://bitcoinmagazine.com/articles/what-are-the-use-cases-for-private-blockchains-the-experts-weigh-in-1466440884> [<https://perma.cc/DT5Y-ENTJ>] (finding the value of private blockchains is in their ability to “provide interesting opportunities for businesses to leverage [their] trustless and transparent foundation for internal and business-to-business use cases”); cf. EUR. UNION BLOCKCHAIN OBSERVATORY & F., BLOCKCHAIN AND THE GDPR 16 (2018), https://www.eublockchainforum.eu/sites/default/files/reports/20181016_report_gdpr.pdf [<https://perma.cc/74D2-QF2X>] (suggesting that permissioned blockchains might need permissionless blockchains in order to be globally interoperable).

80. The computer science literature explores a variety of scenarios in which network participants have greater or lesser trust in other participants. These works occupy the space between assuming an overwhelming majority of participants are honest and assuming preexisting trusted relationships with other network participants (permissioned models). Such constitute a middle ground between permissioned and permissionless approaches to distributed computation, and generally represent a trade-off between scalability and trust. While laudable efforts reduce the level of trust that it is necessary to sacrifice for substantial gains in scalability, these subtleties are not the focus of our work. For treatment of a technique for scaling blockchain computations and for a discussion on other approaches, see Harry Kalodner, Steven Goldfeder, Xiaoqi Chen, S. Matthew Weinberg & Edward W. Felten, *Arbitrum: Scalable, Private Smart Contracts*, 27 PROC. USENIX SEC. SYMP. 1353, 1353 (2018).

81. While Bitcoin can be scripted to perform sophisticated tasks, doing so is challenging. Coders wishing to do so must work against the limited functionality of the platform. Jacqui Frank & Sara Silverstein, *Vitalik Buterin Created One of the World’s Largest Cryptocurrencies in His Early Twenties—Here’s How He Did It and Why*, BUS. INSIDER (Feb. 13, 2019, 6:22 AM), <https://www.businessinsider.com/vitalik-buterin-created-ethereum-one-of-the-worlds-three-largest-cryptocurrencies-2019-1> [<https://perma.cc/TD7K-XJRN>] (relaying Buterin’s belief that Bitcoin is a limited functionality medium).

programmer named Vitalik Buterin proposed and developed Ethereum, a blockchain based computing platform, with an associated cryptocurrency, Ether.⁸² The protocol's explicit goal was to permit enhanced scripting—more complicated logical operations than recording ownership—on a blockchain. Ethereum uses the Ethereum Virtual Machine (EVM)—a software system with predefined rules and operations, which you can think of as a simulated computer.⁸³

The Ethereum Virtual Machine enables programs to store data on the Ethereum blockchain and defined how the original data could be modified.⁸⁴ By the nature of the blockchain, all data is public and therefore replicable; the EVM, however, allows developers to restrict how the data may be modified. Storage and control, in turn, created the platform for transactional scripts. Scripts are programs operating on Ethereum, which when executed (or “called”) affect the ledger itself.

82. On the various kinds of cryptocurrencies and their relationship to monetary policy, see Max Raskin, Fahad Saleh & David Yermack, *How Do Private Digital Currencies Affect Government Policy?* (Nat'l Bureau of Econ. Rsch., Working Paper No. 26219, 2020).

83. Kolber, *supra* note 25. Solidity, the most common language used to program Ethereum transactional scripts is Turing complete, meaning it can in theory perform any program. This is limited only by the restrictions the protocol places on complexity to mitigate denial of service attacks. See Niharika Singh, *Turing Completeness and the Ethereum Blockchain*, HACKERNOON (Feb. 16, 2019), <https://hackernoon.com/turing-completeness-and-the-ethereum-blockchain-c5a93b865c1a> [<https://perma.cc/7EU5-EVXP>]. The EVM draws its inspiration from common models of computer architecture with modifications that ensure the integrity of scripts' code. While this architecture is foreign to the computers on which the software runs, it is simulated by way of the virtual machine. See Preethi Kasireddy, *How Does Ethereum Work, Anyway?*, MEDIUM (Sept. 27, 2017), <https://medium.com/@preethikasireddy/how-does-ethereum-work-anyway-22d1df506369> [<https://perma.cc/5G3V-MCDM>].

84. For a more complete discussion on the limits of the Bitcoin model, see ConsenSys, *Thoughts on UTXO by Vitalik Buterin (Co-Founder of Ethereum)*, MEDIUM (Mar. 9, 2016), <https://medium.com/@ConsenSys/thoughts-on-utxo-by-vitalik-buterin-2bb782c67e53> [<https://perma.cc/78P6-GRDJ>]. As agents are able to manipulate the records stored on the blockchain, transactional scripts are able to transfer value contingent on those records. This is the true source of their utility over other types of programs that interact with distributed databases: the tight integration of an asset storage mechanism (cryptocurrencies and cryptoassets) with a programmatic way to transfer ownership (transactional scripts). The coupling of the two ensures that the value of the asset is conditional on playing by the rules of the game—which in turn provides certain assurances that the contracts will be executed.

The instruction set provided by Ethereum is as powerful and expressive as any other programming language.⁸⁵ If limits are not imposed on transactional script running time and resource requirements, a malicious actor could force validators to perform never ending computations, halting all useful work on the chain. Validators must also be incentivized to spend their computational resources evaluating scripts. Ethereum therefore imposes what we term a *complexity tax*: transaction fees proportional to the computation required by a transaction.⁸⁶ This fee is known as *gas* and is paid in fractional amounts of ether.⁸⁷ The Ethereum protocol specifies a hard limit on how much gas may be consumed by a single transaction.⁸⁸ Thus, the *amount* of gas paid per transaction is determined by the Ethereum protocol, while the *exchange rate* is determined by the user. The user must also pre-pay gas that in their estimation will sufficiently cover costs. If this pre-payment is too low, and there is insufficient gas to completely execute the script, the script will terminate without a refund.⁸⁹

There are 70 different operations understood by the EVM, each of which is associated with a cost, based on the amount of time and energy it takes a validator to execute the operation. Here are only a few, with their associated costs.

85. Ethereum scripts require special purpose languages and tools to deploy. There is active research in designing languages that facilitate better development practices. See generally Mudabbir Kaleem, Anastasia Mavridou & Aron Laszka, *Vyper: A Security Comparison with Solidity Based on Common Vulnerabilities*, 2ND CONF. ON BLOCKCHAIN RSCH. & APPLICATIONS FOR INNOVATIVE NETWORKS & SERVS., June 14, 2020.

86. See Bruno Skvorc, *Ethereum: How Transaction Costs Are Calculated*, SITEPOINT (May 24, 2018), <https://www.sitepoint.com/ethereum-transaction-costs> [<https://perma.cc/NAU4-RTYJ>] (describing the proportionality of gas costs in Ethereum's structure).

87. For a detailed explanation of the internal mechanisms of Ethereum, see generally Kasireddy, *supra* note 83.

88. At the time of writing, this was set at approximately 8 million gas (with an upcoming change to 10 million gas), equivalent to approximately \$2 USD at a gas price of 1 Gwei (billionths of an ETH).

89. Existing literature describes the complexity tax in general terms without elaborating on its exact costs. See, e.g., Thibault Schrepel, *Is Blockchain the Death of Antitrust Law? The Blockchain Antitrust Paradox*, 3 GEO. L. TECH. REV. 281, 292 (2019) (noting that Ethereum rewards successful miners with transaction fees); Sklaroff, *supra* note 19, at 293 n.139 (describing the "supply and demand dynamic" created and its policing effects on "buggy or infinitely recursive code"); *id.* at 294 n.143 ("[T]his solution would be prohibitively expensive from a transaction fee perspective. . .").

Operation Name	Max Cost (Gas)	Cost (\$)/1 million operations ⁹⁰	Effect
ADD	3	~\$30	Adds two numbers together
MUL	5	~\$50	Multiplies two numbers
SSLOAD	200	~\$2000	Load a single number from permanent storage ⁹¹
SSTORE	20000	~\$200,000	Store a single number into permanent storage

The most expensive operation, SSTORE, updates data that constitute the Ethereum ledger's memory.⁹² As it adds to the long-term cost of storing the ledger, the protocol imposes a substantial up-front fee.

Data storage on a public ledger forms a key component of many proposed blockchain use-cases,⁹³ but it is financially costly in comparison to general purpose storage (which is available from commercial providers on a yearly basis at approximately one-ten-millionth of the cost of Ethereum based storage).⁹⁴ Similarly, the processor in a typical laptop performs on the order of billions of operations per second,

90. These figures are calculated at the default gas price of 3×10^{-8} gas/ETH and an exchange rate of \$295 USD/ETH and are rounded up to one significant figure. See Danny Ryan, *Calculating Costs in Ethereum Contracts*, HACKERNOON (Sept. 6, 2017), <https://hackernoon.com/ether-purchase-power-df40a38c5a2f> [<https://perma.cc/6SUH-ADGR>].

91. The maximum value on which the EVM operates on in a single operation is 2^{256} . Outside of technical reasons for this choice, this limit prevents coders from storing all their data represented as one huge number in an attempt to pay lower gas costs.

92. See Ting Chen, Xiqi Li, Xiapu Luo & Xiaosong Zhang, *Under-Optimized Smart Contracts Devour Your Money*, 2017 IEEE 24TH INT'L CONF. ON SOFTWARE ANALYSIS, EVOLUTION & REENGINEERING, Mar. 11, 2017 (indicating that SSTORE is particularly expensive among smart contract operations in the Ethereum Virtual Machine).

93. Even for systems that purport to store the bulk of data off-chain, the cost of storing references to the off-chain data necessitates a cost-benefit approach to assessing the viability of a blockchain based solution. See MEDICALCHAIN, WHITEPAPER 2.1, at 1, 14 (2018), <https://medicalchain.com/Medicalchain-Whitepaper-EN.pdf> [<https://perma.cc/KSC8-AUS4>] (observing that while records themselves are stored off chain, each update to a record necessitates storing a new value on the blockchain); see also ARMAN JABBARI & PHILIP KAMINSKY, BLOCKCHAIN AND SUPPLY CHAIN MANAGEMENT (2018), <http://www.mhi.org/downloads/learning/cicmhe/blockchain-and-supply-chain-management.pdf> [<https://perma.cc/925L-X9R7>] (envisioning large-scale supply-chain data storage on the blockchain).

94. See *Amazon S3 Pricing*, AWS, <https://aws.amazon.com/s3/pricing> [<https://perma.cc/4VHU-77H6>].

which would equate to hundreds of dollars' worth of computations on Ethereum.⁹⁵ Gas costs don't stem from the regular unit cost of storing a chunk of data or performing a computation.⁹⁶ They are artifacts of the replicated work and storage used to maintain and validate consensus.⁹⁷

This complexity tax has limited, but real implications. We concede that neither storage nor complex programming are what blockchains are for, and such special purpose tools shouldn't be evaluated in comparison with traditional computers. But it's important to recognize that for some proposed uses of transactional scripts—such as encoding semantic frameworks like discretion, good faith, and best efforts—public blockchain solutions incur significant costs.⁹⁸

The result of such costs is that it is practically impossible to run some kinds of scripts in public. Many real-world computational tasks (searching and sorting large amounts of data, machine learning, optimization) require non-trivial amounts of computational power and/or data storage. For example, though it would be potentially useful to write a script that used an algorithm to determine if a worker had used her best efforts, and then pay her for her time, that kind of computation is not practical other than by delegating the computation to an agent “off-chain.”⁹⁹ Current scripts thus generally contain only

95. Some sites indicate that modern processors process operations at the tens of billions of floating-point operations per second. On the scale of our illustrations, that is comparable with the basic arithmetic operations of the EVM. See, e.g., *CPU Performance*, ASTEROIDS HOME, https://asteroidsathome.net/boinc/cpu_list.php [<https://perma.cc/39EF-LGL6>].

96. Noted blockchain researcher Emin Gün Sirer correctly notes that costs of storing the blockchain itself are negligible on a per node basis. Emin Gün Sirer (@el33th4xor), TWITTER (Nov. 29, 2019, 1:13 PM), <https://twitter.com/el33th4xor/status/1200477778463907841> [<https://perma.cc/2HDN-CCL5>]. Our concern is how this cost is magnified by current consensus protocols.

97. These particular costs are distinct from the computational cost incurred by Proof-of-Work style consensus algorithms that establish consensus in the first instance. Rather, the need to check the validity of the output from computations and maintain copies of the output (to allow failure recovery), impose a cost that is paid ex ante by a tax on expensive computations.

98. While Sklaroff notes the potential high ex ante cost of developing contracts flexible enough to incorporate legal frameworks, we here quantify those costs, focus on the ongoing taxes imposed by the blockchain paradigm, and expand on coding realities that further drive development costs. Sklaroff, *supra* note 19, at 279.

99. Off-chain transactions are those that move value from the blockchain to be processed externally. Such transactions may eventually be reconciled and integrated back onto the chain. See *Off-Chain Transactions*, BITCOIN WIKI (June 18, 2019, 4:43 AM), https://en.bitcoin.it/wiki/Off-Chain_Transactions [<https://perma.cc/9YD3-VWLQ>].

the simplest of if-then type logic.¹⁰⁰ A useful analysis of the legal significance of our scripts requires a clear-eyed evaluation as to the likely future uses and limitations of the form.

There is a robust technical literature exploring mechanisms to avoid the complexity tax. But most proposals to mitigate costs trade-off between the security and transparency of the on-chain model, and the efficiency attainable under frameworks that assume more on the part of other protocol participants.¹⁰¹ While there are some improvements that can be made to efficiently scale blockchains without compromising the promise of truly trustless exchange,¹⁰² the most effective solutions will inevitably delegate the bulk of computation and storage to smaller subsets of the network or even to non-participants off chain.¹⁰³

The complexity tax thus produces a real hurdle for certain kinds of complex contracting. There are known techniques to reduce the cost of most computations, but finding a way to minimize blockchain data storage costs while ensuring the data is accessible on demand is an open problem.¹⁰⁴ Firms considering off-chain solutions must there-

100. The quarrelsome reader may note that all imperative paradigm programming adheres to an if-then paradigm. We merely note the paucity of sophistication in scripts currently deployed. For a brief discussion on different programming paradigms, see Jing Chen, *A Brief Survey of "Programming Paradigms,"* MEDIUM (Apr. 11, 2019), <https://medium.com/@jingchenjc2019/a-brief-survey-of-programming-paradigms-207543a84e2b> [<https://perma.cc/V3WE-HECF>].

101. For an analysis of various designs facilitating off-chain computation by a smaller number of nodes and reducing the cost imposed by large scale replication of work, see Lewis Gudgeon, Pedro Moreno-Sanchez, Stefanie Roos, Patrick McCorry & Arthur Gervais, *SoK: Layer-Two Blockchain Protocols*, <https://eprint.iacr.org/2019/360.pdf> [<https://perma.cc/N9MZ-CK3U>]. See also Kyle Croman, Christian Decker, Ittay Eyal, Adem Efe Gencer, Ari Juels, Ahmed Kosba, Andrew Miller, Prateek Saxena, Elaine Shi, Emin Gün Sirer, Dawn Song & Roger Wattenhofer, *On Scaling Decentralized Blockchains*, in FINANCIAL CRYPTOGRAPHY AND DATA SECURITY 106, 106–10 (Jeremy Clark, Sarah Meiklejohn, Peter Y.A. Ryan, Dan Wallach, Michael Brenner & Kurt Rohloff eds., 2016).

