Who is Causally Responsible for a Cryptocurrency?

Peder Østbye[*]

Discussion Paper, February 2019
Minor revision March 2019

Abstract

Cryptocurrencies have entered the economy as alternative money, speculation objects, and as utility tokens for digital platforms. Cryptocurrencies are based on cryptography-based asset disposals broadcasted peer-to-peer to be validated in a decentralized way according to consented protocols. The organization and governance of cryptocurrencies disrupts the way things have been done by centralized institutions. This has consequences for responsibility. A centralized institution can be held responsible for its actions; however, in a cryptocurrency scheme, those actions are taken by a crowd of distributed participants not individually necessary for the outcome. This paper explores the participants' causal links to harm produced by a cryptocurrency scheme, and discusses how these causal links translate into responsibility. It is found that despite the decentralized nature of cryptocurrencies, participants cannot conceptually, legally, or morally use lack of causal links as a justification to evade responsibility.
JEL: E42, G28, K24
Keywords: Algorithms, Blockchain, Causality, Cryptocurrencies, Liability, Responsibility

1 Introduction

Cryptocurrencies have entered the economy as alternative money, speculation objects, and as utility tokens for digital platforms. Cryptocurrencies are based on cryptography-based asset disposals broadcasted peer-to-peer to be validated in a decentralized way according to consented protocols. The organization and governance of cryptocurrencies disrupts the way things have been done by centralized institutions. This has consequences for responsibility. A centralized institution can be held responsible for its actions; however, in a cryptocurrency scheme, those actions are taken by a crowd of distributed participants not individually necessary for the outcome. This paper explores the participants' causal links to harm produced by a cryptocurrency scheme, and discusses how these causal links translate into responsibility. It is found that despite the decentralized nature of cryptocurrencies, participants cannot conceptually, legally, or morally use lack of causal links as a justification to evade responsibility.

The possible harms from cryptocurrencies have been well studied in the literature.[1] The legal responsibility of various participants in a cryptocurrency has been studied in some detail in the literature.[2]

[*]Dr. Philos. Special Adviser, Norges Bank. This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the author and do not necessarily reflect those of Norges Bank. For correspondence, please use p.e.oest@gmail.com.

[1]An overview is provided by Østbye (2018b)

[2]Zetzsche et al. (2017) provides a general assessment. See Walch (2018) and Seira (2018) for a discussion as to whether

The present paper also shares some topics with responsibility for algorithms in general, for software, and, in particular, open source software.[3] However, cryptocurrencies are distinguished from traditional algorithms by the involvement of participants in a decentralized manner for the operation. Hence, the responsibility of cryptocurrency participants is not fully captured by the general study of responsibility for algorithms. The present paper makes a novel contribution by specifically addressing causality issues in the responsibility assessments associated with the decentralized operation of cryptocurrencies.

Section 2 will provide a brief overview of cryptocurrency technology and the roles of the various participants in a cryptocurrency scheme. Section 3 will describe possible harms produced by cryptocurrency schemes, and discuss how the legal system generally addresses such harms. Section 4 will discuss the casual link between participants' actions and harms from a cryptocurrency scheme. Section 5 and Section 6 will discuss legal and moral responsibility in light of the causal links. Section 7 will conclude.

2 Cryptocurrencies and the role of the various participants

Bitcoin was launched in 2009, but documentation was available already in 2008. The creator or creators of Bitcoin are unknown to the general public. The Bitcoin white paper, Nakamoto (2008), was written under the pseudonym Satoshi Nakamoto. The intention behind Bitcoin expressed in the white paper is that "[w]hat is needed is an electronic payment system based on cryptographic proof instead of trust, allowing any two willing parties to transact directly with each other without the need for a trusted third party." As a disruptive innovation and from the perspective of competition, it is a welcome potential challenger to banks and other financial service providers.

