Electronic Coins

Craig Warmke, Northern Illinois University

In the bitcoin whitepaper, Satoshi Nakamoto (2008: 2) defines an electronic coin as a chain of digital signatures. Many have since defined a bitcoin as a chain of digital signatures. This latter definition continues to appear in reports from central banks, advocacy centers, and governments, as well as in academic papers across the disciplines of law, economics, computer science, cryptography, management, and philosophy. Some have even used it to argue that what we now call bitcoin is not the *real* bitcoin. But the definition fails and Satoshi likely never endorsed it. In this paper, I explain why it fails and what Satoshi likely endorsed instead. Along the way, I untangle some issues around bitcoin fungibility and clarify some others around the ontology of digital assets.

1. INTRODUCTION

Bitcoin automates and decentralizes two tasks normally entrusted to third parties and centralized institutions. First, bitcoin's peer-to-peer network facilitates transfers of value without relying on trusted intermediaries. So anyone can send bitcoin directly to anyone without trusting payment processors and credit card companies. Second, no central bank controls the supply of bitcoin. Instead, nodes on the bitcoin network run software that encodes a fixed, disinflationary issuance schedule that caps the maximum supply under 21 million around the year 2140.

Thus far, we've said something about how bitcoin *works*. But we've said very little about what bitcoins *are*. On this score, many have identified bitcoins with particular chunks of code. They draw inspiration from Satoshi Nakamoto, bitcoin's pseudonymous inventor. In the bitcoin whitepaper, Satoshi says: "We define an electronic coin as a chain of digital signatures."[1] Like a series of physical signatures on a check, an electronic coin's chain of digital signatures represents a transaction history. So, in the whitepaper, Satoshi identifies an electronic coin with its encoded transaction history.

Many have since replaced the reference to electronic coins in Satoshi's definition and defined each *bitcoin* as a chain of digital signatures. We'll call this the *Chain Definition*. In effect, the Chain Definition identifies each bitcoin with its encoded transaction history. And it continues to appear in reports from central banks,[2] advocacy centers,[3] and governments,[4] as well as in academic papers across the disciplines of law,[5] computer science,[6] economics,[7] cryptography,[8] management,[9] philosophy,[10] and the nascent field of cryptoeconomic systems.[11] Some have even used it to argue that what we now call "bitcoin" is not the *real* bitcoin.[12] However, the Chain Definition fails, and, despite appearances, Satoshi likely never endorsed it.

Many questions in cryptoeconomic systems require skills or evidence from more than one discipline.[13] Practitioners in one discipline may have skills or evidence relevant to a question that

[1] Nakamoto [2008, 2].
[2] ECB [2012].
[3] Van Valkenburgh [2014, 9].
[4] Lastra and Allen [2018, 55].
[5] Akins et al. [2014, 30 n. 30], Zhang [2017, 560], Borroni [2016].
[6] Wu et al. [2017, 3124].
[7] Kroll et al. [2013, 3]. This deserves a qualification, which I address in Section 3.3.
[8] Gao et al. [2018, 27207].
[9] Friedlmaier et al. [2018, 2].
[10] Bjerg [2016, 3].
[11] Khalilov and Levi [2018].
[12] Peter Rizun, at the Future of Bitcoin Conference, in 2017. See https://youtu.be/hO176mdSTG0?t=362 Those remarks occurred before the Bitcoin Cash hard fork, but Rizun repeats the criticism after the hard fork: https://twitter.com/PeterRizun/status/935285146562859008.
[13] Voshmgir and Zargham [ms].

practitioners in other disciplines lack. This uneven distribution also holds with respect to the status of the Chain Definition. Many bitcoin developers and computer scientists already know that the definition fails. But relatively few others do. Even fewer know why. And, in my experience, understanding why the Chain Definition fails does not guarantee that one also understands Satoshi's original claim about electronic coins. With respect to the Chain Definition, I attempt in this paper to narrow the chasm between disciplines closer to the epicenter of bitcoin development and those disciplines further away. As we proceed, we will also untangle some issues around bitcoin fungibility and clarify some others around the ontology of electronic coins.

Here's the roadmap we'll follow. In the next section, we assess the Chain Definition's meaning and justification. In Section 3, I explain why the definition fails. Then, in Section 4, we return to Satoshi's definition of electronic coins. There, we explore the ontology of electronic coins and propose a framework for understanding Satoshi's definition. Finally, in Section 5, we conclude with some brief reflections on the interdisciplinary nature of cryptoeconomic systems.

2. THE CHAIN DEFINITION

The Chain Definition identifies each bitcoin with a chain of digital signatures. But without some idea of what a bitcoin is, or what chains of digital signatures are, we'll be ill-equipped to evaluate the definition of one in terms of the other. So let's begin with a quick review of each.

2.1. Bitcoins

Bitcoin is a highly divisible asset, and its divisibility requires a unit of measurement. Somewhat confusingly, the word 'bitcoin' serves not only as a name for the asset (e.g., 'I have bitcoin'), but also as the unit of measurement for that asset (e.g., 'I sent 3.275 bitcoin') and a count noun for whole number amounts of the asset (e.g., 'He received four bitcoins'). A whole bitcoin also divides into 100 million units called *satoshis*. Hence, 2.37 bitcoin equals 237 million satoshis.