102. One proposal that succeeds in maintaining a trustless network while improving substantially on scaling is the Avalanche Protocol, which proposes an alternative to proof-of-work. See Team Rocket, Maofan Yin, Kevin Sekniq, Robbert Van Renesse & Emin Gün Sirer, *Scalable and Probabilistic Leaderless BFT Consensus Through Metastability* (Aug. 24, 2020) (unpublished manuscript), ARXIV: 1906.08936. However even protocols that minimize computational costs such as Avalanche and Arbitrum still don't realize the promise of effectively free transactional scripts. See Kalodner et al., *supra* note 80, at 1353–54 (giving a description of Arbitrum).

103. Importantly for our analysis of buggy scripts in Part III, *infra*, even where off-chain scaling techniques are adopted, the code remains publicly accessible.

104. See Croman et al., *supra* note 101, at 108–10.

fore ask whether the game is worth the candle. That is, public blockchains have important virtues: primarily, they enable trustless exchange between pseudo-anonymous counterparties. Blockchain-based alternatives may not have those virtues.

D. CODING ON ETHEREUM

We now return to our coders' work and focus on how they write transactional scripts. Let's suppose that a team of coders has authored high-level code and run that code through the compiler, producing bytecode that can now be executed on the EVM.

If they are responsible, they must now test the program; only after testing will it be deployed. During development programmers will go through many, many cycles of coding, compiling and testing, finding bugs and areas that need improvement.¹⁰⁵

105. Cf. Rich Butkevic, *The Lifecycle of a Software Bug*, OPENSOURCE (June 25, 2018), <https://opensource.com/article/18/6/life-cycle-software-bug> [<https://perma.cc/2FMU-YVKE>] (describing how development teams fix bugs).

To make the discussion more concrete, we present the *StagedContract* script, written in Solidity, and explain how it is processed by the virtual machine. It illustrates a simple deal: the owner of the script agrees to pay for work in increments of 1ETH as the recipient completes stages of some off-chain job.¹⁰⁶ The recipient can claim payment for each stage without the intervention of the owner up to the total value of the contract.

```
contract StagedContract {

    // THE JOB IS WORTH 5ETH
    int256 public transferlimit = 5;
    int256 public amounttransferred = 0;
    int256 public discount = 0;

    address payable public owner =
        0x2BcEB4B315Fd823Bd93cb9065A97C7A7d0174E93;
    address payable public recipient =
        0x1161B70D1ddc964785189ab7CFf5006cbbEfab4E;

    // ALLOW THE OWNER TO APPROVE A DISCOUNT
    function approveDiscount() public {
        if(msg.sender == owner){
            transferlimit = transferlimit - discount;
        }
    }

    // ALLOW THE RECIPIENT TO DISCOUNT THEIR WORK BY A FIXED
    // AMOUNT
    function offerDiscount(int256 _amount) public {
        if(msg.sender == recipient){
            discount = _amount;
        }
    }

    function completeStage() public {
        // CHECK IF RECIPIENT IS RUNNING THIS,
        // AND THEY HAVEN'T BEEN PAID THE ENTIRE AMOUNT YET
        if(amounttransferred < transferlimit &&
            msg.sender == recipient) {
            revert();
        } else {
            // PAY THE RECIPIENT 1ETH FOR EACH STAGE OF THE
            // WORK
            // THEY CAN CLAIM IT AT ANY TIME
            amounttransferred = amounttransferred + 1;
            recipient.transfer(1);
        }
    }
}
```

Figure 2: Transactional Script in Solidity

106. Not depicted is a mechanism for the owner to increase the amount of payment available over time.

StagedContract also allows the recipient to offer a discount through the `discount()` function which the owner can approve via the `approveDiscount()` function. Illustrative of the complexity tax, at time of writing, it cost \$1.30 to deploy the script and \$0.08 to execute the `completeStage()` function even once. (Imagine using this script for payment of hundreds of thousands of gig workers.)

Like most computer programs, Solidity requires the availability of fundamental computational structures:¹⁰⁷ basic arithmetic operations, operations to access and modify some form of memory (the “state”), a way to input/output data, and crucially, a mechanism to choose between different possible execution paths based on the state of the computer. These are the key elements implemented by the EVM.

Even this very simple script requires expertise to conceive of, implement and deploy. It requires still more expertise to recognize that it contains a serious bug. Noting that if an integer grows or shrinks beyond the limit set by the EVM (overflow and underflow respectively), the integer resets, a sneaky recipient could mount the following attack:¹⁰⁸

1. The recipient offers a discount that reduces the cost of the job to nothing, which the greedy owner readily accepts.
2. The recipient offers an additional discount in advance of future work and waits until just before the owner approves it.
3. The recipient quickly engorges their discount to the maximum size permitted by the EVM.
4. The unsuspecting owner fails to notice the change and approves the oversized discount, reducing the sum available for payment so much that it underflows—resetting the amount available for payment to the *maximum* value permitted by the EVM.
5. The recipient now executes `completeStage()` repeatedly, emptying the script of its Ether, without running into the limit.

107. These, along with the ability to execute loops, are also the elements that permit the EVM to support a Turing complete language. See Singh, *supra* note 83.

108. While the scenario is admittedly contrived, it serves to illustrate the difficulty of writing correct code.

This is but one of many ways in which the specifics of the EVM can trip up otherwise code-literate individuals.

The particular problem for scripts is that when they are already on blockchain and parties have begun committing assets, they can't be easily modified (at least without pushing out a unilateral update which might cause serious reputational blowback).¹⁰⁹ Therefore, wise developers would first try it out on "testnets," small scale blockchains replicating the behavior of their larger siblings, but without the goal of preserving immutability or value long term.¹¹⁰ Testnets are also used to trial changes to the core protocol of a blockchain, without impacting existing users.¹¹¹ Testnet transactions are generally free.¹¹²

When a developer is ready to test their transactional script, either on the real network or on the testnet, they submit a special transaction to the network. The transaction includes the bytecode of the transactional script, an amount of gas to pay for the deployment, and a wallet from which the gas will be transferred out. The transactional script is then stored on the blockchain, and its component functions and storage can be accessed.

Modern programming languages, including Solidity and the Ethereum EVM instruction set, separate segments of code that perform discrete operations into "functions," each of which require a developer using the function to provide certain inputs. Once a function finishes executing, it produces an output which is provided to either the user or, in instances where a function is executed ("called") within another function (the "caller"), the output is available for use within the caller. To call a function, the caller must also include gas sufficient to pay for the execution of the callee (and in a recursive fashion for any functions the callee might call).

109. See Andrew Chow, Comment to *Why Is It Impossible to Modify a Transaction in the Blockchain*, STACKEXCHANGE: BITCOIN (Mar. 30, 2018, 5:02 PM), <https://bitcoin.stackexchange.com/questions/73229/why-is-it-impossible-to-modify-a-transaction-in-the-blockchain> [<https://perma.cc/H65D-AGLA>] ("Overall, modifying transactions already in the blockchain requires remining blocks, and after a transaction already has a few confirmations, doing this requires immense amounts of computing power.").

110. *Testnet*, BITCOIN WIKI (Apr. 14, 2020), <https://en.bitcoin.it/wiki/Testnet> [<https://perma.cc/5743-J2KZ>].

111. See MyEtherWallet, *Understanding Blockchain Changes: Testnets and Mainnets*, MEDIUM (May 26, 2019), <https://medium.com/myetherwallet/understanding-blockchain-changes-testnets-and-mainnets-c2171a8e835f> [<https://perma.cc/C5ZL-J7DN>] (indicating the potential uses of testnets).

112. See *id.* ("[I]n Ethereum's case, the gas payment for testnet computations does not cost any 'real' money . . .").

This all assumes the coders have implemented each step of this process correctly. But, as the *StagedContract* example illustrates, errors can be subtle. There simply is no foolproof method to generate software that matches its initial specification.¹¹³ Further, even as software is patched to remove old bugs many new bugs creep in, and software does not converge to a bug-free state.¹¹⁴

Penetration testing and security auditing are other important components of sophisticated software development. Specialized security engineers attempt to find and exploit security flaws and assess the quality of code as relevant to security.¹¹⁵ However, security failures bedevil even the best audited software packages.¹¹⁶ The real test for code, particularly when designed to operate adversarial environments, comes only when it is deployed and used. It is often only when the code meets the road that developers find the bulk of bugs, improving on their code by constant iteration.¹¹⁷ Even then, vulnerabilities may remain latent for long periods of time prior to discovery.¹¹⁸

113. See Wenbo Guo, Dongliang Mu, Xinyu Xing, Min Du & Dawn Song, *DEEPVSA: Facilitating Value-Set Analysis with Deep Learning for Postmortem Program Analysis*, 28 PROC. USENIX SEC. SYMP. 1787, 1787 (2019) (remarking on the inevitability of flaws in software).

114. See Saender A. Clark, *The Software Vulnerability Ecosystem: Software Development in the Context of Adversarial Behavior* (2017) (Ph.D. dissertation, University of Pennsylvania), <https://repository.upenn.edu/cgi/viewcontent.cgi?article=4019&context=edissertations> [<https://perma.cc/G87G-JWR6>].

115. Due to the level of specialization required, such engineers or testers are contracted from boutique firms, with fees comparable to those of high-end lawyers. This puts many well-performing service firms outside the price range of the majority of transactional script developers. See ekotysh, *How Much Does a Smart Contract Audit Cost?*, REDDIT (July 24, 2017, 11:03 PM), https://www.reddit.com/r/ethdev/comments/6pdgvd/how_much_does_a_smart_contract_audit_cost [<https://perma.cc/GL7T-ATHS>].

116. OpenSSL, one of the most scrutinized software packages (and which underlies much of the cryptography used to secure the Internet), has disclosed over eleven high-severity vulnerabilities since 2014, despite having been created prior to 2002. See *Vulnerabilities*, OPENSSL, <https://www.openssl.org/news/vulnerabilities.html> [<https://perma.cc/97CK-XMD5>].

117. Some experts estimate the density of bugs is 10 to 50 bugs per KLOC (1000 lines of code). Even Microsoft with its sophisticated development practices still releases code with a bug density of ~0.5 per KLOC. This is equivalent to tens-of-thousands of bugs in a modern operating system compiled from tens-of-millions of lines of code. See STEVE MCCONNELL, *CODE COMPLETE: A PRACTICAL HANDBOOK OF SOFTWARE CONSTRUCTION* 652 tbl.27-1 (2d ed. 2004), <http://aroma.vn/web/wp-content/uploads/2016/11/code-complete-2nd-edition-v413hav.pdf> [<https://perma.cc/9QXQ-NQVL>] (showing the range of defect densities).

118. See Clark, *supra* note 114, at 74–75.

These characteristics of the vulnerability life cycle pose challenges for a transactional script developer, where there are limited opportunities for safe testing on a live blockchain or patching *ex post*. Worse, the semantics of the Solidity coding language in which almost all transactional scripts are written are substantially different from traditional software development, leading to overconfident developers ignoring potential pitfalls such as reentrancy vulnerabilities that are not present in most coding environments.¹¹⁹

The need for security also imposes a second, more metaphysical “complexity tax:” the more complex a transactional script, the harder it is to preserve both the coder’s intention (and the intention of the person hiring the coder and so on), while not creating any insecurities; more development time must be expended to shore up the code.¹²⁰ The bigger the transactional script ship, the more effort must be expended to patch the leaks.

E. SUMMARY

A recent survey of script developers suggests that coders are becoming increasingly aware of the problematic real-world consequences of how Solidity interacts with development.¹²¹ Developers, many of whom were working for free,¹²² had ideological motives: to

119. Reentrancy vulnerabilities are a class of flaw where a function calls another function that is external to the given script, and that callee unexpectedly calls the original function before the original function has finished executing. See *Known Attacks*, GITHUB, https://consensys.github.io/smart-contract-best-practices/known_attacks [<https://perma.cc/FFR3-EWAU>]. Programming language considerations that contribute to the difficulty of developing secure and correct Solidity code include the potential for “integer overflow” (where the size of numbers exceeds the space available to store them), a lack of support for decimal numbers and incomplete formal specification of the language. See throwies11, *What Are the Main Security Problems Associated with Solidity Language?*, REDDIT (Jan. 18, 2018, 5:44 PM), https://www.reddit.com/r/ethdev/comments/7rdocn/what_are_the_main_security_problems_associated [<https://perma.cc/V9R9-XZ5G>].

120. See Yonghee Shin, Andrew Meneely, Laurie Williams & Jason A. Osborne, *Evaluating Complexity, Code Churn, and Developer Activity Metrics as Indicators of Software Vulnerabilities*, 37 IEEE TRANSACTIONS ON SOFTWARE ENG’G 772, 772 (2010) (expounding on the difficulties of finding vulnerabilities in code).

121. See Amiangshu Bosu, Anindya Iqbal, Rifat Shahriyar & Partha Chakraborty, *Understanding the Motivations, Challenges and Needs of Blockchain Software Developers: A Survey*, 24 EMPIRICAL SOFTWARE ENG’G 2636, 2652 (2019).

122. *Id.* at 2644 (indicating that developers reported they were not directly compensated).

“create a decentralized currency that cannot be manipulated by a central authority.”¹²³ That is, “removing power from banks and governments.”¹²⁴ These were motivated, committed, blockchain proponents.

But the survey respondents noted that the above-mentioned security concerns and high stakes made coding difficult. “In most . . . [non-blockchain projects] when a bug appears, it will be fixed and soon forgotten. But in blockchain projects some bugs can be very costly and never forgotten.”¹²⁵ Similarly, some complained that erroneous ledger entries are “almost impossible” to fix.¹²⁶ These unique blockchain problems, when coupled with the decentralized VM on which software operates, “makes it difficult to build robust software. Unreliable connections, unexpected latency, and malicious nodes create a hostile production environment.”¹²⁷ There are also few production-ready tools to work through errors in scripts, particularly a “reliable and user-friendly decompiler.”¹²⁸

In sum: it’s simply impossible to create perfect software the first time through, and existing tools to pre-test scripts before deployment are inadequate or extremely costly. Irreducible features of Ethereum (and other blockchains designed using the same logics) will render transactional scripts buggy. One recent study, looking only at the very simple ecosystem of scripts on Ethereum, found 100 errors per 1000 lines of code.¹²⁹ This error rate likely will increase as developers pursue ever-more-ambitious Ethereum projects. As bugs accrue and create real-life losses, parties will turn to tribunals and to the law for recourse. In the next section, we offer some concrete examples of this insight.

123. *Id.* at 2650.

124. *Id.*

125. *Id.* at 2652.

126. *Id.*

127. *Id.* at 2650.

128. *Id.* at 2661.

129. See Peter Vessenes, *Ethereum Contracts Are Going to Be Candy for Hackers*, VESSENES (May 18, 2016), <https://vessenes.com/Ethereum-contracts-are-going-to-be-candy-for-hackers> [<https://perma.cc/RWK5-METY>] (“My review of Ethereum Smart Contracts . . . shows a likely error rate of something like 100 per 1000, maybe higher.”); see also Ivica Nikolić, Aashish Kolluri, Ilya Sergey, Prateek Saxena & Aquinas Hobor, *Finding the Greedy, Prodigal, and Suicidal Contracts at Scale* (Mar. 14, 2018) (unpublished manuscript), ARXIV: 1802.06038 (finding a near 3.5% vulnerability rate across an analysis of one million scripts); Ludovica Marchesi, Michele Marchesi & Roberto Tonelli, *ABCDE—Agile Block Chain Dapp Engineering* (Dec. 19, 2019) (unpublished manuscript), ARXIV: 1912.09074 (“The feeling of many software engineers about such huge interest in Blockchain technologies is that of unruléd and hurried software development . . .”).