Many of the cryptocurrencies introduced in the aftermath of Bitcoin seek to improve upon shortcomings in Bitcoin. For instance, scale and increased anonymity have been popular features to improve upon.[4] Lately, there have also been initiatives to create cryptocurrencies with more stable value.[5] Some cryptocurrencies have been created by known natural or legal persons, and some are more or less centralized in terms of governance and permission-based access. For instance, Ripple is intended to improve the efficiency of settlements between financial institutions.[6] Many cryptocurrencies serve as utility-tokens to fuel service platforms. Ethereum is an example, providing a complete programming language on the platform, which can be used for smart contracts (automated contracts).

Cryptocurrencies are based on two main principles: cryptography-based asset disposal and distributed ledgers. Cryptography-based asset disposal means that cryptographic keys are used to sign transactions and verify ownership.[7] The transaction sender signs a transaction with a secret private key, and a corresponding public key can be used to validate that the transaction has been signed by the corresponding

protocol developers in public cryptocurrencies can be considered fiduciaries. See also Østbye (2018a) for a discussion of antitrust liabilities.

[3] Tjong Tjin Tai (2018) provides a review of liability for algorithms. Choi (2018) studies liability for software. Bahn and Dressel (2006) provides some aspects on liability associated with open source software.

[4] For instance, Litecoin seeks to improve scale and speed relative to Bitcoin. Dash, Cloakcoin, Monero, and Zcash, among others, seek to improve privacy. See Duffield and Diaz (2014) and Cloak (2018) for documentation of Dash and Cloakcoin, respectively. Both also improve scalability. Sasson et al. (2014) is the original whitepaper for Zcash. Improved anonymity is achieved by various sorts of coin-mixing arrangements that prevent transparency with respect to the sender and receiver of coins.

[5] Such as the cryptocurrencies Dai and Basis. See Dai (2017) for the Dai whitepaper and see Al-Naji et al. (2018) for the Basis whitepaper. See also Østbye (2018b) for a critique of the mechanisms relied upon by Al-Naji et al. (2018). See also Blockchain Luxemburg (2018).

[6] See https://ripple.com/.

[7] Cryptography-based asset disposal is not an invention to be credited to cryptocurrencies. Public-key cryptography has been available for decades and has been suggested in variants of digital cash since the 1980s.

private key.[8] The cryptographic-asset disposal also allows for various mechanisms for conditional disposal, such as the execution of smart contracts. As it is private keys and not personal identities that determine control of assets, and there is no need to link real-world identities with private keys, the systems are pseudo-anonymous.[9]

However, digital assets are easy to copy, entailing a double-spending risk. A traditional solution is to rely on trusted third parties to maintain registers. The main invention associated with cryptocurrencies is elimination of the need for a trusted third party by letting decentralized operators validate transactions and maintain the integrity of the register. This is called distributed ledger technology (DLT). DLT protocols are designed to maximize the incentives of the decentralized operators to maintain the integrity of the ledger in compliance with the protocol governing the cryptocurrency. The DLTs in various cryptocurrencies are designed such that they facilitate:

- Detection: the transparency of the ledger facilitates detection of dishonest behavior.
- Punishment of dishonest behavior: Profits from validation will probably be lost in case of dishonest behavior. For many cryptocurrencies, the protocol allows a reward for validation in terms of transaction fees and newly minted coins in the validated cryptocurrency.

In simplified terms, we can say trust is maintained by the participants' interest in future trust in the cryptocurrency. By such a design, operators given the authority to validate transactions have incentives to do so honestly to maintain the value of the reward.[10]

To gain more insight into DLT, the DLT-framework provided by Rauchs et al. (2018) is useful. A DLT can be considered to consist of three different layers: the protocol layer, the network layer, and the data layer. The data layer deals with the representation data; this will not be discussed in detail here.