Those who endorse the Chain Definition understand Satoshi's reference to "electronic coins" as a reference to bitcoins (in the count noun sense). But they have just as much reason to think the reference applies to satoshis. We can conceive of every bitcoin transaction as a transfer of a whole number amount of satoshis. Furthermore, bitcoin transactions are even denominated in satoshis within the *blockchain*, the ledger of bitcoin transactions. Services which represent transactions denominated in bitcoin have converted the raw transaction amounts denominated in satoshis to the presently more readable amounts denominated in bitcoin.[14] But instead of evaluating the claim that satoshis are chains of digital signatures, we will continue to focus on the Chain Definition, which concerns bitcoins and remains popular. The main arguments below apply equally well to both claims.

2.2. Digital Signatures

Digital signatures help the bitcoin network automate two jobs we often entrust to banks. When Alice writes a physical check to Bob, she specifies an amount to credit Bob's account from her own and signs the check to authorize the transaction. When Bob deposits the check, the bank then performs two tasks. First, the bank verifies Alice's signature. By verifying signatures, banks help ensure that no one but the owner(s) of an account transfers funds from it. Second, after the bank ensures that Alice's account has sufficient funds, the bank clears the transaction by debiting Alice's account and crediting Bob's with the specified amount. So the bank serves as a trusted intermediary by verifying signatures and clearing transactions. The bitcoin network verifies signatures and clears transactions without trusted intermediaries.

Let's first review how the bitcoin network disintermediates the verification process. We will simplify matters here as much as possible and abstract away from unnecessary details. Suppose Alice has a digital address with a bitcoin she has yet to spend. When Alice sends that bitcoin to Bob's address, she uses a software application to write a candidate transaction. Though we will soon cover

[14] Antonopoulos [2017, 121-122].

transactions in more detail, we now highlight four important details included in Alice's candidate transaction:

(i) a spending address,[15]
(ii) an amount of bitcoin to spend,
(iii) the previous transaction(s) whose output(s) credited the address in (i) with enough bitcoin for (ii), and
(iv) a receiving address.

Not just anyone can spend an address's bitcoin, though. For Alice's transaction to appear in the ledger, she needs the spending address's private key, a string of characters that functions like the address's password. This private key enables Alice to produce a digital signature that is otherwise practically impossible to provide.

Why would Alice need the private key to produce the signature? The digital signature results from feeding a special function two chunks of information. Roughly, the first chunk is the information embedded in (ii)-(iv). The second chunk is Alice's own private key.[16] The resulting signature is important for two reasons. First, producing the signature without the private key is practically impossible. Second, anyone can easily verify that Alice's private key helped produce the signature without access to the private key. To verify a signature over a proposed transaction, the verification function only needs the signature and the proposed transaction, all publicly available information. Computers running the bitcoin software called *full nodes* use this function to verify signatures and reject candidate transactions whose signatures fail to verify. Full nodes also reject attempts to spend previously spent bitcoin, as well as other kinds of ill-formed attempts.

In addition to disintermediating the verification of signatures, the bitcoin network also disintermediates the clearing of transactions. Full nodes send valid transactions to *miners*, computers running the bitcoin software that compete to publish those transactions in the ledger's next block. Miners compete by using processing power to solve a well-defined mathematical puzzle approximately every ten minutes. If miners marshal more or less power and solve puzzles more or less quickly, the puzzle difficulty automatically adjusts so that solution time remains close to ten minutes. A miner with a solution for the next block sends the candidate block with the solution back to the full nodes to verify. Once the full nodes verify the block and append it to their own individually stored copy of the ledger, miners begin competing for the next block.

We can only consider a transaction to be more or less cleared depending on how many blocks have been published since. Why? Due to the way blocks are ordered, undoing a transaction—and its block—requires creating an alternative chain that branches off from a previous block. The alternative chain then begins adding blocks from that point forward. But bitcoin's proof-of-work consensus mechanism endorses as the official version of its blockchain whichever chain furnishes proof of having used the most processing power to solve puzzles. Therefore, to undo a transaction, one would have to build an alternative chain that the network eventually judges as having used more processing power than the original, even as the original continues to grow. Hence, the further back a block appears in the blockchain, the safer it is from tampering because one would have to solve puzzles at a faster pace than the miners incentivized to work on the original chain. This is the brilliance of bitcoin: the mining reward lures people to compete in a way that secures the value of that very reward. But because someone could theoretically use enough power to undo a past block, transactions are cleared only in a probabilistic sense.

2.3. The Definition's Appeal

Now that we've covered how the bitcoin network automates the verification and clearance functions, we can finally see the appeal of the Chain Definition. The table below abstracts away from

[15] Technically, Alice provides not the spending address but the public key which returns the address after being fed through a compound hashing function. See https://en.bitcoin.it/wiki/From_address.
[16] See Antonopoulos [2017, 139 ff.] for more details.

crucial details involved in transactions, and we'll cover those sorts of details later. But, for now, the table includes enough detail to trace a bitcoin's trajectory through a chain of signatures across the blockchain. The simplified transactions in the table should be read as having the following form: *Sender Signature* [sender | receiver | amount | source of unspent bitcoin]. Transactions are, of course, more complicated than this, but those complications need not concern us now. And we will introduce more detail in Section 3.2.