II. TRANSACTIONAL SCRIPTS IN THE REAL WORLD

Now that we have in hand a better understanding of what writing a transactional script entails—and where error might creep into that process—let's consider three typical use cases of scripts in the current blockchain ecosystem.

A. TOKENS

A basic use of a transactional script is to change a blockchain record to debit a cryptocurrency from a single address' entry and credit the entries of other addresses. However, the flexibility of the EVM allows for more sophisticated types of trades, both of ether (the Ethereum-defined base currency) and script-defined cryptoassets.¹³⁰

ERC-20 is the Ethereum technical standard that provides a template for creating a fungible, tradeable asset, known as a token.¹³¹ Such tokens are the most common cryptoasset.¹³² Token balances are not stored by the owner of the token. Instead, the transactional script that created the asset updates an internal ledger of addresses and their corresponding token balances. Below we provide some of the code that creates such tokens.

130. See *Ethereum Whitepaper*, ETHEREUM, <https://ethereum.org/en/whitepaper> [<https://perma.cc/LV92-E3HC>] (last updated July 9, 2020) (indicating that Ethereum has the functionality to not only process debit transactions, but also more complex smart-contract transactions via script).

131. The standard requires a compliant script to include a method to determine the total number of such tokens in circulation, the number of tokens owned by a given address, and functions that facilitate the transfer of tokens. See *ERC20*, BITCOIN WIKI (Oct. 29, 2018, 7:46 AM), <https://en.bitcoinwiki.org/wiki/ERC20> [<https://perma.cc/4Z28-3NS4>].

132. See Jake Frankenfield, *Crypto Tokens*, INVESTOPEDIA (June 30, 2020), <https://www.investopedia.com/terms/c/crypto-token.asp> [<https://perma.cc/5PE9-7PZ7>] (describing the popularity of crypto tokens as derivatives of the overarching cryptocurrency infrastructure).

In the following script excerpt the variable *balances* serves as the script's internal ledger of account balances. While any human can manually inspect the ledger, an Ethereum script can only access the balances via the `balanceOf()` function which outputs the balance for a given address. This includes the token script itself, which uses `balanceOf` within the transfer function to check if there is sufficient balance available to debit, before subtracting that amount from one address and crediting to another. These transfer and balance functions are depicted in the graphic below.

Figure 3: Token Script

By limiting direct access to the balances, the script functions like a metaphorical sealed vault of a novel commodity, with a corresponding public ledger determining ownership over the contents to different parties. The vault allows users to query and alter the ledger only by the mechanisms controlled by a series of buttons on the vault's exterior. The buttons correspond to the functions that a given script makes available, and illustrate how even with a transparent vault, control over and ownership of the commodity ultimately lies in the mechanism underlying the buttons. Likewise, exchanges of a token can only be effectuated through the interface provided by its parent script. In the above example, this is captured by the transfer function.

While a secondary script may layer on supplementary terms of an exchange, the actual transfer of ownership is mediated through the script that maintains the balances variable for that asset—no matter the rules or rituals one constructs around the operation of the aforementioned vault, the ledger remains under the sole control of the button mechanisms. Any flaw in the mechanism is therefore propagated to all users of the asset.

Consider what would happen were the coder to forget the check ensuring that adequate funds were available. For one, a user could transfer more tokens than present in their balance, allowing them to accrue more tokens, i.e., cryptoassets, than they would otherwise be entitled. Additionally, any other script using the asset would also be affected as all token transfers are mediated by this contract that maintains and controls the balance variable.

Of course, tokens are not merely technological artifacts. They take their value from social consensus: their holders must think they will eventually provide some utility (even if merely being trading instruments).¹³³ To obtain that consensus, tokens are typically described and marketed with natural language text, written by people who may, or may not, have coded the tokens' scripts. In previous work, we examined the ERC-20 tokens created as a part of initial coin offerings in 2017.¹³⁴

Such offerings, loosely modeled on initial public offerings, typically involve the exchange of bitcoin or another form or cryptocurrency for a set of rights embodied in a transactional script.¹³⁵ For example, an organization called Kik raised \$98M in 2017 by offering for sale some of 10 trillion "Kin" tokens it had created.¹³⁶ According to Kik's White Paper, thirty percent of the total sale proceeds were earmarked for "startup resources, technology, and a covenant to integrate with the Kin cryptocurrency and brand."¹³⁷ Kik could, and did, embed these promises in a transactional script. We found that on a variety of measures, Kik's marketing documents and code matched exactly.¹³⁸

Tokens—which are already of significant practical import—thus, pose at least two sorts of problems for jurists. First, what if the code

133. See Hasu, *Unpacking Bitcoin's Social Contract*, MEDIUM (Dec. 3, 2018), <https://medium.com/s/story/bitcoins-social-contract-1f8b05ee24a9> [<https://perma.cc/6WQG-J2K2>] (explaining the import of social contract theory, and thus social consensus, in the appraisal of the value of cryptocurrencies).

134. See Cohnsey et al., *supra* note 2, at 619–25.

135. *Id.* at 593–95.

136. Khari Johnson, *Kik Raises \$98 Million in Kin Cryptocurrency Token Sale*, VENTUREBEAT (Sept. 26, 2017, 8:07 AM), <https://venturebeat.com/2017/09/26/kik-raises-98-million-in-kin-cryptocurrency-token-sale> [<https://perma.cc/EYA9-TN27>].

137. Cohnsey et al., *supra* note 2, at 628.

138. *Id.* at 673 app.B. That is not to say Kik is in the clear. It remains enmeshed in a fight with the SEC about whether its tokens were securities.

itself is somehow flawed, meaning that their buyers receive something different than they expected? Second, what if the code fails to match the natural language promises that purport to describe it?¹³⁹

B. EXCHANGES

Sometime on November 15, 2018, someone placed a buy-offer for a token called “Free Coin” on the TokenStore decentralized cryptocurrency exchange at ten times the prevailing market rate.¹⁴⁰ This created a substantial arbitrage opportunity for an enterprising trader who subsequently purchased Free Coin at the market rate and resold it at the inflated rate, realizing a profit of 0.79ETH or \$267 USD, while paying a complexity fee of \$5.00.¹⁴¹ Wishing to ensure that the entire sequence of trades was completed, the second trader batched the transactions using a script guaranteed to complete both the buy and sell trades, or neither.¹⁴²

This series of events is normal on digital marketplaces that trade cryptocurrency. Such marketplaces today are a major source of liquidity for cryptoassets, and consequently the most practically important public face of the transactional script commercial ecosystem.¹⁴³ Indeed, they may be the only fora where it is obvious that transactional scripts are pragmatically important mechanisms of exchange. Cryptocurrency exchanges take multiple approaches to custody, settlement, and order matching, which we now explore.

Unlike the “Free Coin” trade, the majority of cryptocurrency trades currently occur on *centralized* exchanges.¹⁴⁴ Users send either fiat currency or cryptocurrencies to an account controlled by the exchange, and in return, the exchange promises to (and normally does)

139. See Complaint at 13, SEC v. REcoin Grp. Found., L.L.C., No. 17-CV-5725 (E.D.N.Y. Sept. 29, 2017) (alleging securities liability due to REcoin’s whitepaper making representations about charitable giving for which “there is no program code”).

140. Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach & Ari Juels, Flash Boys 2.0: Frontrunning, Transaction Re-ordering, and Consensus Instability in Decentralized Exchanges 1, 4 (Apr. 10, 2019) (unpublished manuscript), ARXIV: 1904.05234.

141. The fee was 113,265 gas at a price of 134 Gwei, or about \$5. *Id.* at 5.

142. *Id.*

143. See generally Andrea Pinna, Simona Ibba, Gavina Baralla, Roberto Toneli & Michel Marchesi, *A Massive Analysis of Ethereum Smart Contracts Empirical Study and Code Metrics*, 7 IEEE ACCESS 78,194, 78,202–06 (2019).

144. See Nathan Reiff, *What Are Centralized Cryptocurrency Exchanges*, INVESTOPEDIA (June 25, 2019), <https://www.investopedia.com/tech/what-are-centralized-cryptocurrency-exchanges> [<https://perma.cc/7D56-QPX5>] (“[Centralized cryptocurrency exchanges] are the most common means that investors use to buy and sell cryptocurrency holdings.”).

promptly transfer an equivalent amount of a requested asset.¹⁴⁵ Such trades almost always occur off-chain and are settled using the exchange's internal ledgers.¹⁴⁶ Exchanges do generally offer custodial "wallets" that store the keys to a user's cryptoassets for easy trading.

These centralized exchanges could, but in most cases do not, use transactional scripts.¹⁴⁷ Rather, they act as market makers, facilitating trades between two parties for whom it is the custodian. There is thus, good reason to assume that trades between such parties are governed by ordinary contracts law.

The trend, however, is towards decentralized models for cryptocurrency exchanges.¹⁴⁸ Collapses of centralized exchanges have left users without access to balances stored on the exchange.¹⁴⁹ Centralized exchanges have also suffered for lack of liquidity across rarer asset types, as exchanges compete for order flow.¹⁵⁰ These flaws, combined with the blockchain community's ideological opposition to centralization, provided fertile ground for the development of decentralized exchanges,¹⁵¹ or DEXes.

TokenStore, a DEX, did not automatically match trades in their order book. Rather, sellers of assets would post an order and buyers would digitally sign their intent to match the order, forwarding the transaction to the DEX's on-chain contract, which provided for a 0.3% payment allotted to TokenStore. This is a common setup for DEX systems today.¹⁵² The service provided in this instance can be viewed as

145. *See id.* (explaining the idea of centralization).

146. Note these are not ledgers in the blockchain chain but merely the exchange's own internal record keeping mechanisms.

147. *Cf. id.* (noting that a crucial difference between centralized and decentralized exchanges is the use of intermediaries and custodians).

148. While trades on decentralized exchanges are still of comparatively low volume, the bulk of new exchanges are adopting decentralized models.

149. Mt. Gox, notable for its 2014 collapse, handled as much as 70% of all Bitcoin transactions prior to its demise. *See* Josh Constine, *The Plot to Revive Mt. Gox and Repay Victims' Bitcoin*, TECHCRUNCH (Feb. 6, 2019, 8:00 PM), <https://techcrunch.com/2019/02/06/the-plot-to-revive-mt-gox-and-repay-victims-bitcoin> [<https://perma.cc/79JU-49KX>].

150. *See* Nathan Sexer, *State of Decentralized Exchanges, 2018*, CONSENSYS MEDIA (Jan. 31, 2018), <https://media.consensys.net/state-of-decentralized-exchanges-2018-276dad340c79> [<https://perma.cc/94KS-AJLT>].

151. Despite the allure, the term "decentralized exchange" was long somewhat of a misnomer, as earlier generations still relied on a single contract for market making and settlement.

152. An explanation of similar architectures and the costs they impose is provided by Iddo Bentov, Lorenz Breidenbach, Phil Daian, Ari Juels, Yunqi Li & Xueyuan Zhao, *The Cost of Decentralization in 0x and EtherDelta*, HACKING, DISTRIBUTED (Aug. 13, 2017,

two separable components:¹⁵³ a listing service for open orders, and a platform to automatically consummate signed trades. Each component is mediated by a different piece of software. The order book is maintained by a centralized database, and interaction with it is by a website and its accompanying interfaces, both controlled by TokenStore.

The consummation component is managed by a minimal transactional script. To perform a trade, users first place cryptoassets in the custody of the transactional script using a deposit function. A trader wishing to match an order then uses the web interface to generate a signed transaction and submits it to the `trade()` function in the transactional script. The script performs a number of checks to ensure that the trade is valid and then updates the balances of both users.

In the case of our poor trader, once the trade was entered into the order book, any counterparty could force its execution. TokenStore appears to lack either a whitepaper or any substantive formalization in natural language of their system architecture. Its website contains no link to a terms and conditions, and even their medium blog is sparsely populated with only eight posts, with fully half being merely launch announcements.

The natural language content surrounding TokenStore is limited to a handful of tweets, blog posts, commit messages, and code comments, which yield only modest insight into the purported offering. A medium post entitled “Advantages of token.store ETH—Summarized” touts the security of the platform claiming “token.store doesn’t hold any of your funds; the trader deposits his funds into a smart contract This makes the experience safe and secure by its very nature” and directs users to a link to “check the code.”¹⁵⁴ On Twitter, the project bragged that “[f]unds at [http://]token.store ETH and EOS are

1:45 PM), <http://hackingdistributed.com/2017/08/13/cost-of-decent> [https://perma.cc/U79S-XE86].

153. This separation is justified by noting that a DEX could build a platform compatible with a rival’s signed transactions, using its own contract to consummate trades. TokenStore itself appears to have added support for transactions originating from “0x,” another popular DEX. TokenDev, *Tokenstore/Contract*, GITHUB (July 24, 2018), <https://github.com/tokenstore/contract/pull/11> [https://perma.cc/34HB-8YW6] (announcing an expansion allowing users to take 0x orders).

154. Harry Birch, *Advantages of Token.Store ETH—Summarized*, MEDIUM (Aug. 27, 2018), <https://medium.com/token-store/advantages-of-the-token-store-summarized-9164c4bab41> [https://perma.cc/KF8F-HX8H].

held in smart contracts: only users who hold the private key to the wallet which deposited them can withdraw them.”¹⁵⁵

Of direct relevance to trading error, the code controlling trade execution is preceded by the following code comment:

Note: Order creation happens off-chain but the orders are signed by creators,
// we validate the contents and the creator address in the logic below.¹⁵⁶

The notion that TokenStore “validates” the contents of an order prior to fulfillment leaves open the question of what validation a non-code-reading user ought to expect. We return to these issues in Section III.

C. ORACLES

In their default setting, transactional scripts are unable to interact with events occurring or data outside of the blockchain. Consider a transactional script that pays a shipper so long as the temperature within the shipping container stays below a certain threshold. While the logic is simple (if the temperature never went above X, pay the contract price), the quandary for the coder is how to ensure that the transactional script knows what the temperature inside the container is. This problem is solved using an “oracle,” a computerized agent that periodically submits the needed external data to the blockchain to be used within contracts.¹⁵⁷

155. token.store (@TokenDotStore), TWITTER (June 22, 2019, 11:32 AM), <https://twitter.com/TokenDotStore/status/1142470315777363968> [<https://perma.cc/TY3W-G64V>]. Of note, the script is upgradable via an opt-in process, meaning that while TokenStore currently has no way to access user assets, a future version of the contract certainly could. Though TokenStore has delisted a number of tokens from its order book and user interface, individuals who have delisted tokens stored in the transactional script can still access them by manually submitting a withdraw transaction to the blockchain. See token.store (@TokenDotStore), TWITTER (Nov. 15, 2018, 11:46 AM), <https://twitter.com/TokenDotStore/status/1063126115290615808> [<https://perma.cc/E4VK-DK4T>].