The protocol layer of a DLT consist of the basic rules of the system, including the governance structure. This layer set out the mechanisms relied upon for the operations on the other layers. This includes what cryptographic techniques to rely upon, how data are represented, who can validate transactions and the reward for such validation, what kind of consensus is needed for a transaction to be part of the ledger, and procedures for protocol changes. A main distinction between various cryptocurrencies is whether they are public or permission-based – that is, whether there are rules restricting participation. Such rules may include who can make protocol changes. For instance, protocol development may be reserved to a person, a firm, or a consortium. A common trait of many public cryptocurrencies is that the protocols are implemented in open source software according to open source licenses, and the development share many commonalities with open source development in general.[11] Some organizational structure is needed to coordinate the development, as is the case in for example Bitcoin, where the protocol includes mechanisms to communicate improvement proposals via the blockchain. Sometimes institutions are established to coordinate the development, such as the Ethereum Foundation. Pre-mined coins and ICOs can fund such establishments. Some cryptocurrencies are trying to improve the integrity of protocol development by establishing "constitutions" and other advanced governance structures for protocol development.[12] Cryptocurrency protocols may also include structures to finance development, such as voting procedures

[8] The public key is generated from the private key with a non-invertible function, which is supposed to make this system secure. Non-invertibility is meant in a practical, not mathematical, sense.

[9] However, as so-called network analysis can be used to infer identities from limited real-world information, several cryptocurrencies seek to improve anonymity by variants of mixing to hide the senders and receivers of transactions. See, for instance, Conti et al. (2017).

[10] For a lengthier description of cryptocurrency technology, see Narayanan et al. (2016).

[11] Open source licenses are supposed to preserve the open source characteristic of the products. Different licenses are more or less restrictive when it come to preserving the open source of the further development. Another important purpose is to disclaim liability.

[12] See https://medium.com/coinmonks/a-deep-dive-into-eos-governance-49e892eeb4a2 for a discussion of the constitution behind the cryptocurrency EOS.

among validators to finance development with newly minted coins.[13] As with open source development in general, single persons, often the initiator or an early developer, may gain an informal leader role in the project.

However open-source development is not exclusive - this is the nature of open source development based on open source licenses, and is what distinguish it from proprietary development. Disagreements with respect to development seems inevitable in many open source projects. Such disagreement may result in forks, where separate fractions of developers go each their way with competing variants of the product. This has happened with cryptocurrencies as well. Bitcoin has forked several times, where disagreements over protocol development has been the issue. The fork that created Bitcoin Cash in 2017 is an example of such a disagreement. Another example is The DAO, which was an investor-directed venture capital fund implemented with smart-contracts on Ethereum. The DAO failed spectacularly in 2016 after a hacker exploited a bug to get hold over much of the fund. The validators and protocol developers of Ethereum cooperated to reverse the exploit. This was controversial as it can be considered to be in conflict with the decentralized nature of Ethereum. As a consequence, Ethereum forked into Ethereum and Ethereum Classic, where the latter didn't reverse the exploit.

With the protocol as a basis, the operation of a cryptocurrency happens at the network layer. In permission-based cryptocurrencies, the protocol might put some constraints on who can perform the various functions, and who will get access to certain information. However, in public cryptocurrencies there are no such constraints. A main component of this layer is the communication of transactions. Transactions are broadcasted by a node on the network and propagated according to peer-to-peer technology. Normally, the default-software of nodes will prevent invalid transactions being propagated further. In privacy-enhancing cryptocurrencies, the nodes may also facilitate various kinds of coin-mixing to blur the details of the transaction. Another component of the network layer is the validation of transactions and the adding of transactions to the ledger. There is a great variety of check-and-balances mechanisms to facilitate integrity. As mentioned above, they are based on detection and punishment of dishonest behavior and there is an ever-increasing number of proposals for doing this efficiently informed by game-theory.

For Bitcoin and many other cryptocurrencies, so-called proof-of-work (PoW) is used to facilitate integrity. Competitive block validators collect transactions in a block to add into the ledger in a block. Each new block is pointing to a hash[14] of the previous block.[15] Hence, the blocks are chained together in a blockchain.[16] To be allowed to propose a candidate block to the blockchain, the validator must be the first to solve a computationally costly puzzle. This puzzle consist of assembling the hash of the previous block, a hash of the transactions[17] in the candidate block, some other inputs, and a freely chosen nonce into a hash-function, such that the resulting hash falls below a certain threshold.[18] This nonce is hard to

[13] The cryptocurrency DASH has established an advanced system for governance, see https://docs.dash.org/en/stable/governance/understanding.html

[14] A hash function generates a non-invertible fixed-length output from an input in the same manner as a public key is generated from a private key

[15] To be precise, the header of the previous block.