Table I. A Chain of Digital Signatures

Block	Trans. ID	Transaction
#36	txid 105	*A* [A's Address | B's Address | 1 BTC | txid 102]
#43	txid 237	*B* [B's Address | C's Address | 1 BTC | txid 105]
#97	txid 358	*C* [C's Address | D's Address | 1 BTC | txid 237]

D's address has the bitcoin by Block #97. Where did that bitcoin come from? Well, let's look at the transaction in which D's address received it. It bears C's digital signature, which allowed C to unlock the bitcoin received in txid 237. Then, txid 237 bears B's digital signature, which allowed B to unlock the bitcoin received in txid 105. And txid 105 bears A's signature, which allowed A to unlock the bitcoin received in txid 102. We've now described a chain of digital signatures that partially represents a bitcoin's transaction history.

The entries in bitcoin's ledger represent the sources and destinations of various amounts of bitcoin. But unlike a ledger whose entries represent beers owed at the bar or gold owned in a vault, the entries on the bitcoin ledger represent nothing "out in the world." So if bitcoins aren't outside the ledger, perhaps they are somehow embedded in the ledger itself. This thought partially motivates the Chain Definition, which identifies each bitcoin with a transaction history scattered across the ledger.

Yet not just any transaction histories will do. In 2012, the European Central Bank published a report on virtual currencies. After quoting Satoshi's definition of electronic coins as chains of digital signatures, the report claims that the very bitcoins which are "divisible to eight decimal places" are such that each "carries the entire history of the transactions it has undergone, and any transfer from one owner to another becomes part of the code."[17] The European Central Bank has correctly inferred that if bitcoins are transaction histories, they are entire transaction histories. How, then, should we understand the notion of a bitcoin's entire transaction history?

All bitcoin first appears in a *coinbase transaction*, the special transaction that rewards a winning miner's address with newly minted bitcoin. And all bitcoin remains unspent at the address to which it was most recently spent. So if each bitcoin has an entire transaction history, that history should consist of a chronologically ordered series of transactions from the originating coinbase transaction through all and only the subsequent transactions in which that bitcoin is signed over to an address. For any given bitcoin, then, if a series of transactions excludes a transaction in which that bitcoin is signed over, the series isn't that bitcoin's *entire* history. And if a series of transactions includes a transaction in which that bitcoin isn't signed over, the series isn't *that* bitcoin's entire history. Therefore, *if* bitcoins are chains of digital signatures, they are chains of digital signatures that represent entire transaction histories.

We may now offer an official formulation:

> CHAIN DEFINITION. A bitcoin is a chain of digital signatures that represents its entire transaction history.

Many have endorsed either the above definition, something that implies the above definition, or something near enough. We've already mentioned the report from the European Central Bank. In

[17] ECB [2012, 23].

computer science, Wu et al. [2017, 3124] say of the very bitcoins which will someday total near 21 million that "Nakamoto defines them as chains of digital signatures." Similarly, and closer to my own field of philosophy, Bjerg [2016, 3-5] claims of the very bitcoins which will someday number under 21 million that each "consists of a unique chain of digital signatures." In law, after specifying that 'bitcoin' refers to the unit of account, Zhang [2017, 556, 560] says that "each electronic bitcoin consists of a 'chain' of 'digital signatures'..." In economics, Kroll et al. [2013, 1-6] offer a slightly weaker version of the definition. They say that the very "Bitcoins" which were valued at $130 at the time of the paper's writing, and the very things whose number halves every four years in the mining reward, are such that each is "represented as a chain of digital signatures over the transactions in which the coin was used." We could provide more examples across more disciplines. But these will suffice for now.

In my view, the mistaken Chain Definition has achieved enough interdisciplinary influence to merit a public refutation. However, though it's a mistake to infer that bitcoins are chains of digital signatures from Satoshi's definition of electronic coins as digital signatures, the mistake is quite understandable. In fact, I conclude in Section 5 that given certain features of Satoshi's writings and the nature of interdisciplinary research, it was almost inevitable that many would make this mistake. But before we get there, we must first cover why bitcoins are not chains of digital signatures.

3. WHY THE CHAIN DEFINITION FAILS

To help us see why the Chain Definition fails, we'll follow three stages of bitcoin transactions. To describe these transactions, we'll adopt the simplifying assumption that some transactions send bitcoin from one address to another.[18] Although bitcoin transactions don't exactly work this way at the lowest level of detail, they do work this way at a certain level of representation.[19] And, importantly, the argument below doesn't hang on whether we make the simplifying assumption or not. So we might as well save some time and trouble.

Stage One. In this first stage, addresses 1 and 2 (A1 and A2) each send one bitcoin to the previously empty A3:

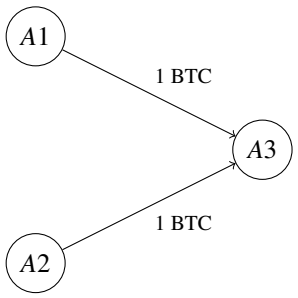

Figure 1. Two bitcoins from different addresses to the same address.