156. *Block Explorer for Ethereum Mainnet*, ANYBLOCK ANALYTICS, <https://explorer.anyblock.tools/ethereum/ethereum/mainnet/address/0xE17dBB844Ba602E189889D941D1297184ce63664> [<https://perma.cc/QB8N-QDJN>] (last modified July 6, 2020).

157. Incorporating external data directly onto the chain undermines the consensus mechanism, as it leaves no principled way for network participants without direct access to the external data to validate its correctness. See Alexander Egberts, *The Oracle Problem: An Analysis of How Blockchain Oracles Undermine the Advantages of Decentralized Ledger Systems* 5 (Dec. 12, 2017) (M.A. thesis, EBS University of Business and Law), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3382343 [<https://perma.cc/JW5E-UHJP>].

An oracle normally consists of two parts: a script and a way to populate the script's data store.¹⁵⁸ The figure below excerpts from the oracle for a currency exchange called Synthetix, showing a portion of the function that incorporates updates to currency exchange rates onto the script's storage.

Figure 4: Oracle Script from Synthetix

¹⁵⁸ Oracles can exist in software or hardware, with the latter purportedly offering some security guarantees stemming from the increased difficulty of compromising hardware.

As shown above, to execute the update function, the program must provide as input the data it wishes to incorporate. Notably, with each update a gas fee must be paid to store the new data. Oracles are thus not immune from the challenges associated with the complexity tax.

The data from oracles is typically provided by a third-party. This highlights a significant trade-off: oracles require trust in the data source (and further in the soundness of the oracle script).¹⁵⁹ For some, the intermediation and trust reintroduced by oracles highlights the “oracle paradox”: the more relevant the oracle data, the less one should be willing to trust the provider.¹⁶⁰ There are a number of ingenious protocol designs that achieve to tweak these trade-offs.¹⁶¹

But oracles break down despite the best intentions of all involved. For example, on June 25, 2019, one of the commercial data feeds from which Synthetix receives its USD/KRW (Korean Won) exchange rate “began to intermittently report a price 1000x higher than the current rate.”¹⁶² An earlier outage had taken one of the other two data sources offline, and the code to disregard outliers averaged the remaining feeds was also buggy. For a script operating a market for synthetic cryptoassets, tied to real world assets, a mismatch between the synthetic asset and the underlying asset was the worst-case scenario.

159. Federated oracles attempt to tackle this trust problem by distributing it: they require that a small number of parties independently provide the same information. The transactional script validates that the information it received is consistent across each parties’ submission before it permits the data to be used. Such an approach is no guarantee that a series of bugs won’t cause the overall failure of the system.

160. See Jesus Rodriguez, *The Middleman of Trust: The Oracle Paradox and Five Protocols that Can Bring External Data into the . . .*, HACKERNOON (July 31, 2018), <https://hackernoon.com/the-middleman-of-trust-the-oracle-paradox-and-five-protocols-that-can-bring-external-data-into-the-df39b63e92ae> [https://perma.cc/N3SJ-FFJG].

161. Decentralized oracle protocols (such as Chainlink) provide networks of individual oracles along with incentive structures that promote trust in the network. See, e.g., *Welcome to Chainlink*, CHAINLINK DEVS., <https://web.archive.org/web/20191228052855/https://docs.chain.link/docs/welcome-to-chainlink> [https://perma.cc/6XZB-VB8S]. An alternative solution operated by Augur maintains a prediction market which syncs to off-chain events by maintaining a fee pool paid out to participants that report on the state of the off-chain world. See Jack Peterson, Joseph Krug, Micah Zoltu, Austin Williams & Stephanie Alexander, *Augur: A Decentralized Oracle and Prediction Market Platform* (Feb. 3, 2018) (unpublished manuscript), ARXIV: 1501.01042. Amusingly, a participant can disrupt the market by placing a self-referential bet on Augur itself.

162. *Synthetix Response to Oracle Incident*, SYNTHETIX (June 25, 2019), <https://blog.synthetix.io/response-to-oracle-incident> [https://perma.cc/KEQ9-CKTP].

A bot was able to detect the mismatched exchange rate and took advantage of the arbitrage opportunity, acquiring tokens with a market value of \$1b with a mere \$1000 investment.¹⁶³ Fortunately for Synthetix, the bot owner reversed the trades in return for a bug bounty payout. The chain indicates the reversal was achieved through a subsequent trade at an exchange rate between an sETH (tokenized ETH) and sKRW (tokenized KRW) many orders of magnitude below the actual rate.¹⁶⁴ This suggests cooperation between Synthetix (which would have been able to manually adjust the exchange rate) and the bot owner who willingly converted their sETH below market rate. While the on-chain transactions suggest that the bot owner lost their investment (along with their profits) in the reverse transaction, this was likely remediated through a side payout.

We previously showed a script by which Synthetix incorporated various exchange rates. As with our previous examples, that script was accompanied by natural language text, which provided the opportunity for gaps between intention and outcome. The top of the file included the following description of its function:

[This is a] contract that any other contract in the Synthetix system can query for the current market value of various assets, including crypto assets as well as various fiat assets.

This contract assumes that rate updates will completely update all rates to their current values. If a rate shock happens on a single asset, the oracle will still push updated rates for all other assets.¹⁶⁵

Alongside extensive code commentary, Synthetix makes representations as to the functionality of its platform in the README accompanying the code, in the help section of its webpage, and in its marketing materials.¹⁶⁶ None of those commentary documents references the possibility of an error at the data source. The README file,

163. For information on the exchange of ETH used to take advantage of the arbitrage opportunity, see *Transaction Details*, ETHERSCAN (June 17, 2019, 9:04 AM), <https://etherscan.io/tx/0x6f6ee43ee07013503df786532493a3c405465f91e3ce8bb4ba8717a715db1caa> [<https://perma.cc/3B39-V7X3>].

164. This was deduced by following the chain of transactions with an exchange of 37m sETH for 362 sKRW. *Transaction Details*, ETHERSCAN (June 25, 2019, 2:54 AM), <https://etherscan.io/tx/0xc3fc19c63e1090eb624212bad71a27cd3dc7afcd0cf9063d24bfc47b5d036ae2> [<https://perma.cc/R52W-HKZF>].

165. *Contract*, ETHERSCAN, <https://etherscan.io/address/0x70c629875dadbe702489a5e1e3baae60e38924fa#code> [<https://perma.cc/VX3Y-56FR>].

166. The webpage does include a link to terms and conditions, but they govern use of the website rather than the blockchain platform. *Terms of Use*, SYNTHETIX, <https://www.synthetix.io/terms-of-use> [<https://perma.cc/8GCH-9BAK>] (“By accessing the website at <http://synthetix.io>, you are agreeing to be bound by these terms In no event shall Synthetix . . . be liable for any damages . . . arising out of the use or inability to use the materials on Synthetix’s website . . .”).

however, notes that the fees are governed by an exchange rate that is “derived by looking at a basket aggregate of currencies,” that the rates will be not be “stale” and that the oracle will be “trusted.”¹⁶⁷ And the MIT License for the Synthetix software, which is provided as a file on its Github page, states that the software is provided “AS IS.”¹⁶⁸ How to compile these various statements is the subject of our next Section.

III. SCRIPTS AND STACKS

As we explored above, it simply is not practical to create scripts that perfectly embody their coders’ intent. The point is generally true for all coded exchange. A 2019 decision issued by the Singapore International Commercial Court, applying the common law of Singapore (itself derived from English common law), offers a unique window into how jurists might resolve the contractual consequences of code gone awry.¹⁶⁹

The case involved a lopsided trade of cryptocurrency. Quoine operated a centralized currency exchange platform that primarily enabled trading of cryptocurrencies.¹⁷⁰ B2C2 was an “electronic market marker,” which had developed a trading algorithm written by its president several years before the events of the case.¹⁷¹ In April 2017, Quoine, as it always did, had a program in place to monitor users who had borrowed collateral with which to trade.¹⁷² When Quoine’s program identified an imbalance in the reserves, it forced the sale of collateralized assets at the best available price.¹⁷³

Unfortunately, the program that Quoine had written to ensure a liquid market temporarily failed.¹⁷⁴ The result was its prices were out of sync with the global market.¹⁷⁵ B2C2, whose goal was to capture returns from the bid-ask spread,¹⁷⁶ offered to fill seven particular orders

167. *Oikos Contracts*, GITHUB, <https://github.com/oikos-cash/oikos#oikos-tron-contracts> [<https://perma.cc/9CME-PC2E>] (last modified Aug. 11, 2020).

168. *MIT License*, GITHUB (Nov. 29, 2018), <https://github.com/Synthetixio/synthetix/blob/8f3b95d1205f2b4d6b62124bd07f593773800743/LICENSE> [<https://perma.cc/Q2PX-46T7>].

169. *B2C2 Ltd. v. Quoine Pte. Ltd.* [2019] SGHC (I) 03.

170. *Id.* at 1.

171. *Id.*

172. *Id.* at 12.

173. *Id.*

174. *Id.*

175. *Id.*

176. *Id.* at 5.

at a price around 250 times that available on the broader market, making it a quick profit of several million dollars.¹⁷⁷

The next day, Quoine reversed the trades in its order book.¹⁷⁸ B2C2 sued, arguing that the reversal violated Quoine's terms and conditions, which provided that "once an order is filled, you are notified via the Platform and such an action is irreversible."¹⁷⁹

Quoine offered two principal defenses to its supposed obligation to complete the trades with B2C2.¹⁸⁰

First, it pointed out that its risk disclosure statement, uploaded prior to the trade but not expressly incorporated into the terms and conditions, permitted a reversal of trades if "market circumstances shift dramatically or something else happens."¹⁸¹ The court held that because the statement was not expressly incorporated into the contract (that is, into the terms and conditions), it did not override the express language of the contract itself.¹⁸² This decision seems perfectly sensible on its face—an application of the usual idea that an integrated agreement ought not be contradicted by extrinsic statements.¹⁸³

Second, and more interestingly, the court considered whether the facts established a *mistake*, allegedly because the parties would not themselves have executed the trade in a hypothetical world where they talked about it in person in real time.¹⁸⁴ That is, the parties' agents—the programs—had executed a deal which, though an accurate expression of the code's instructions, somehow failed to capture their "real intent."¹⁸⁵ As Quoine argued, the platforms were "really complex platforms," in which "a lot of things [could] go wrong."¹⁸⁶ The mistake in question was not a typographical error in the code, but rather an "oversight in the design of the system."¹⁸⁷ Quoine asked the

177. *Id.* at 12.

178. *Id.*

179. *Id.* at 55.

180. *Id.* at 58.

181. *Id.* at 64.

182. *Id.* at 73.

183. For this point, the tribunal cited a British treatise and supporting Singaporean authorities, but the result would be no different in most U.S. jurisdictions. *See id.* at 69; *infra* text accompanying notes 271–72.

184. *B2C2 Ltd.* [2019] SGHC (I) 03, at 75.

185. *See generally* 1 TIMOTHY MURRAY, CORBIN ON CONTRACTS: FORMATION OF CONTRACTS § 4.11 (rev. ed. 2020) (describing old cases on telegraphs and mistakes in transmission of messages).

186. *B2C2 Ltd.* [2019] SGHC (I) 03, at 30.

187. *Id.*

court to apply a “pragmatic and judicious stance” and void a “clearly erroneous trade.”¹⁸⁸

The tribunal, adopting a narrower reading of unilateral mistake derived from British common law, held that to avoid liability, the person who was not mistaken must have actually known of her counterparty’s error and “was shutting her mind to the obvious.”¹⁸⁹ In deciding whether B2C2 had that requisite state of mind, the court noted conceptual difficulties:

[A]pplying the law to . . . algorithmic trading . . . raise[s] new questions. What mistakes have been made and to what extent are they fundamental? How does one assess knowledge or intention when the whole operation is carried out by computers acting as programmed? Whose knowledge is relevant? At what date is this knowledge to be assessed?¹⁹⁰

Given that the program executed a fixed trade—i.e., the design of the code bespoke the programmer’s intention to strip it of any discretion—the court analogized the program to a “mere machine carrying out actions which in another age would have been carried out by a suitably trained human. [It is] no different [from] . . . a kitchen blender relieving a cook of the manual act of mixing ingredients.”¹⁹¹

The court thus held that for the purposes of ascertaining mistake, it would focus on the intent of “the person on whose behalf the computer placed the order in question [B2C2].”¹⁹² This, the court explained, required examining the “state of mind of the programmer of the software of that program at the time the relevant part of the program was written.”¹⁹³ Assessing the relevant evidence, the court decided that the programmer had not inserted code intending to trade at lopsided rates or knowing that doing so could only result from the counterparties’ error.¹⁹⁴ Thus, it rejected the claim of mistake.¹⁹⁵

Quoine is not a case about a transactional script gone wrong. But it does offer a few glimpses into the future of how scripted exchanges will be resolved, or at least one possible approach to such problems.

188. *Id.* at 79.

189. *Id.* at 80; *cf.* RESTATEMENT (SECOND) OF CONTS. § 153 (AM. L. INST. 1981) (requiring the non-mistaken party to have “had reason to know of the mistake”).

190. *B2C2 Ltd.* [2019] SGHC (I) 03, at 86.

191. *Id.* at 89.

192. *Id.* at 87–88.

193. *Id.* at 89–90.

194. *Id.* at 99.

195. *Id.* In 2020, the Singapore Supreme Court affirmed the judgment in an opinion that was equally lengthy but broke no new conceptual ground. *Quoine Pte. Ltd. v. B2C2 Ltd.* [2020] SGCA (I) 02, [https://www.supremecourt.gov.sg/docs/default-source/module-document/judgement/-2020-sgca\(i\)-02-\(v-4\)-pdf](https://www.supremecourt.gov.sg/docs/default-source/module-document/judgement/-2020-sgca(i)-02-(v-4)-pdf) [<https://perma.cc/7LF2-HY2V>].

The first is the importance of the role that old-fashioned contract law—and old-fashioned contracts—will play in the disposition of the parties’ legal rights. The *Quoine* tribunal privileged the natural language contract embodied in the terms and conditions over the code. Though it ultimately declared that the risk disclosures outside the terms and conditions were not binding, it apparently would have enforced those disclosures, had they been incorporated, over software code that enabled the trade. It was only in the absence of contract terms governing the deal that the court turned to what the code permitted, and why. Thus, the coded rules of exchange were, in the court’s view, largely irrelevant.

Second and relatedly, the *Quoine* court’s focus on the intent of the programmer is a natural outgrowth of contract caselaw. Interestingly, *Quoine* focused on intent at the time of the original programming. Knowing what B2C2 intended required the court only to take the testimony of one coder, and to apply the normal judicial tools (assessing demeanor, consistency with prior statements and the documentary record) to determine the “truth” of the coder’s intent.¹⁹⁶ Such a simple story is unlikely to replicate when we interrogate more complex coding problems.

We think the *Quoine* decision offers a rough guide to the sorts of problems that transactional scripts will raise, and a sense of how common law courts will be motivated to resolve them. Indeed, we think the tribunal got it mostly right in context, but that its reasons won’t scale. Given what we’ve learned about the technological environment generating scripts, we make two basic arguments about how the law ought to consider problems related to scripting.