The two bitcoins at A3 are not perfectly fungible—i.e., they are distinguishable from each other. But what distinguishes them? They don't have distinguishable names because bitcoins don't have names. Instead, every quantity of unspent bitcoin is tied in a unique way to its most recent transaction. And

[18] On why this is a simplifying assumption, see n. 15.

[19] See, for example, Christian Decker's answer at https://bitcoin.stackexchange.com/questions/7838/why-does-gettransaction-report-me-only-the-receiving-address.

each unspent bitcoin at A3 is tied to a different transaction. To appreciate this point, we dive a little deeper into the technical details of transactions.

Bitcoin transactions have both inputs and outputs. An output in a transaction specifies an amount of bitcoin and an address as the recipient of that amount. Transactions have one or more outputs, and each has an index number within its transaction. The first or only output in any transaction is 0, the second is 1, and so on. Consequently, we can identify any output in the blockchain by referring to its index number and associated transaction ID. Transactions also have one or more inputs, and each input tags an output from a previous transaction to spend by its index number and the transaction ID in which it appears. When we put it all together, each transaction specifies, in its inputs, the bitcoin to spend by referencing outputs from previous transactions, and, in its outputs, how much of that bitcoin goes where.[20]

The bitcoins at A3 are distinguishable because each arrived at A3 through a different transaction. Because the bitcoins at A3 are distinguishable, the possessor of A3's private key may spend the bitcoin from A1 rather than the bitcoin from A2 (and vice versa). To spend the bitcoin from A1, the user would specify the transaction and output in which A1 sent A3 a bitcoin. Alternatively, to spend the bitcoin from A2, the user would specify the transaction and output in which A2 sent A3 a bitcoin.

Stage Two. Instead of spending one bitcoin rather than another, A3 then sends both to A4 in a single transaction:

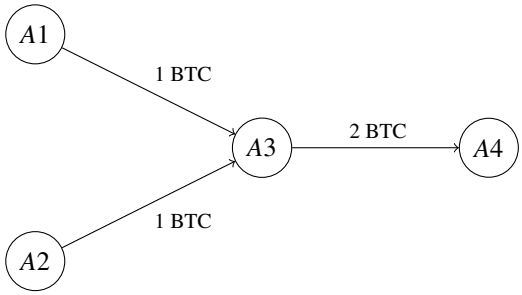

Figure 2. A single transaction with two bitcoins from one address to another.

Previously, in Stage 1, the bitcoins at A3 were distinguishable because they arrived at A3 in different transactions. In Stage 2, however, both bitcoins arrive at an address in a single transaction. So the kind of feature which previously distinguished the bitcoins at A3 no longer distinguishes them at A4. Nothing else distinguishes them either. No names, no tags—nothing. The bitcoins are now perfectly fungible with one another, and it simply isn't possible to tag one rather than the other to spend in a new transaction. Of course, we can spend an amount of one bitcoin out of two.[21] But we cannot single out an individual bitcoin for spending because there are no individual bitcoins to be singled

[20] In this way, bitcoin uses a "transaction-based" rather than an "account-based" ledger, the kind banks ordinarily use. Hal Finney [2008] makes the distinction and categorizes bitcoin correctly two weeks after the whitepaper's publication. For an accessible explanation of the distinction, see Akcora et al. [2018, 2-3]. Although the choice between these two kinds of ledgers has important tradeoffs, they differ less than some let on. For a translation scheme between the two kinds of ledgers, see Zahnentferner [2018]. Because of translation schemes like this, it is wholly appropriate to liken transaction inputs in the ledger to debits and transaction outputs in the ledger to credits as Antonopoulos [2017, 18] does.

[21] In bitcoin transactions, one spends the entirety of the previously unspent bitcoin from a previous transaction. But one can send a specified amount back to the same address (which is discouraged for security reasons) or to another address for which one has the private key.

out and spent. In other words, A4 has an *amount* of two bitcoin but not two *individual* bitcoins. So there are no individual bitcoins with which anything like transaction histories can be identical.

The point is not that bitcoins are fictional and that fictional entities lack identity criteria, or anything like that. For the sake of argument, let's grant that bitcoins are fictional entities and that some fictional entities like Sherlock Holmes have identities within their fictional universes. Rather, the point is that the "two bitcoins" in A4 are not differentiated entities, even within any supposed fiction. Even within the supposed fiction, we have a quantity of substance without having a number of distinguishable individuals whose sum equals that quantity. For these reasons and others, I've argued elsewhere that bitcoin is a fictional substance.[22]

An analogy may help. Suppose you deposit in your previously empty savings account a $50 check from your parents and a $50 check from your in-laws. Your account now has $100. Asking which bitcoin in A4 came from A1 and which came from A2 is like asking which of the hundred dollars came from your parents and which came from your in-laws. The question falsely presupposes that each digital dollar is marked as an individual in the bank's ledger. Similarly, the Chain Definition falsely presupposes that each bitcoin is marked as an individual in the bitcoin ledger. Just as there is no fact of the matter about the source of individual digital dollars in your account, there is no fact of the matter about the source of the individual bitcoins in A4. In both cases, we have quantities without individuals.

Stage Three. It gets worse for the Chain Definition. Not only are there no individual bitcoins to pair with entire transaction histories, there is often no entire transaction history to pair with a purported individual bitcoin. To bring this point into relief, we will proceed to the next stage of transactions. In this third stage, A4 sends one bitcoin to each of A5 and A6:

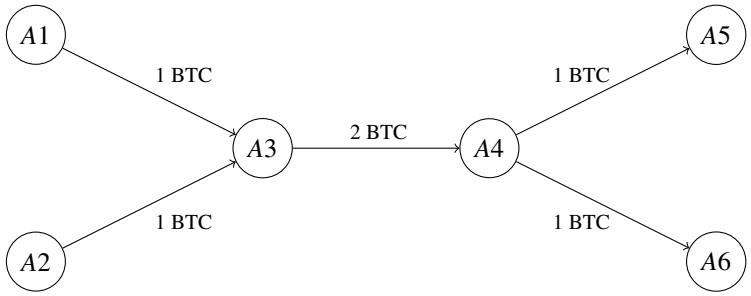

Figure 3. A transaction which halves an amount of 2 BTC across two addresses.

As we covered in Stage 2, there are not two individual bitcoins in A4 but only a quantity of two bitcoin. As a result, there is no fact of the matter about *which* bitcoin arrives in A5 rather than A6. So there is no fact of the matter about whether the bitcoin in A1 ultimately arrives at A5 or A6. Likewise, there is no fact of the matter about whether the bitcoin in A2 ultimately arrives at A5 or A6. Now, if the bitcoin ledger did mark bitcoins as individuals, we could discover whether the bitcoin at A5 took the A1-A3-A4-A5 path or the A2-A3-A4-A5 path. But there is simply no fact of the matter about whether the bitcoin in A5 took either path. Similar remarks apply to the bitcoin in A6. In the idiom of the bitcoin community, this series of transactions has *mixed* the bitcoin initially sitting separately at A1 and A2. Mixing produces a metaphysical and not merely epistemic indeterminacy. We don't just not know which path a bitcoin has taken. There simply is no such path. In the present case, neither the bitcoin in A5 nor the bitcoin in A6 has an entire transaction history. Hence, neither bitcoin has an entire transaction history with which it could be identical.

[22] Warmke [ms].

Mixing is important not only because it falsifies the Chain Definition. It also enhances bitcoin fungibility. The features that distinguish quantities of bitcoin from each other pre-mixing get smeared across the quantities of bitcoin post-mixing. Now, in the Stages above, mixing occurs over a series of transactions. But it also frequently occurs within single transactions. The CoinJoin method developed by Gregory Maxwell [2013] has been advertised as enhancing privacy because it blurs the connection between users and their addresses. In such a transaction, multiple users combine amounts of unspent bitcoin and split them into chunks having the same amount back to themselves across new addresses. According to Maxwell, this blurring between user and address "is what makes CoinJoin possible." CoinJoin enhances privacy precisely because it enhances fungibility by smearing the sources of transaction inputs across all the mixed outputs.

Within CoinJoins and other multi-output transactions, inputs do not pair up explicitly with outputs as if to say "Input 2 is the source of bitcoin in Output 3." Although every input claims one or more *previous* transaction outputs as its source of bitcoin, outputs do not similarly claim inputs *within* transactions as their source of bitcoin. Hence, we could track bitcoin from a particular input through a particular output only if amounts of bitcoin themselves had traceable identifiers. But not a single satoshi has such a traceable identifier. So we cannot track a single satoshi in a well-constructed CoinJoin transaction from an input address to an output address. Coinjoins increase the fungibility of bitcoin because it smears the transaction history of each input across every output. But this smearing is a general feature of mixing and can also occur over a series of non-CoinJoin transactions, as we witnessed in Stages 1 through 3.

3.1. Assessment

The Chain Definition fails because it falsely presupposes that the bitcoin ledger marks bitcoins as individuals. But without individual bitcoins, we don't have entire transaction histories either. What could a bitcoin's entire transaction history be except for the path that a specific bitcoin takes through a series of transactions? In general, a history of an individual piggy-backs ontologically on that individual. No individual, no history. So if there are no individual bitcoins, there are no entire transaction histories either. Consequently, the Chain Definition faces double jeopardy. The Chain Definition not only falsely presupposes that the bitcoin ledger marks bitcoins as individuals. It also identifies them with transaction histories of the sort that bitcoins often fail to have precisely because bitcoins are not individuals.

Now, perhaps the Chain Definition goes slightly too far in identifying bitcoins with entire transaction histories. Kroll et al. [2013] instead say that "a bitcoin is a fixed-value cryptographic object *represented* as a chain of digital signatures over the transactions in which the coin was used."[23] Something can *represent* something else without being *identical* to it. Instead of identifying bitcoins with transaction histories, perhaps we should read Satoshi as offering a kind of simplifying assumption of the sort we often find in math, philosophy, science, and so on. Theorists often use simplifying assumptions when they use sets of objects to model the meanings of words, sets of possible worlds to model propositions, n-tuples of real numbers to model locations in n-dimensional real space, sets themselves to model numbers, etc. Many who model in this way would reject or at least refuse to accept the claim that the thing modeling is the thing modeled. Satoshi might have had something similar in mind.