First, judges should not privilege “contract” over “code” but rather ought to ascertain and harmonize meaning across a *contract stack*. Code—when read with its natural language comments and commit logs—has communicative meaning that courts should seek to ascertain and enforce. Second, conditional on integrating the stack, the search for meaning should focus on expressions that best reflect the best public-facing account of the parties’ shared intent at the time that they committed to the deal. We work through these principles by suggesting a list of canons of scripted law.

196. *B2C2 Ltd.* [2019] SGHC (I) 03, at 89–90.

A. THE CANONICAL STACK

As Jason Allen has recently argued, transactional scripts are the latest in a series of “‘contractware’, i.e., technological artefacts designed to embody and perform contracts.”¹⁹⁷ A chip in a credit card embodies the concept well. There is a natural language credit agreement between you and your card company, updated and modified at the issuers’ will against the background of regulation, which define the circumstances under which credit may be extended. Those terms parallel ones agreed to between the issuer (or its agents) and merchants, which enable your card (the artifact resulting from your natural language agreements) to effectuate a pending sale by insertion into the merchant’s reader. That is, the card/reader performs contracts that hover in the air around them.

All contractware has this property: it is wrapped within an ordinary, legal contract. Absent a legitimate connection to those contracts, use of a credit card can still affect the world—your account will be debited, the merchant’s credited—but such changes can be quickly reversed. Of course, not all contractware requires physical manifestations. Allen argues that transactional scripts are an example of contractware in which “the subject matter of the contract is an immaterial object which can be manipulated directly by the [code].”¹⁹⁸

Allen concludes that transactional scripts are layers in the “contract stack.”¹⁹⁹ The stack is a useful metaphor to describe the various elements of “contract as a complex entity.”²⁰⁰ A stack might include, perhaps, an oral contract (or other indicia of social agreement), a written text, and the legal rules, which give effect to the relationship between the two. In a transactional script, the “written text,” i.e. natural language terms like the statement of terms and conditions in *Quoine*, is “complemented (or supplanted) by code which is also, incidentally, wholly or partially executable by a machine.”²⁰¹ Each additional step—the compiling of machine-readable code from human readable code—adds another layer of complexity to the stack, but the stack is intended to operate, and therefore ought to be read, as a whole.²⁰²

Many working in this space acknowledge that contract stacks will be the mechanism through which commercial projects will deliver

197. Allen, *supra* note 33, at 313.

198. *Id.* at 318.

199. *Id.* at 330.

200. *Id.*

201. *Id.*

202. *Id.* at 331.

scripted performance.²⁰³ This is the apparent impetus, for example, for the International Swaps and Derivative Association's attempt to insert explicit references to scripts, which operationalize interest payments into the ISDA master contract.²⁰⁴ But other sophisticated projects, like the OpenLaw cooperative, are explicitly developing stacked deals.²⁰⁵

But we'd go further. *All litigated scripts will exist in contract stacks.* That's not to say that there will always be a 100-page master agreement, or even clicked-through terms and conditions of sale. Rather, our claim is that there will *always* be *some* non-code statements—ranging from the highly formalized terms and conditions, to less formalized white papers and code commentary, to quite informal promises (Twitter)—which will inform tribunals' understandings of what the parties intended to exchange.

This is true both as a matter of practice, and a matter of logic. Code is not self-descriptive. Any scripts that have practical relevance will have some non-code language surrounding them. Any set of communications relevant to the exchange that are visible to the parties are at least presumptively a part of the stack. (Whether they all count equally is a harder question.)

Thus, although some in the literature have asked if a "smart contract" is really a contract, standing alone, we doubt the practical relevance of that question. The code standing alone may not fully specify an executory contract, unilateral or otherwise, because it simply accomplishes performance. But, unlike the cheese, at least at the moment, the code never stands alone. As our examples in Part II illustrate, it is typically deployed in a social context.

Though the *stack* is a relatively recent term, the idea that transactions are accompanied by many sorts of potentially legally-operative promises is not. You buy a car, led to the dealership by a commercial, and see a price at the entrance to the lot. The dealer makes representations to you. You reply with your own admissions. Finally, you sign a written agreement, which often seeks to exclude the prior

203. See, e.g., DE FILIPPI & WRIGHT, *supra* note 33, at 77 (arguing that "hybrid" contracts will be useful especially when particular components cannot be reduced to code).

204. INT'L SWAPS & DERIVATIVES ASS'N, LEGAL GUIDELINES FOR SMART DERIVATIVES CONTRACTS: THE ISDA MASTER AGREEMENT 7 (2019), <https://www.isda.org/2019/02/19/legal-guidelines-for-smart-derivatives-contracts-the-isda-master-agreement> [<https://perma.cc/HSN4-2A54>].

205. See *General Questions*, OPENLAW, https://app.openlaw.io/faq#first_draft_automate [<https://perma.cc/KZ43-6LMC>] ("On OpenLaw, you can execute smart contract code by embedding a smart contract call in any template.").

representations as outside of the operative deal. The whole exchange is a contract stack. Some of it is legally operative, and some is not. When you buy a cup of coffee at the store after seeing a sign at the door, the process is the same—though if you use an app to fulfill the purchase, some of the code that accomplishes payment is obscure to you and is, perhaps, not part of the contract. The point is that what counts as the “contract” results from the multifarious set of policy choices that we call contract doctrine. Sometimes that construction maps onto a single, written, document, but often does not.

We thus suggest our first canon—a maxim—to help courts make sense of the adjudication of disputes arising from scripted exchange. Like all of the rules that we suggest, it builds on existing caselaw. It is also a default: the parties can (and as we explore below) have changed it by contract.

Canon 1: All shared communications of intent, including the code, comprise the legally-operative stack.

In working through this canon, an example may help. Consider a token white paper makes a promise about governance rights, but the script contains no reference to that promise.²⁰⁶ That was the norm in the 2017 ICO craze.²⁰⁷ In previous work, we conducted an exhaustive audit of the match between semantic disclosures (including those in White Papers, Twitter, Instagram, Reddit, and Medium) and scripted code in tokens. We found that most programs had not, in fact, created code that conformed to their promises.²⁰⁸ “For over 20% of ICOs in our sample where promoters promised cryptoasset supply restrictions, and 35% of promised token burning, we could not observe corresponding restrictions [written into transactional scripts].”²⁰⁹ Worse, we did not find code-based vesting restrictions in twenty-five of the thirty-six ICOs where promoters promised to adhere to such restrictions.²¹⁰ Finally, of twelve ICOs for which our audit revealed that a central party could modify the functionality of the cryptoasset’s code, only four disclosed that ability in their promotional materials.²¹¹

206. Cf. Cohnney et al., *supra* note 2, at 640–44 (discussing market solutions to missing disclosures).

207. *Id.* at 640.

208. *Id.* at 639.

209. *Id.* at 640.

210. *Id.*

211. *Id.* at 639–40.

To the extent that an action for a legal breach of such a contract stack is brought, how do we know what was promised? Some might argue that the only legally operative promises are those made in natural language documents outside of the code; others, that the code provides the only relevant set of rules.²¹² In a way, this is quite similar to the start of a conventional parol evidence problem where multiple documents are candidates for inclusion as the litigated “contract.”²¹³ When the bounds of contract are fuzzy—that is, in the absence of integration—the stack of contractware is capaciously constructed. This problem is well illustrated by the conventional 2-207 problem, where the parties have exchanged non-matching forms and the court must compile a “contract” *ex post*.²¹⁴

Courts appropriately compile into the stack those promissory statements, which seem to indicate transactional intent. Thus, terms and conditions, because they clearly indicate the parties’ intent to contract, are almost certainly part of the stack. So too are most published white papers, which make numerous promises about what’s to be delivered, even though in some ways they are offers directed to the world. The reason is obvious: white papers are the best evidence of what the code supplier intends to create, and the way that counter-

212. Some have argued that the parol evidence rule would make it difficult “for the parties to prove their intent to contract by pointing to other circumstances, such as prior dealings or negotiations.” Anna Duke, *What Does the CISG Have to Say About Smart Contracts: A Legal Analysis*, 20 CHI. J. INT’L L. 141, 159 (2019). This is extremely puzzling, as the rule itself only applies to fully integrated agreements. Others suggest that when the parties have “signed and verified that the contract had been accurately translated into computer code,” it will be difficult for them to later argue that there were additional terms. Alan Cohn, Travis West & Chelsea Parker, *Smart After All: Blockchain, Smart Contracts, Parametric Insurance, and Smart Energy Grids*, 1 GEO. L. TECH. REV. 273, 281 (2017). This too is confusing, since the typical way that parties would verify that the contract is correct would be to agree to that stipulation in a natural language agreement.

213. See Gregory Klass, *Parol Evidence Rules and the Mechanics of Choice*, 20 THEORETICAL INQUIRIES L. 457, 464 (2019) (“[A] writing is integrated if and only if the parties together intended that it would serve as a final statement of some or all terms of their agreement.”).

214. See John D. Wladis, *The Contract Formation Sections of the Proposed Revisions to U.C.C. Article 2*, 54 SMU L. REV. 997, 1011 (2001) (“[T]he exchange of non-matching records rarely involves parol evidence issues because there is usually no one record that is a final expression of the parties’ agreement.”).

parties are enticed to invest, contribute funds, or otherwise transact.²¹⁵ It might be (as we will explore) that terms and conditions provide better evidence of intent than do white papers, when both are present, but as a starting point, the white paper ought to be considered as part of the set.

More transient promises, like those made on social media posts, on Reddit, or via video, are also plausible candidates for the stack. Jurists will ask if they make promises sufficiently definite and certain to enable a fact finder to generate the grist for obligation.²¹⁶ This will turn on the nature of the promises made.²¹⁷ Does a post on Reddit stating the precise amount of a capped supply count?²¹⁸ We think so. But one that states that “people will have an extended period during which they can burn” without more detail seems to be too vague to generate obligation.²¹⁹

The hardest problem is the contractual status of the script, including its commentary. It’s hornbook law that the parties can use ciphers to express themselves, and that courts ought to enforce such coded meanings “however we may marvel at the caprice.”²²⁰ In deciding whether to give effect to private meanings, courts traditionally engage in a hypothetical inquiry: *if* the parties had been queried at the moment of contracting about the meaning of a particular term, what would they have jointly said? Courts will thus self-consciously adopt

215. See MURRAY, *supra* note 185, § 2.12[1] (describing a “webpage whose contractual nature is not obvious” as one where courts will ask if “the recipient actually knows or has reason to know of it [when] she assents to it”).

216. See *id.* § 1.11 (defining an offer as “an act whereby one person gives to another the legal power of creating the relation called contract”).

217. For examples of the sorts of marketing promises made in ICO disclosures, see Shaanan Cohny, David Hoffman, Jeremy Sklaroff & David Wishnick, *Coin-Operated Capitalism Appendix C*, COLUM. L. REV., <https://columbialawreview.org/content/coin-operated-capitalism-appendix-c> [<https://perma.cc/QGE7-H7AF>].

218. See, e.g., u/helloicon, *ICON Technical Q&A Summary*, REDDIT (Sept. 18, 2017, 1:06 AM), https://www.reddit.com/r/helloicon/comments/70t56h/icon_technical_qa_summary [<https://perma.cc/5XJZ-APCB>] (answering questions regarding the technicalities related to ICON).

219. See aetrnty, *Burning Token*, REDDIT (Sept. 10, 2017, 1:42 PM), http://www.reddit.com/r/Aeternity/comments/6za07b/burning_token [<https://perma.cc/K3QQ-PJBL>].

220. 5 MARGARET N. KNIFFIN, CORBIN ON CONTRACTS: INTERPRETATION OF CONTRACTS § 24.9 (Joseph M. Perillo ed., rev. ed. 1998); see, e.g., *Hurst v. W.J. Lake & Co.*, 16 P.2d 627, 629 (Or. 1932) (holding the term “minimum 50%” to encompass “as low as 49.5%”); *Smith v. Wilson* (1832) 110 Eng. Rep. 266, 266; 3 B. & Ad. 728, 728 (holding that “parol evidence was admissible to sh[ow] that . . . the word thousand, as applied to [the contract], denoted twelve hundred”).

the parties' expressed "vernacular."²²¹ Unilateral and uncommunicated meaning bear as little on the problem of contractual interpretation as the public meanings that the parties meant to cast aside.²²²

A traditional requirement about non-traditional communications is that the party against whom they would operate reasonably understands them to be contractual in nature.²²³ Context matters: the more impermanent the medium, the less obvious it should be that a bargain results from the posting.²²⁴ This is why, in an old case, a court found that the language on the back of coat check tickets was unenforceable: such scraps "did not arise to the dignity of a contract."²²⁵ Promises made in terms and conditions feel like contracts (more or less), while those made on Twitter may not.²²⁶

But how about the code? The primary objection to including scripts as presumptive source of contractual intent is that code *does*, it makes no *promises to do*. True, those in the know can learn the coder's sophistication from her script's elegance and economy, just as connoisseurs of font will make inferences based on this Article's typesetting. But that's not the same as the sort of promissory communication that normally generates obligation. Moreover, because the code is inherently buggy—as we have explored—including it into the stack necessarily means that courts are adopting ambiguous evidence of what the parties "really" mean to accomplish.

Of course, natural language is buggy too, as centuries of experiences with contract interpretation problems make clear. And the code is festooned with natural language comments and commit logs, which do, implicitly or explicitly, state what the coders are trying to accomplish. Most importantly, in light of the social conventions of the script-

221. KNIFFIN, *supra* note 220, § 24.13.

222. *Id.* § 24.6.

223. MURRAY, *supra* note 185, § 2.12[1] (discussing non-contractual documents that are not enforced for lack of inquiry notice).

224. In some contexts, statements directed at the whole world will be deemed to be advertisements. *Id.* § 2.4. But since white papers and the like aren't considered in the abstract here, but rather when accompanying the scripts which put them into effect, the better analogy is the supermarket circular distributed at the grocery itself. *Id.* § 2.7.

225. Healy v. N.Y. Cent. & Hudson River R.R. Co., 153 A.D. 516, 519–20 (N.Y. App. Div. 1912).

226. See Berkson v. Gogo L.L.C., 97 F. Supp. 3d 359, 382 (2015) (arguing that electronic contracts "require clearer notice than do traditional retail"). *But cf.* Kristen Chiger, *When Tweets Get Real: Applying Traditional Contract Law Theories to the World of Social Media*, 3 ARIZ. STATE U. SPORTS & ENT. L.J. 1, 6 (2013) (finding that tweets can create contractual obligation).

ing industry today, it is reasonable to conclude that participants *believe* that the code-and-comments create obligations, i.e., that the script is contractual in nature. The very term “smart contracts” justifies this conclusion. The rhetoric of exchanging in *contracting* should be constitutive.

This conclusion is strengthened in light of our focus on public blockchains, where all of the relevant parties have access to the scripted rules, can inspect them, and can read the associated code commentary. By definition, we have examined scripts that are open to bilateral inspection. It would be as if we were considering a contract about the importation of chicken from abroad, between parties from different countries, that expressly incorporated a foreign language dictionary, and then denied that dictionary’s relevance in understanding the parties’ meaning. Course of performance is typically seen as a strong signal of meaning; the script is yet a clearer signal. It is the expressed performance itself, fixed contemporaneously with agreement. (If the code were not public, as we’ll discuss below, its relevance to discerning the parties’ agreement would be more obscure.)