However, the very features of bitcoin transactions which sink the Chain Definition also sink the claim that each bitcoin is "represented as a chain of digital signatures over the transactions in which the coin was used." Without names or markers for individual bitcoins, the blockchain does not encode the sort of transaction histories with which we could represent individual bitcoins. Mixing patterns like those depicted in Figures 2 and 3 illustrate this point well.

With the exception of coinbase transactions, the bitcoin protocol ensures that each transaction's outputs preserve the total amount of bitcoin from its inputs, much like the laws of thermodynamics ensure that, in an isolated system, energy is not created or destroyed. But, in preserving these

[23] The emphasis is mine.

amounts, the network uses transactions that transfer quantities of bitcoin, not individuals that are bitcoins. In so doing, the bitcoin network resolves what we may call the *divisibility dilemma*. To prevent double spending, one might have thought that each smallest unit in an electronic cash system needs an identity on the ledger. But, then, if we want transactions with highly divisible amounts, transactions would require many millions of signatures per transaction.[24] The strain on the network would be tremendous. A system that tracks the identities of individual units faces a trade-off between efficiency and divisibility. Satoshi achieved a balance between efficiency and divisibility by adopting a distributed ledger of the kind described by Wei Dai [1998] that tracks quantities instead of individuals. So the Chain Definition ultimately obscures one of Satoshi's smarter engineering decisions.

4. INSTRUMENTS AND QUANTITIES

If bitcoins are not chains of digital signatures, how should we understand Satoshi's definition of electronic coins as chains of digital signatures? The definition did not occur in a vacuum. Bitcoin built on previous attempts at electronic cash, and these attempts help illuminate the reference to electronic coins in the bitcoin whitepaper.

Bitcoin's prehistory involves a number of proposals for electronic money. These proposed versions of electronic money naturally inherited features from their non-electronic counterparts. To illustrate, consider a signed check for $20. The check is a *financial instrument* that signifies a quantity of $20. And, except in special cases, what signifies isn't identical to what's signified. The name 'Einstein' signifies the man, Einstein. And the name isn't the man. The same generally holds for financial instruments and the quantities they signify. In the case of our $20 check, the signifying instrument and the signified quantity have different properties. Whereas the check is an individual and bears a traceable identifier, the quantity of $20 itself has no traceable identifier. Different properties, different identities.

One might object that the signified quantity does have a traceable identifier. It inherits the identifier of the signifying check. But the connection between the quantity and the identifier is contingent, and we can easily sever it. Imagine depositing the check in a previously empty bank account and then tearing up the check. The quantity survives in the bank account. But the check is no more. And the check's identifier no longer follows the quantity that the check signified.

The distinction between signifying instrument and signified quantity does not disappear when we digitize the instrument. For example, in the 1980s, David Chaum created the eCash system for transferring electronic money in the form of Cyberbucks.[25] In this system, Alice creates a random string of symbols to serve as a numeric note analogous to a paper check with a check number. Then, she sends the note to a bank for a digital signature that determined the note's monetary value. If she wanted a note for 20 Cyberbucks, for instance, the bank would withdraw that amount from her account and sign the numeric note with the private key reserved specifically for signing 20-Cyberbuck notes.

Like checks and their signified quantities, each note differed from the quantity of Cyberbucks signified. The notes were strings of symbols and had identifiers in the form of their randomly chosen note numbers. They were digital financial instruments. But whereas the digitized instrument bore a traceable identifier in the form of its note number, the signified quantity of Cyberbucks had no such traceable identifier. Now, any particular quantity of Cyberbucks might have had a contingent tie to the note that signified it. But these ties were contingent and easily severed. We could deposit the note in a previously unused bank account where the quantity of Cyberbucks would persist but the numeric note and its identifier would not.

The distinction between the signifying financial instrument, on the one hand, and the signified quantity, on the other, holds for bitcoin's other predecessors, too. For example, the RPOW tokens

[24] As Tatsuaki Okamoto [1995, 439] correctly observes, "a system in which a coin worth $5367 consists of 5367 $1 coins is a rather unwieldy and inefficient divisible cash system." Compare Chow [2007, 151].

[25] Chaum [1983, 1985, 1992], Chaum et al. [1988].

developed by Hal Finney [2004] were financial instruments in the form of a string of symbols. Whereas physical tokens signify a quantity with numerals stamped on their sides, Finney's RPOW tokens signified a quantity with symbols encoded in their bits. Unfortunately, Finney didn't use a special unit of account like the Cyberbuck. So let's call the unit an *rpow*. Then, we can say that Finney's RPOW tokens signified quantities of rpow. Again, quantities of rpow were not themselves the RPOW tokens, not even if an RPOW token signified a lone rpow.

The distinction between instrument and quantity also holds for bitcoin. But if the blockchain represents quantities of bitcoin, what are the signifying financial instruments? The financial instruments are unspent transaction outputs or *UTXOs*. UTXOs are transaction outputs that remain unspent. They are like physical checks that have yet to be signed and deposited. Unsurprisingly, quantities of bitcoin and UTXOs differ in important ways. Whereas each UTXO has an identifier in the form of its index number and the ID of the transaction in which it appears, no quantity of bitcoin has such an identifier. Furthermore, while spending a UTXO "destroys" it, its total quantity of bitcoin persists in one or more quantities signified by one or more new UTXOs.