Incentives provide a final reason to include code in the stack. If lawyers are on notice that the code has legal relevance, that it can create or destroy obligation, or help jurists to interpret them, they will begin to pay attention. The scripting industry at the moment is covered by few specific rules, because of uncertainty about the nature of the underlying asset type. Private regulation through lawyering—having lawyers work to understand exactly what the script says and inquire how it might go wrong—is a way to help to civilize this wild west. That is, bringing the script into the contract stack will motivate lawyers to care about coding, and coders, and consequently reduce the likelihood of promissory fraud.²²⁷

B. TENSIONS WITHIN THE STACK

What happens if there are differences in what’s promised between elements of the stack? Allen argued that questions of intent are particularly difficult when interpreting code, as the formal language might have meta-logics—internal and intentionally chosen goals (like, for example, compactness).²²⁸ He thus, suggests that courts may need to modify traditional canons of interpretation to think about how to best capture the parties’ meaning when working through layers of a

227. Some may worry that lawyers will seek to clean code of commentary altogether to reduce interpretative overhang. But it will be extremely difficult, if not impossible, to do so without leaving inculpatory evidentiary traces.

228. Allen, *supra* note 33, at 341.

stack.²²⁹ He gives no further details, and this next part elaborates on the problem.

Courts will often say that interpretation ought to make sense of the “contract as a whole,” that is, “the entire deed, and not merely upon disjointed parts of it.”²³⁰ That’s particularly true when the parties express their agreements in multiple documents.²³¹ In determining meaning when there is ambiguity, extrinsic evidence—that is, anything other than the contract itself—is commonly admitted.²³² This leads us naturally to a canon seeking harmonization.

Canon 2: Where Possible, Interpret the Stack to Harmonize Meaning.

We start by returning to the ICO example. At least until recently, it was common for white papers to make promises about governance in ICO “smart contracts,” but neglect to write actual scripts containing such coded rules.²³³ In our view, the stack as a whole ought to be read to be making a promise: the absence of protection in the code should not be dispositive. If, for example, the projects’ founders were to take assets in violation of a textual vesting promise, an action ought to lie for breach. That’s true even though an informed reader with an understanding of Solidity would have readily observed that the tokens contained no promises. (Most, in fact, were essentially the unmodified ERC-20 code.) The point is that a *legally* informed reader could believe those promises to be enforceable based on the white paper alone, and their repetition within the script unnecessary.

Conversely, when the natural language disclosures say nothing about the ability to modify rights described as having been created in

229. *Id.*

230. KNIFFIN, *supra* note 220, § 24.21 (quoting Blackstone).

231. See RESTATEMENT (SECOND) OF CONTS. § 202(2) (AM. L. INST. 1981) (“[A]ll writings that are part of the same transaction are interpreted together.”).

232. Note that New York courts, like many others in a modern formalist trend, take a different approach. See, e.g., *Gilbane Bldg. Co./TDX Constr. Corp. v. St. Paul Fire & Marine Ins. Co.*, 97 N.E.3d 711, 714 (N.Y. 2018) (holding that extrinsic evidence regarding intent should be admitted only when the parties’ agreement is unclear).

233. For example, the Monaco project promised that its supply would be capped in a transactional script: “The MCO smart contract will stop accepting commitments at 888,888ETH hard cap.” MONACO, MONACO WHITEPAPER 8 (2018), <https://whitepaperdatabase.com/wp-content/uploads/2018/03Monaco-MCO-Whitepaper.pdf> [<https://perma.cc/7GKB-9D2D>]. But our audit disclosed no such scripted commitment. See Cohn et al., *supra* note 2, app.B.

a token, but the script appears to permit modification, we might conclude that the contract stack permits unilateral changes.²³⁴ In the search to understand what the contract stack promises, neither code nor script ought to prevail over the other, again unless the parties otherwise indicate. The goal is to determine what the parties intended, expressed in whatever cipher they chose. We should not *a priori* dismiss rights provided in code.

The TokenStore problem offers another setting for the harmonization principle. Recall that TokenStore had a vestigial legal wrap in place: a handful of twitter and medium posts, making vague gestures about the exchange's commitment to being a hands-off-enterprise. The code permitted trader errors—indeed, it did nothing to prevent them. However, the code *commentary* stated that “we [i.e., TokenStore's operators] validate the contents and the creator address” of the “orders.”

It's not clear what the writers of this comment intended it to mean. In programming communities, validate can take on a narrow meaning—someone can enter an “input in a form that is not expected,” leading to “altered control flow, arbitrary control of a resource, or arbitrary code execution.”²³⁵ But it can also take on a broader meaning, i.e., that the code produces commercially reasonable results.²³⁶ And perhaps such narrow programming meanings are only relevant in commercial markets, so that “validate” ought to take on an ordinary meaning, i.e., “to make legally valid.” This would mean that the exchange bears the risks of obvious errors. Such a reading would be helpful to a trader's action for rescission based on mistake (which will turn in part on what the contract says about risk) as well as that for ordinary breach of contract. In our view, the correct approach again would seek to make sense of the gestalt project, treating the code and its natural language comments as guides to the parties' joint intent.

234. See Cohn et al., *supra* note 2, at 630–34. Of course, a court could find that, in light of an industry-wide practice of describing rights as “immutable,” silence in one layer of the stack should be interpreted against a commercial background denying modifiability. *Id.* at 615 n.114.

235. *CWE-20: Improper Input Validation*, COMMON WEAKNESS ENUMERATION (Sept. 19, 2019), <https://cwe.mitre.org/data/definitions/20.html> [<https://perma.cc/C3GA-FM6Y>].

236. See *Data Validation: OSWAP Guide to Building Secure Web Applications and Web Services*, COMMON OPEN WEB APPLICATION SEC. PROJECT (Jan. 5, 2006), <https://searchsoftwarequality.techtarget.com/news/1156594/Data-validation-Chapter-12-OWASP-Guide-to-Building-Secure-Web-Applications-and-Web-Services> [<https://perma.cc/P8GG-VG7F>] (“Business rules [e]nsure that data is not only validated, but business rule correct.”).

Harmonization becomes difficult when parts of a contract conflict.²³⁷ When pieces of a deal counterpose, courts traditionally first seek to ascertain the parties' principal purpose, and then to advance it by deciding out which pieces of evidence to privilege.²³⁸ For example, handwritten terms prevail over typewritten terms, and specially typed provisions control over pre-printed forms.²³⁹ Courts foreground those provisions that they believe are best indicators of what the parties "really" meant.²⁴⁰

Scripted exchange will sometimes also pose problems of inconsistent intent. Given that pieces of the stack are written at different times, by authors with distinct professional backgrounds, and intended for different readers, we should anticipate conflicts in meaning. As a default rule, we'd propose a hierarchy of meaning, which privileges natural language wrapping text over code when they conflict.

Canon 3: In cases of conflict, privilege natural language promises over coded ones: i.e., wrapping text, commitment messages and code commentary over code, high level code over byte code.

By wrapping text, we mean to include any text that is outside the code itself, but within the stack. When such text conflicts with code, we think the parties' contract—what they can sue on—most likely turns on their natural language expression.²⁴¹ This is a pragmatic choice.²⁴² As with the *Quoine* tribunal, most judges are going to have a natural affinity for text that they can read without the aid of an expert

237. Cf. Surden, *supra* note 32, at 657 ("A primary unresolved tension may occur in future scenarios where there is both a written and data-oriented representation of the same contractual expression, with interpretations that differ.").

238. See KNIFFIN, *supra* note 220, § 24.20 ("When the principal purpose of the parties becomes clear, further interpretation should be guided thereby.").

239. See generally *id.* § 24.23, at 236, 251 (discussing various rules that courts use to reconcile conflicting contract terms).

240. See *id.* § 24.9, at 59 ("A court's purpose in using extrinsic evidence to interpret a contract is discernment of the parties' intentions.").

241. See Rohr, *supra* note 33, at 85 ("[C]ode [that automates a larger agreement] is likely to be viewed as a component of performance that one party will attempt to prove is nonconforming.").

242. For a defense of a search-costs based theory of parol evidence and other rules that limit "idiosyncratic understandings," see Joshua A.T. Fairfield, *The Search Interest in Contract*, 92 IOWA L. REV. 1237, 1265–67 (2007).

translator.²⁴³ They will argue that “no one reads smart contracts.”²⁴⁴ (The fact that no one reads regular contracts is equally true, though it feels easier to blame them for it.)²⁴⁵ Of course, just like contract text, code *can* be “read,” and often results from a similar iterative drafting process as old-fashioned contracts.

Why, then, privilege English over Code? One reason sounds in the classic worries about opportunism and bad faith that drives many judicial treatments of adhesion contracts. If parties could avoid a contractual promise by negating it in code that you had no reason to think would be read, we would rightly worry about promissory fraud or other forms of bait-and-switch behavior. Though today, many users of transactional scripts are sophisticated—even to access scripts, you usually install specialized software on your computer—that may not be the case going forward. Given that English is easier to read than Solidity and other high-level programming languages, courts will likely privilege natural language promises wherever they can.

But at a deeper level, our intuition is that courts imagine they are looking for something they call “real” intent, which is really more like what the parties expressed about their intent to the world.²⁴⁶ Just as with other forms of commercial transactions drafted and entered in stages, discerning real intent is often a fool’s errand.²⁴⁷ That inquiry falsely implies that the parties gave the problem some thought. When courts speak about intent, they are engaging in a hypothetical and imaginative exercise, which entails significant degrees of analytic freedom. But imaginative exercises must be explainable in public judicial

243. In consumer-facing transactions, courts also may worry about exploitation when parties use language that is hard for adherents to understand. *See, e.g., Frostifresh Corp. v. Reynoso*, 274 N.Y.S.2d 757, 759 (N.Y. Dist. Ct. 1966) (finding that the adherents were “handicapped” by contract terms written “in a language foreign to them”).

244. Cohney et al., *supra* note 2, at 598.

245. *See* Tess Wilkinson-Ryan, *A Psychological Account of Consent to Fine Print*, 99 IOWA L. REV. 1745, 1751 (2014) (“[O]ne of the truisms of empirical contracts research is that ‘nobody reads.’”).

246. *See* Rohr, *supra* note 33, at 78 (discussing how courts can derive traditional contract formation concepts, like intent, from a vending machine transaction).

247. *See* Douglas G. Baird & Robert Weisberg, *Rules, Standards, and the Battle of the Forms: A Reassessment of § 2-207*, 68 VA. L. REV. 1217, 1219 (1982) (noting that the law cannot resolve the battle of the forms by inquiring into the parties’ intent); GRANT GILMORE, *THE DEATH OF CONTRACT* 46–47 (2d ed. 1995) (“If . . . ‘the actual state of the parties’ minds’ is relevant, then each litigated case must become an extended factual inquiry into what was ‘intended,’ ‘meant,’ ‘believed’ and so on.”).

opinions, and thus rely to a degree on text that can be read by the widest audience, and that is susceptible to the cheapest judicial oversight.²⁴⁸

As we've shown, code is irreducibly buggy, and the normal ways that coders handle error—by iterating better versions—may not translate well to scripted exchange. Code simply isn't a very straightforward way to express the parties' intent. By contrast, parties have had hundreds of years of experience contracting in English (or French, or Esperanto, turning on the court's and parties' native tongue). Unless there is good evidence that a particular line of code was made salient—for example, if it is referred to by line number in the natural language contract itself—courts should conclude that text trumps code.²⁴⁹ A combination of realism and efficiency, at the end of the day, will privilege publicly accessible meaning.²⁵⁰

Whether the text wrapping layer should displace case commentary and commit logs is a much harder problem. Here, the concerns about publicly accessible meaning drop away, since code commentary is generally written in English (or at least a coding dialect that can be grokked). The remaining issue is whether commentary intending to explain a cipher is as good evidence of what the exchange was intended to accomplish as the wrapping text's more legalistic frame.

On the one hand, the commit logs and commentary are integral to the code itself, expressed contemporaneously with its fixation and with the goal of revealing its intent.²⁵¹ It is the "crown jewel" of the

248. One analogy is the courts' treatment of disputes about meaning based on differences in the parties' respective native tongues. *See, e.g.*, *Frigalment Importing Co. v. B.N.S. Int'l Sales Corp.*, 190 F. Supp. 116 (S.D.N.Y. 1960). In such cases, the courts may adopt the broadest and widest-shared meaning available.

249. Rohr points out that courts that have analyzed vending machine contracts were drawn to meeting of the mind analogies, even when they plainly were inapt. Rohr, *supra* note 33, at 80 ("Vending machine cases are . . . predictive of the types of issue that . . . [will] arise as judges attempt to apply foundational common law contract principles to smart contracts going forward.").

250. For a different defense of publicly accessible meaning, see generally Aaron D. Goldstein, *The Public Meaning Rule: Reconciling Meaning, Intent, and Contract Interpretation*, 53 SANTA CLARA L. REV. 73 (2013), which argues that extrinsic evidence should be limited to public and shared meaning to avoid gamesmanship.

251. *See* Daniela Steidl, Benjamin Hummel & Elmar Juergens, *Quality Analysis of Source Code Comments*, in 2013 IEEE 21ST INT'L CONF. ON PROGRAM COMPREHENSION 83, 83 ("A significant amount of source code . . . consists of comments, which document the implementation and help developers to understand the code Comments are the second most-used documentary artifact for code understanding, behind only the code itself.").

code, laying bare its “inner secrets.”²⁵² Code commentary and commit logs are thus like the definition section of an ordinary contract: the very best evidence of meaning.²⁵³ Consequently, commentary and commit logs should be privileged over the code it explains.

That said, some might worry that code commentary (and, to a lesser extent, commit logs) is intended to be disposable²⁵⁴—it might signal what a coder hoped to achieve but is not likely to have been written with particular care.²⁵⁵ For example, as we discussed above, most open source code today is reused from script to script. Thus, it’s not necessarily (or even usually) the case that the commentary was written with a singular project’s goals in mind. Blindly adopting such commentary as gospel risks being misled as to what the parties wanted, just as (for example) adopting boilerplate can, over time, lead parties to use terms that even they do not understand.²⁵⁶

Even considering these risks, commentary and commit messages should have the same interpretative weight as natural language contract terms. True, they might not provide clear evidence of promissory intent. But the same objections can be made about boilerplate that travels from deal to deal. Moreover, though it’s true that some coders treat commentary and logs as disposable, well-counseled projects,

252. Jeffrey D. Sullivan & Thomas M. Morrow, *Practicing Reverse Engineering in an Era of Growing Constraints Under the Digital Millennium Copyright Act and Other Provisions*, 14 ALB. L.J. SCI. & TECH. 1, 17 (2003) (“Reverse engineering does not lay bare a program’s inner secrets. Indeed, it *cannot*. The inner secrets of a program, the real crown jewels, are embodied in the higher levels of abstraction material such as the source code commentary and the specification.”).

253. There may be examples where coders make an explicit attempt to synthesize the semantic contract within the code. See, e.g., LEXON, <http://demo.lexon.tech/apps/editor> [<https://perma.cc/DY2R-4X4D>] (embedding human readable semantic contract within compliable code that exports to Solidity).

254. Andrew Johnson-Laird, *Software Reverse Engineering in the Real World*, 19 DAYTON L. REV. 843, 857 (1994) (“[Source code commentary] is the equivalent of marginal annotations and is intended to assist the original programmer or those that follow in understanding why the program was crafted in a particular way, or to explain a particularly complex flow of logic. There are no restrictions on what must or must not be written in comments, but inevitably they are the repository of all the knowledge that the programmer has in his or her head as the code is being created. One also frequently sees a certain irreverence in the commentary which is a by-product of the exuberance of programmers and is best not taken too seriously . . .”).