With the distinction between signifying instrument and signified quantity now in hand, we may begin to diagnose the Chain Definition. Given bitcoin's prehistory, we might expect Satoshi's definition of electronic coins to concern digitized instruments, like Chaum's notes and Finney's RPOW tokens. Then we could diagnose the Chain Definition with substituting the original reference to the coins, the digitized instruments we call UTXOs, with a reference to individual bitcoins. Since, as I've argued, the blockchain does not represent bitcoins as individuals, and since digitized instruments differ from the quantities they signify, we would expect such a substitution to fail. If only things were so simple.

In my view, 'unspent transaction output' is often ambiguous between the chunk of code that signifies a quantity of unspent bitcoin and the signified quantity itself. Let's reserve 'UTXO' for the chunk of code. And let's call the signified quantity of bitcoin an *unspent quantity*. As I previously mentioned, spending a UTXO destroys it. One spends and destroys a UTXO by providing the appropriate digital signature in another transaction's input. So UTXOs don't persist through a chain of digital signatures. A chains of digital signatures coincides with a trail of destroyed UTXOs. Though UTXOs don't persist through a series of transactions, quantities of bitcoin sometimes do.

Like pouring a cup's contents into another without spilling, early bitcoin users could often transfer a UTXO's entire unspent quantity without spending any of it in a transaction fee. Feeless transactions occur rarely now, so a UTXO's unspent quantity is now typically a flash in the pan. Even so, each UTXO represents an unspent quantity which has at least vacuously persisted through a chain of at least one digital signature. So, with some success, we can model a UTXO's particular quantity of unspent bitcoin with the chain of digital signatures that has preserved that quantity back to the transaction in which it resulted by combining smaller quantities, by splitting a bigger quantity, or by serving as a mining reward. Consequently, charitably interpreting Satoshi leads us to the conclusion that the electronic coins in Satoshi's definition are probably best understood as unspent quantities of bitcoin—not the UTXOs which contingently and often temporarily signify them. However, we should note that this interpretation departs from the usual ways of speaking about earlier attempts at electronic money because it applies a label typically reserved for a financial instrument to the kinds of quantities that those instruments signify. This departure doesn't seem to me a serious objection against the proposal, since Satoshi sometimes uses terminology inconsistently. Plus, the main alternative to my proposal—that the electronic coins are UTXOs—makes much less sense. But, whatever the case may be, these electronic coins are definitely not individual bitcoins or satoshis.

5. CONCLUSION

In the bitcoin whitepaper, Satoshi says that "electronic coins" are chains of digital signatures. Many have since endorsed the Chain Definition, inferring from Satoshi's claim that bitcoins are chains of digital signatures. But the inference fails. As I argued in Section 3, the Chain Definition falsely presupposes that the bitcoin ledger marks bitcoins as individuals. The ledger tracks quantities of bitcoin, not individual bitcoins. But if the electronic coins in Satoshi's definition are not bitcoins,

what are they? Many would say that they are UTXOs. But we must distinguish UTXOs from the quantities of bitcoin that they signify. And I've argued that the electronic coins in the whitepaper are probably best understood as these quantities of bitcoin and not the UTXOs that signify them.

Unfortunately, the Chain Definition has begun to spread like an interdisciplinary wildfire. Like many wildfires, the Definition's influence has been both destructive but understandable. It has been destructive not only because it has sown confusion across many disciplines, but also because some used it as a key premise to defend the contentious bitcoin cash hard fork in 2017. Nonetheless, the definition's influence has been understandable since Satoshi uses "coin" equivocally very early on. For example, as early as the v0.1 Bitcoin software release, Satoshi refers to the 21,000,000 maximum supply of bitcoins as a cap on "coins."[26] This equivocal use of terminology made some confusions almost unavoidable, especially when so many academics with interests in bitcoin operate at an altitude far above its technical machinery.

In a new and highly interdisciplinary field like cryptoeconomic systems, more understandable confusions await us. The field will not succeed without legal scholars, mathematicians, economists, and computer scientists speaking across disciplinary boundaries.[27] But this very condition for success will draw some to trespass into disciplines for which they have little or no training.[28] It will lead others to talk past each other with superficially similar terminology. And some will attempt to exploit the confusion for personal gain. Going forward, those who specialize in clarifying concepts and drawing distinctions could play an invaluable role.[29]

ACKNOWLEDGMENTS

For comments and discussion, thanks to Wassim Alsindi, Nic Carter, and participants in a workshop at the University of Groningen, especially David Dick, Amin Ebrahimi, Francesco Guala, Frank Hindriks, Benjamin Neeser, Asya Passinsky, and Joakim Sandberg.

REFERENCES

2012. *Virtual currency schemes*. European Central Bank. Frankfurt am Main.

Cuneyt Gurcan Akcora, Matthew F. Dixon, Yulia R. Gel, and Murat Kantarcioglu. 2018. Blockchain Data Analytics. *Intelligent Informatics* (2018), 4.

Benjamin W Akins, Jennifer L Chapman, and Jason M Gordon. 2014. A whole new world: Income tax considerations of the Bitcoin economy. *Pitt. Tax Rev.* 12 (2014), 25–56.

Andreas Antonopoulos. 2017. *Mastering Bitcoin: unlocking digital cryptocurrencies* (2nd. ed.). O'Reilly Media, Inc. First published in 2014.