255. Cf. Haque et al., *supra* note 52, at 58 (“Another [developer] may constantly contribute a high volume of lines over a long period of time, but still be a functionary whose work is wholly non-essential and could easily be replaced by others.”).

256. See generally Stephen J. Choi, Mitu Gulati & Robert E. Scott, *The Black Hole Problem in Commercial Boilerplate*, 67 DUKE L.J. 1, 2 (2017) (describing *pari passu* clauses as “a standard provision in sovereign debt contracts that almost no one seems to understand”).

knowing the rule that we propose, would be well-positioned to expand on commentary before deploying a script, and impose discipline over commits that are merged into the project. The result could be another opportunity to surface and correct bugs, while aligning the parties' expectations with what they receive. And, of course, this is just a default rule: the parties may express a different rule by contracting for it.²⁵⁷

Finally, we think that the source code generally is a better source of meaning than the compiled byte code. That's so because the source code is, in broad strokes, readable by humans with the exercise of reasonable effort, meaning that all parties to a transaction can gain some insight as to what they've agreed. The alternative, which holds that byte code is the "real" contract, seems likely to lead to embarrassing results. For example, consider this "online user agreement" that a law student confronted:

257. Cf. Surden, *supra* note 32, at 652 ("A 'data-meaning threshold agreement' provides specific interpretations [for computable contracts].").

Figure 5²⁵⁸

This is not bytecode, but rather an erroneous encoding of the bytes into the glyphs that represent them. Obviously, clicking on the agree box wouldn't bind anyone—the gibberish communicates no meaningful information to humans (even if some computer, somewhere, could make sense of it). Byte code is likewise at the very bottom of our interpretative hierarchy, as decoding it requires the exertion of effort and expertise likely beyond the capacity of the contracting parties.

These interpretive principles ought to give way depending on context. Thus, we imagine that ephemeral contractual promises—a

258. Samuel P. Morse (@SamuelPMorse), TWITTER (Sept. 13, 2019, 8:10 AM), <https://twitter.com/SamuelPMorse/status/1172497664799363075> [<https://perma.cc/3KZ2-TY4F>].

stray tweet by a project manager promising a particular outcome of the script—would likely not displace a well-curated GitHub repository. Or, a piece of code commentary that makes a joke. Here, again, our argument rests on a pragmatic judgment as to what courts will do, which in turn relates to the parties' reasonable, commercially-informed expectations.²⁵⁹ The more permanent, considered and reliable the evidence of a promise—the more, in other words, it would be reasonable to rely upon it—the more a court is likely to consider it a reliable basis of the bargain.

Finally, consider a problem of interpretation that has no easy non-scripted analogue. What if the programmers' intent is internally contradictory *and is recorded as such*? Of course, courts will often say that a party's private meaning ought to be discounted when ascertaining the shared intent of the transaction. Thus, simply because one lawyer on a team thought that "black" means "white" doesn't mean that the term takes on that meaning, unless that lawyer communicated her meaning to the other side in a way that seemed authoritative.²⁶⁰ This is why some have argued that visible metadata ought to bear on meaning.²⁶¹

Evidence of inconsistent drafter intent in the transactional script context is different. Given the use of version control systems, both the expression of code and the identity of who wrote each line of human-readable code may be knowable at the moment the counterparty inspects the terms, and certainly at the moment of formation. So is the commentary. Because the code is public, all parties can see the conflicting evidence of meaning at the moment the contract is entered: it is like the parties' final executed document included visible redlined changes.²⁶² Thus, a party might encounter a piece of code that has multiple drafters who appear to be communicating different goals. What to do in this scenario?

259. MURRAY, *supra* note 185, § 4.12 (discussing the role of reasonable interpretations in determining meaning).

260. See *Cendant Corp. v. Commonwealth Gen. Corp.*, No. 98C-10-034 HLA, 2002 WL 31112430, at *6 (Del. Super. Ct. Aug. 28, 2002) (holding that a contract's "Material Adverse Change" created an issue of material fact because each party had plausible, but different, interpretations of the clause).

261. See, e.g., Thomas H. White, *Parol Metadata: New Boilerplate Merger Clauses and the Admissibility of Metadata Under the Parol Evidence Rule*, 4 CASE W. RES. J.L. TECH. & INTERNET 237, 267 (2012) ("If the metadata is visible on the final version of the contract, this should be conclusive weight in favor of its admissibility [under the parol evidence rule].").

262. For an argument that non-resolved tracked changes ought to be included as part of the integrated document, see Elizabeth A. Janicki, *Contracts as Speech Acts: Bringing Jakobson to the Conversation*, 107 GEO. L.J. 201, 218 n.113 (2018).

The simplest answer—harmonization with the rest of the agreement and with the social context—is probably best. But its application is subtle and rooted in the sociology of the operative coding community. We ought to prefer (and render operative) the meanings of coders who advance the larger agenda of the project, to the extent it can be reconstructed.²⁶³ Failing that, we should prefer later in time to earlier in time pieces of code, and code which hews closely to the commentary that surrounds it. These match the background rules of contract interpretation, which generally seek to find and give priority to the best available evidence of the parties' expressed and integrated intent.

Overall, this interpretative hierarchy, which promotes supervised coding and last-in-time syntax, arises out of the same intuitions that generate the parol evidence rule. Recall Corbin's famous, though not universally accepted, rationale for the rule.²⁶⁴ It was not, he argued, primarily about controlling self-serving frauds by excluding convenient *ex-post* evidence of meaning.²⁶⁵ Rather, the rule was intended to give priority to the parties' fixed agreement: to discard those earlier agreements, negotiations and understandings which had not made it into the final and binding contract.²⁶⁶ On the question of whether the parties intended a particular expression to be the final expression, "no relevant evidence, [] parol or otherwise, is excluded."²⁶⁷ But even if admitted, courts could clearly weigh such evidence and find that "the more bizarre and unusual an asserted interpretation is, the more convincing must be the testimony that supports it."²⁶⁸

263. See Hunn, *supra* note 31, at 281 ("[I]ncreasing attention is given to architecting smart legal contracts 'in a language that is both human-intelligible and machine readable, whose text incorporates an algorithm which automates some or all of the performance of the agreement.'" (quoting J.G. Allen, *Wrapped and Stacked: 'SmartContracts' and the Interaction of Natural and Formal Language*, 14 EUR. REV. CONT. L. 307, 313 (2018))).

264. See generally Arthur L. Corbin, *The Interpretation of Words and the Parol Evidence Rule*, 50 CORNELL L.Q. 161, 189 (1965) ("[A]ntecedent agreements . . . are rendered inoperative [evidence] by having been discharged by a subsequent agreement that has been duly proved and interpreted.").

265. See PETER LINZER, CORBIN ON CONTRACTS: PAROL EVIDENCE AND IMPLIED TERMS § 25.2, at 6 (rev. ed. 2010) ("[Corbin's] statement of the [parol evidence] . . . had nothing to do with the reliability of oral or written evidence of intent; it had to do only with what the parties actually wanted to be the final word of their agreement.").

266. *Id.*

267. *Id.* at 9. But see Klass, *supra* note 213 (describing some theorists' view that integration also controls interpretation evidence).

268. LINZER, *supra* note 265, § 25.4, at 32.

On this understanding, if a transactional script's promoters seem to advance multiple potential purposes at odds with one another, a court ought to identify those terms that best match the *reasonable* expectations of the parties at the moment that the counterparty committed to the particular exchange. The existence of prior conflicting terms would still be potentially relevant, but they ought to be de-emphasized.

In applying this canon, consider the Synthetix example discussed above. One party deployed a bot that took advantage of a corrupted third-party oracle to execute a trade which would have garnered it potentially a billion dollars in profit. Presumably because that amount would have been uncollectable, the hacker settled for some undisclosed bug bounty, paid off-chain. Neither party tested what contract law would have had to say. But we can speculate somewhat as to its contents.

Canon 1 reminds us to take up all the sources of meaning. A license, deployed by Synthetix, disclaimed all warranties as to the Code's correctness. But its commentary promised that the quotes were the "current market value." Thus, if a court were to ask if the parties intended this result, the answer would turn, we think, on the hierarchy in *Canon 3*. In our view, as between two kinds of wrapping natural language, the commentary to the code ought to take precedence over the ambiguous license, meaning that Synthetix would have a difficult time arguing that whatever the oracle delivered was, for the parties' purposes, actionable "market values." At the very least, Synthetix should have borne a heavy burden of proving an alternative.²⁶⁹ Thus, Synthetix, which presumably would have argued mistake, probably fairly bears the burden of a third-party data source risk.

Canon 4: Where Natural Language Contracts Refer to Code, Integration Clauses Should Be Read Narrowly.

We have suggested a set of default rules for interpreting scripts. But what if the parties want to vary such rules by agreement: should we give these scripted integration clauses force?²⁷⁰ Offline, courts

269. The case would have been different with an appropriate disclaimer in the terms of use, which we have not yet found.

270. See Klass, *supra* note 213, at 466–71 (discussing when integration clauses are legally enforceable). See generally Gregory Klass, *Intent to Contract*, 95 VA. L. REV. 1437, 1442–43 (2009) (discussing the legal relevance of parties' intent with respect to legal enforcement).

have generally permitted integration clauses to control which pieces of prior or contemporaneous contracting are included within the litigated deal, at least between sophisticated firms.²⁷¹ But courts have sometimes been dubious about attempts to integrate fuzzy stacks and to limit evidence of meaning that appears otherwise relevant.²⁷²

A problem here is that sophisticated projects will usually directly refer to the script in the natural language terms and conditions. In such cases, we think it is impossible to exclude the script entirely—in other words, we don’t think that script can be both pointed at and also treated as extrinsic evidence. That would be much like saying that an addendum to a contract, which contains a key description of the relevant subject matter, is not a part of the deal.

What parties might do is to try to use the natural language contract to determine meaning. Natural language terms and conditions would state that the only operative promises are those found in the natural language itself, and that parties should not read the commentary in the script to make additional or contradictory promises about what it accomplishes. This would not deny that the code is a part of the bargain, but rather would attempt to limit what can be inferred from the natural language text it contains. Or, the converse: that is, the natural language may deny its own efficacy and privilege code.

The choice of whether to defer to such attempts to control meaning turns on whether the court generally adopts a more contextual or more formalist approach to interpretation.²⁷³ Contextualism, which discourages opportunistic drafting and thus protects consumers, has

271. See *Klass*, *supra* note 213, at 475–78 (arguing for a “hard express integration rule for firm-to-firm negotiated contracts”).

272. Cf. *LINZER*, *supra* note 265, § 25.7, at 61 (“[T]he essence of integration is whether they intended a document to be the final word, and evidence of this intention should be found from all sources, not just the words of the contract.”).

273. Compare *Wells Fargo Bank, N.A. v. Cherryland Mall Ltd. P’ship*, 812 N.W.2d 799, 810 (Mich. Ct. App. 2011) (admitting usage of trade notwithstanding contractual clause stating that “no trade practices . . . shall be used to contradict, vary, supplement or modify any term of this guaranty agreement”), with *S. Concrete Servs., Inc. v. Mableton Contractors, Inc.*, 407 F. Supp. 581, 584 (N.D. Ga. 1975) (“The court recognizes that all ambiguity as to the applicability of trade usage could be eliminated by a blanket condition that the express terms of the contract are in no way to be modified by custom, usage, or prior dealings.”). See generally Joshua M. Silverstein, *Contract Interpretation Enforcement Costs: An Empirical Study of Textualism Versus Contextualism Conducted via the West Key Number System*, 47 *HOFSTRA L. REV.* 1011 (2019) (providing an empirical study of judicial and academic debate over textualism and contextualism in contract interpretation); Lisa Bernstein, *Custom in the Courts*, 110 *NW. U.L. REV.* 63, 71 (2015) (“The enforceability and effectiveness of a general clause opting out of all trade usages is at best unclear.”).

much to commend it in markets where sharp dealing is more prevalent.²⁷⁴ Fraud has defined many blockchain products to date, as has incoherent transactional lawyering. This is not a space producing formalism's best factual predicates.

To the extent these issues seem fanciful, consider the DAO hack.²⁷⁵ The DAO was a token-mediated platform that allowed small investors to enter jointly into a venture capital pool.²⁷⁶ The entity's "terms," apart from disclaiming various legal rights, stated that "[t]he "use of The DAO's smart contract code . . . carries significant financial risk, including using experimental software." ²⁷⁷ However, it also stated:

“

Explanation of Terms and Disclaimer

The terms of The DAO Creation are set forth in the smart contract code existing on the Ethereum blockchain at 0xbb9bc244d798123fde783fcc1c72d3bb8c189413. Nothing in this explanation of terms or in any other document or communication may modify or add any additional obligations or guarantees beyond those set forth in The DAO's code. Any and all explanatory terms or descriptions are merely offered for educational purposes and do not supercede or modify the express terms of The DAO's code set forth on the blockchain; to the extent you believe there to be any conflict or discrepancy between the descriptions offered here and the functionality of The DAO's code at 0xbb9bc244d798123fde783fcc1c72d3bb8c189413, The DAO's code controls and sets forth all terms of The DAO Creation.

Figure 6: The DAO's Terms

274. See Silverstein, *supra* note 273, at 1018 (noting that contextualist courts consider both an agreement's language and extrinsic evidence to determine ambiguity).

275. See Haque et al., *supra* at 52, at 39–45 (providing an overview of the DAO hack); see also Laila Metjahic, *Deconstructing the DAO: The Need for Legal Recognition and the Application of Securities Laws to Decentralized Organizations*, 39 CARDOZO L. REV. 1533, 1536 (2018) (analyzing the corporate legal theories under which the creators of The DAO might be held liable for the 2016 hack).

276. Cf. Vitalik Buterin, *Bootstrapping a Decentralized Autonomous Corporation: Part I*, BITCOIN MAG. (Sept. 20, 2013), <https://bitcoinmagazine.com/articles/bootstrapping-a-decentralized-autonomous-corporation-part-i-1379644274> [<https://perma.cc/6FNC-93QV>].

277. *Terms: Explanation of Terms and Disclaimer*, DAOHUB, <https://daohub.org/explainer.html> [<https://web.archive.org/web/20160704190119/https://daohub.org/explainer.html>].

The code repository contained an even more incoherent set of disclaimers, including a readme file that claimed using the software “does not, in and of itself, create a legally binding contract,” and that “in order for you to form a legally binding contract . . . you shall seek legal advice from an appropriately qualified and experienced lawyer”²⁷⁸ This set of disclaimers appears to be an attempt to shield the entity from liability by at once embracing and rejecting contract law.²⁷⁹

After users had contributed funds, but before the DAO’s own investments had begun, someone noticed a flaw in its code which allowed siphoning of \$55 million (of around \$170 million total assets) out of the pool. Ethereum’s then developers promulgated a proposed software update to the entire blockchain—a hard fork—which was adopted by some, but not all, holders of the original tokens.²⁸⁰

The DAO’s creators had some warning of the vulnerability.²⁸¹ Before the hack, a commentator posted about the vulnerability (titled, “Protect against recursive withdrawRewardFor attack”) and suggested a seemingly easy change (reversing the ordering of two lines of code) which would close it.²⁸² The DAO made the change, reassuring

278. Stephan Tual, Posting to */blockchainsllc/DAO: Updated Readme*, GITHUB (Apr. 11, 2016), <https://github.com/slockit/DAO/commit/aceec3efcc8afd4277396ebc42628f2e5ca8dff2#diff-04c6e90faac2675aa89e2176d2eec7d8> [<https://perma.cc/5HHY-5W4V>].