Nathan Ballantyne. 2019. Epistemic Trespassing. *Mind* 128 (2019), 367–395.

Ole Bjerg. 2016. How is bitcoin money? *Theory, Culture & Society* 33, 1 (2016), 53–72.

Andrea Borroni. 2016. Bitcoins: Regulatory Patterns. *Banking & Finance Law Review* 32, 1 (2016), 47.

David Chaum. 1983. Blind signatures for untraceable payments. In *Advances in cryptology*. Springer, 199–203.

David Chaum. 1985. Security without identification: Transaction systems to make big brother obsolete. *Commun. ACM* 28, 10 (1985), 1030–1044.

David Chaum. 1992. Achieving electronic privacy. *Scientific american* 267, 2 (1992), 96–101.

David Chaum, Amos Fiat, and Moni Naor. 1988. Untraceable electronic cash. In *Conference on the Theory and Application of Cryptography*. Springer, 319–327.

Sherman SM Chow. 2007. Running on karma–P2P reputation and currency systems. In *International Conference on Cryptology and Network Security*. Springer, 146–158.

[26] Nakamoto [2009].

[27] Voshmgir and Zargham [ms].

[28] Ballantyne [2019].

[29] Though I have philosophers foremost in mind, see Walch [2016, 2017, 2019] for examples from a legal perspective.

Wei Dai. 1998. b-money. (1998). http://www.weidai.com/bmoney.txt

Hal Finney. 2004. Rpow: Reusable proofs of work. (2004). https://nakamotoinstitute.org/finney/rpow/

Hal Finney. 2008. Bitcoin P2P e-cash paper. (2008). https://www.metzdowd.com/pipermail/cryptography/2008-November/014848.html

Maximilian Friedlmaier, Andranik Tumasjan, and Isabell M Welpe. 2018. Disrupting industries with blockchain: The industry, venture capital funding, and regional distribution of blockchain ventures. In *Venture Capital Funding, and Regional Distribution of Blockchain Ventures (September 22, 2017). Proceedings of the 51st Annual Hawaii International Conference on System Sciences (HICSS)*.

Yu-Long Gao, Xiu-Bo Chen, Yu-Ling Chen, Ying Sun, Xin-Xin Niu, and Yi-Xian Yang. 2018. A secure cryptocurrency scheme based on post-quantum blockchain. *IEEE Access* 6 (2018), 27205–27213.

Merve Can Kus Khalilov and Albert Levi. 2018. A survey on anonymity and privacy in bitcoin-like digital cash systems. *IEEE Communications Surveys & Tutorials* 20, 3 (2018), 2543–2585.

Joshua A Kroll, Ian C Davey, and Edward W Felten. 2013. The economics of Bitcoin mining, or Bitcoin in the presence of adversaries. In *Proceedings of WEIS*, Vol. 2013.

Rosa Maria Lastra and Jason Grant Allen. 2018. Virtual Currencies in the Eurosystem: Challenges Ahead. *Brussels, Belgium: ECON Committee, European Parliament* (2018).

Gregory Maxwell. 2013. CoinJoin: Bitcoin privacy for the real world. (2013). https://bitcointalk.org/?topic=279249

Satoshi Nakamoto. 2008. Bitcoin: A Peer-to-Peer Electronic Cash System. (2008). http://bitcoin.org/bitcoin.pdf

Satoshi Nakamoto. 2009. Bitcoin v0.1 released. (2009). https://satoshi.nakamotoinstitute.org/emails/cryptography/16/#selection-9.0-9.21

Tatsuaki Okamoto. 1995. An efficient divisible electronic cash scheme. In *Annual International Cryptology Conference*. Springer, 438–451.

Peter Van Valkenburgh. 2014. Comments to the Conference of State Bank Supervisors on the Draft Model State Regulatory Framework for Virtual Currency. (2014).

Shermin Voshmgir and Michael Zargham. ms. Foundations of cryptoeconomic systems. (ms.).

Angela Walch. 2016. The path of the blockchain lexicon (and the law). *Review of Banking & Financial Law* 36 (2016), 713–765.

Angela Walch. 2017. Blockchain's Treacherous Vocabulary: One More Challenge for Regulators. *Journal of Internet Law* 21, 2 (2017), 9–16.

Angela Walch. 2019. Deconstructing 'Decentralization': Exploring the Core Claim of Crypto Systems. In *Crypto Assets: Legal and Monetary Perspectives*. Oxford University Press.

Craig Warmke. ms. What is Bitcoin? (ms.).

Qianhong Wu, Xiuwen Zhou, Bo Qin, Jiankun Hu, Jianwei Liu, and Yong Ding. 2017. Secure joint Bitcoin trading with partially blind fuzzy signatures. *Soft Computing* 21 (2017), 3123–3134.

Joachim Zahnentferner. 2018. Chimeric Ledgers: Translating and Unifying UTXO-based and Account-based Cryptocurrencies. *IACR Cryptology ePrint Archive* 2018 (2018), 262.

Yilu Zhang. 2017. The Incompatibility of Bitcoin's Strong Decentralization Ideology and Its Growth as a Scalable Currency. *NYUJL & Liberty* 11 (2017), 556.