279. For a trenchant analysis of these issues, see Drew Hinkes, *A Legal Analysis of the DAO Exploit and Possible Investor Rights*, BITCOIN MAG. (June 21, 2016), <https://bitcoinmagazine.com/articles/a-legal-analysis-of-the-dao-exploit-and-possible-investor-rights-1466524659> [<https://perma.cc/XJK6-TCKB>].

280. Reyes, *supra* note 30, at 388; see *The Ethereum Classic Declaration of Independence*, ETHEREUM CLASSIC, https://ethereumclassic.org/ETC_Declaration_of_Independence.pdf [<https://perma.cc/F6D5-7FSR>] (last updated July 2019) (declaring the holders’ intent to “continue the original Ethereum blockchain” and decrying the hard fork as a violation of the blockchain’s “core tenets”).

281. See, e.g., Matthew Leising, *The Ether Thief*, BLOOMBERG (June 13, 2017), <https://www.bloomberg.com/features/2017-the-ether-thief> [<https://perma.cc/N8RJ-F7RE>] (“Gün . . . had already been tracking and publicizing flaws in the DAO’s design. . . . [He] appears to be the first to pinpoint the flaw that put the money in jeopardy.”).

282. LefterisJP, Posting to *bloackhainsllc/DAO*, GITHUB (June 12, 2016), <https://github.com/slockit/DAO/commit/f01f3bd8df5e1e222dde625118b7e0f2bfe5b680?diff=split> [<https://perma.cc/87ZY-FD5V>].

users that, “The important takeaway from this is . . . this is NOT an issue that is putting any DAO funds at risk today.”²⁸³ The updated version was called The DAO 1.1 “milestone.” In the code, the in-line comment preceding the new revision on line 580 stated its explicit purpose:

```
// we are setting this here before the CALL() value transfer to
// assure that in the case of a malicious recipient contract trying
// to call executeProposal() recursively money can't be transferred
// multiple times out of the DAO284
```

Assurance notwithstanding, someone then executed the famous hack—in part because The DAO’s fix was incomplete—transferring money multiple times out of The DAO.

Thus, here we again have a stack of meaning about what the parties to the contract—The DAO’s creators and its investors—expected.²⁸⁵ But the relevant documents are contradictory. The actual code did not accomplish what the comment or white paper promised, a fact that soon became obvious to all. Moreover, the organizers of The DAO had specifically told users that the code governed, even as they disclaimed the legal enforceability in the code itself.²⁸⁶ Let’s use the canons to offer a solution to the private law contracting problems that the DAO occasioned.

As *Canon 1* instructs, we ought to consider all the relevant pieces as evidence of meaning, and of what was promised. The code, terms and conditions and readme files are all a part of the stack.

Canon 2 suggests that the DAO’s counterparties reasonably could have believed that they would not face the risk of recursive money

283. Stephan Tual, *No DAO Funds at Risk Following the Ethereum Smart Contract ‘Recursive Call’ Bug Discovery*, SLOCK.IT BLOG (June 12, 2016), <https://blog.slock.it/no-dao-funds-at-risk-following-the-ethereum-smart-contract-recursive-call-bug-discovery-29f482d348b> [<https://perma.cc/2P4D-5QGB>].

284. LefterisJP, *supra* note 282, at line 580, <https://github.com/slockit/DAO/blob/d48ee5c49f9dc3b9548623aa6985cbc3c9528b67/DAO.sol#L580> [<https://perma.cc/9KTB-7KWJ>].

285. See Tanaya Macheel, *The DAO Might Be Groundbreaking, but Is It Legal?*, AM. BANKER (May 19, 2016, 3:12 PM), <https://www.americanbanker.com/news/the-dao-might-be-groundbreaking-but-is-it-legal> [<https://perma.cc/98RV-C7SB>] (narrating Stephan Tual’s statement, a primary organizer, that “[n]o one benefits from it except the people that support it. Even we, the ones who invented it, get nothing”).

286. See generally Kolber, *supra* note 25, at 221 (“Note: Although the word ‘contract’ is used in The DAO’s code, the term is a programming convention and is not being used as a legal term of art.”).

transfer.²⁸⁷ That promise, embodied in a marketing announcement and in the code comments itself, is, it is true, absent in the operational code. It's as if a door which contained on its front face a sign stating "Private Property: Locked Door" was freely openable with a key hanging nearby. If someone were to have been harmed by relying on that set of statements, when suing in fraud or tort they would face questions about how reasonable their precautions had been.²⁸⁸ But in a contract lawsuit, *Canon 3* teaches that code should bow to comments and commit logs in a way that best embodies what the programmers intended to accomplish and therefore to what they ought to be held.²⁸⁹ *Canon 4* tells us to discount the attempts at non-integration as, at best, confused.

We thus disagree with The DAO's attacker, who argued that the fork violated its rights because those actions were literally permitted by the code.²⁹⁰ Yes, code drafters ought to bear the interpretative risk of error.²⁹¹ But the non-drafting counterparties whose funds were taken could not reasonably be expected to know that the code had that bug, given the commentary promising the opposite result. Had they not received their money back, The DAO's investors should have been

287. See Rohr, *supra* note 33, at 89 (noting analogies to vending machine cases to find that "the agreement includes only those terms that were reasonably available to DAO Token holders prior to purchase" and impliedly concluding that this would include the Terms of Use but not the permissive code).

288. One problem we leave to future work is the relationship between putting statements inside the contractual stack and the tort-contract line in litigation.

289. Thus, we reject the idea that the only operative promises are those written in code. Cf. Lawrence Lessig, *Code Is Law: On Liberty in Cyberspace*, HARV. MAG., Jan.-Feb. 2000, <https://harvardmagazine.com/2000/01/code-is-law-html> [<https://perma.cc/59LX-BRU5>] (noting that code implements values and that coders decide how cyberspace regulates).

290. A Guest, *An Open Letter*, PASTEBIN (June 18, 2016), <https://pastebin.com/CcGUBgDG> [<https://perma.cc/Y739-MR8R>] (claiming that the fork "would amount to seizure of my legitimate and rightful ether, claimed legally through the terms of a smart contract").

291. In this context, perhaps all participants, including investors and programmers, could be considered simply partners. See Stephen Palley, *How to Sue a Decentralized Autonomous Organization*, COINDESK (Mar. 20, 2016, 3:17 PM), <http://www.coindesk.com/how-to-sue-a-decentralized-autonomous-organization> [<https://perma.cc/B8L6-M7RM>] ("A DAO is an organization that's self-governing and that isn't influenced by outside forces. . . . DAOs are formed by groups of like-minded individuals. . . ."). In that event, agency law would presumably add complexity to the interpretative defaults.

able to bring an action for breach against its developers, or even, perhaps, against the attacker if its action amounted to participating in the nexus of contracts that The DAO had proposed.²⁹²

C. RECAPITULATING THE CANONS: *QUOINE* AND NON-PUBLIC SCRIPTS

We have confined our analysis to the definition of transactional scripts we offered in the introduction. Such scripts present a relatively tractable set of problems for contract jurists. Because the scripts are public, and participants typically are knowingly participating in a specialized form of commerce, the sorts of concerns we might have about knowledge, opportunism, and black-box contracts are, by and large, muted.²⁹³

How to apply the canons when parties are *not* sophisticated and cannot easily parse the code at the moment of the exchange, presents a distinct and difficult set of questions. In the *Quoine* case, for example, it is not obvious that either side had easy access to each other's code. Where parties cannot access code, it can't communicate meaning, no matter how well drafted its commentary. Thus, it can't be part of the stack that comprised the grist for the bargain.

But what happens when non-programming parties can access the code but there is no social expectation that it expresses the intent to contract: that is, what about "smart contracts" with transparent code, but involving ordinary consumers who might not even be aware that their contract's execution occurs automatically? When natural language promises conflicts with coded rules, should we read them together? One possibility is that *Canon 3*, which treats natural language describing and commenting on code as on the same level as contract terms themselves, might need recalibration. Or it might not. Because ordinary contract terms themselves are typically unread, more work is necessary to consider which sorts of buried terms count in making bargains. We leave these problems to future research.

292. See SEC. & EXCH. COMM'N, EXCHANGE ACT RELEASE NO. 81207, REPORT OF INVESTIGATION PURSUANT TO SECTION 21(A) OF THE SECURITIES EXCHANGE ACT OF 1934: THE DAO 1, 7–8 (2017), <https://www.sec.gov/litigation/investreport/34-81207.pdf> [<https://perma.cc/Q4HV-SZXE>] (suggesting that DAO curators faced potential liability); see also Hinkes, *supra* note 279 (arguing that investors could potentially bring claims against other investors and DAO curators).

293. See Gregory Klass, How to Interpret a Vending Machine: Smarts Contracts and Contract Law (unpublished manuscript) (on file with author), for an in-depth discussion of these issues.

IV. THE FUTURE OF THE CONTRACT STACK

To date, transactional scripts haven't delivered revolutionary change to either the world or our small, legal, corner of it. In the big picture, the ecosystem is marginal: billions of dollars of investment in a trillion-dollar world economy. And yet the intellectual footings of the script project are expanding at an astounding rate. Every day, new projects (like Facebook's Libra, or JP Morgan's fiat coin) launch with scripted roots, and the technical community gains experience and competence with each failure. We simply have no idea what the future of coded exchange will look like.²⁹⁴

The relationship between this burgeoning, but still highly speculative, ecosystem and law are typically described as antagonistic. Thus, for noted commentator Nick Szabo, the primary virtue of "smart contracts" is that they ostensibly don't need law.²⁹⁵ For others, scripted deals "will subject the provision of justice to market forces and break the state's monopoly over the court system."²⁹⁶ This is an ideological call to arms against the civilizing and constraining role that contract jurists have traditionally had in commercial life.

Our approach is different. First, unlike some skeptics, who think that transactional scripts are toys with no real use cases, we believe that they are a potentially valuable new contracting technology. It might be that scripts will reduce back-end transactional costs by reducing the need for transactional lawyering.²⁹⁷ In certain settings, the

294. See generally Allen, *supra* note 33, at 322 (warning against too heavily discounting the likelihood of successful deployment of legal contracts in code).

295. See, e.g., Nick Szabo (@NickSzabo4), TWITTER (Oct. 14, 2018, 5:51 PM), <https://twitter.com/NickSzabo4/status/1051606530108190720> [<https://perma.cc/N5JY-YVED>] ("Worrying about whether a smart contract is 'legally enforceable' reflects a profound misunderstanding. The main relation of smart Ks to traditional courts is that smart Ks control burden of lawsuit."); see also CleanApp, *Why's Szabo Afraid of "Smart Contract" Critiques?*, MEDIUM (Oct. 16, 2018), <https://medium.com/cryptolaw-review/whys-szabo-afraid-of-smart-contract-critiques-669ef9e63fc0> [<https://perma.cc/75CG-5NA3>] (statement of Nick Szabo) ("The parties can if they choose write a traditional K to backstop a smart K, although in many situations where a smart contract is useful the exercise would be pointless . . .").

296. Raskin, *supra* note 27, at 335.

297. See Sklaroff, *supra* note 19, at 275–77 (noting that transactional scripts can reduce accounting, due diligence, monitoring, and enforcement costs); see also Werbach & Cornell, *supra* note 26, at 322, 348 (noting the ability of transactional scripts to automatically verify, facilitate, and remedy).

benefits from cutting enforcement costs will be worth the costs of up-front specification.²⁹⁸ By reducing monitoring costs on the margin, transactional scripts may make it more likely for parties to enter into the exchange. Similarly, in regimes where institutional trust is at a nadir and centralized trading repositories are unreliable, scripts can provide significant value.²⁹⁹

And yet, this value will have natural limits, defined by a tradeoff between the value of trustless computing and the added costs imposed by the complexity tax, whose scope we are the first to make concrete. Most solutions to the complexity tax require either the intermediation of third parties or consensus protocols running across a smaller set of validators.³⁰⁰ Thus, at least for now, complex organizational solutions will remain within “real” contracts, while particular, discrete, computable aspects can be put on public blockchains. This conclusion fits well with the most sophisticated guidance currently available.³⁰¹

The future of scripts is thus about hybrids, where code and natural language must work together. At the bottom, legal scholarship about computable contracts simply hasn’t fully grappled with the irreducibly buggy nature of coding. Errors in coded exchange will result in the parties’ outcomes stubbornly failing to match their goals. In our view, the persistence of error, and the hazards of determining intent, makes recourse to third-party decisionmakers inevitable. Transactional scripts are not, and will not be, self-executing, at least not all the time. At that point, such decisionmakers will be well-advised to look at traditional contract principles to help resolve disputes.

298. Cf. Sklaroff, *supra* note 19, at 283–84 (“Agreements that effectively specify relevant commercial context are . . . easier to resolve on summary judgement, reducing incentives to litigate.”).

299. See generally Eric Tjong Tjin Tai, *Force Majeure and Excuses in Smart Contracts*, 26 EUR. REV. PRIV. L. 787, 787 (2018) (noting that the main advantage of smart contracts is guaranteed performance due to the absence of human and legal intervention).

300. Other partial solutions to the complexity tax involve more exotic forms of cryptographic proofs (such as Proof-of-Retrievability for storage) that can reduce redundancy and thus costs. See, e.g., Andrew Miller, Ari Juels, Elaine Shi, Bryan Parno & Jonathan Katz, *Permacoin: Repurposing Bitcoin Work for Data Preservation*, 2014 IEEE SYMP. ON SEC. & PRIV. (proposing a modification to Bitcoin that repurposes its mining resources to distribute storage of archival data). The ultimate efficacy and commercial adoption of such schemes remains an open question.

301. Cf. INT’L SWAPS & DERIVATIVES ASS’N, *supra* note 204 (providing “high-level guidance on the legal documentation and framework that currently governs derivatives trading”).

Our approach would bring scripts within the traditional world of contract law through a constitutive legal act: the compilation of a contract stack. This solution is not merely useful for the current form of transactional script. It has relevance for other sorts of digital and algorithmic contracting, whether now contemplated or yet to be imagined. Code can communicate executory intent to contract and can thus be the grist for legal analysis. But, because it is imperfect, code-mediated transactions will often fail to achieve what their promisors intend, even as they are surrounded by communications in “real” languages, intended to be relied on by real people. In such cases, law will confront—and must surmount—two temptations: ignoring the code altogether as a mere instrument of performance or enforcing it as an exculpatory clause written in ciphered text.

We argue for an alternative approach, which first, claims transactional code, commentary, and logs as a part of the contractual stack, capable of expressing meaning about the parties’ intent. We’ve also defined a hierarchy of meaning that situates natural language disclosures above artificial ones. The canons we’ve proposed, built on an informed understanding of the coding ecosystem, predictably enforce the parties’ reasonable, publicly communicated, intent, and will forbid opportunistic exploitation of subtle errors in coding. These ought to be the law’s goals in compiling contracts executed on blockchains, as well as whatever forms of coded exchanges the future delivers to us.