[16] Alternative implementations of DLT, not based on blockchains, have also been developed, as means to maintain the integrity of a distributed ledger. One alternative is to represent the ledger as a directed acyclic graph (DAG). IOTA is an example of a cryptocurrency using DAG for maintaining the distributed ledger, as described in the whitepaper Popov (2017). The consistency of the ledger is preserved by letting the transactions form a DAG of consistent transactions. Inconsistent transactions, such as double spends, will not be included in the consensus graph. The integrity is maintained by assigning transactions weights dependent on their centrality in the graph. The principle is that fraudulent transactions are not likely to be given sufficient weight to remain in the consensus graph.

[17] Organized as a so-called Merkle tree, to make the search for specific transactions efficient.

[18] This can be described as solving an equation where a valid hash is the solution and the nonce is the unknown. Let $b_t = (h(b_{t-1}), m(T), n, ..)$ represent the header of the candidate block; h, a hash function; m, the hash of a Merkle tree of the transactions, T, to be included in this candidate block; and n, the nonce. To be allowed to add a candidate block to the

find, but its validity is easy to verify. The reward for such validation is that the validator can include a fixed amount of newly minted bitcoins to a chosen address (normally of the validator itself or a mining pool in which the block validator participates) and transaction fees set at the discretion of the senders. Since the block-validator is rewarded newly minted coins, the block-validators are commonly referred to as miners. The newly minted coin reward and the transaction fees are lost if the block does not become part of the consensus chain.

Why should a validator be honest and propose a valid block to the blockchain? The first finder of a valid nonce gets the privilege of adding its candidate block to the blockchain. However, it is not guaranteed to be a part of the blockchain. Whether it becomes part of the consensus chain depends on future block validators building their blocks on this particular block. Assuming that future block validators will be honest and only build upon honest blocks, a validator has strong incentives to be honest and follow the protocol. Attempts to violate the protocol rules will render the block abandoned and the potential reward lost. According to the Bitcoin-protocol and the protocol in many other cryptocurrencies, the longest chain is the valid one. A consequence of this is that if a validator wants to include transactions not consistent with the previous blocks in a new block, the validator would then need to alter the whole chain, back to a block consistent with the fraud, possibly the genesis block, to get hashes consistent with the present block of transactions. In addition, the attacker would need to compete with the honest chain to get the longest chain. This would be very costly, and more costly the more computing power is devoted to honest validation.[19] An example of such an attack is the double-spending attack, where the attacker first spend some coins and then after the performance of the counter-party replaces the ledger with a blocks without the transaction.[20]

Various alternatives to PoW exist that may be used in combination with PoW. One commonly applied scheme is proof-of-stake (PoS).[21] PoS means, simplified, that a stake in the currency impacts the influence on validation. Such influence may be implemented by letting the stake reduce the difficulty of solving a puzzle as in PoW, or by letting the stake influence voting power in a byzantine fault tolerant (BFT) consensus scheme. Normally, validators must put their stakes at risk of being lost if the validator's proposal block is not accepted by the consented ledger. A high stake provides incentives to maintain the value of the currency. As long as the stakes in the currency are not too concentrated, a dishonest block-validator will face a risk, as an invalid block is not likely to be included in the chain by the honest validators. More sophisticated validation schemes also take into account the age of the coins held. Other

blockchain, the validator must be the first to solve the computationally costly puzzle $h(b_t) < C$, where the nonce, n, is the unknown and C (positive) is the threshold. To solve this puzzle, the candidate block validator must perform many trials, as the hash function is not invertible and each trial contains minimal information about the solution. In other words, the hash algorithm is made such that one calculation gives as little information as possible regarding the solution to the puzzle. The lower the threshold, the harder it is to find a solution. To maintain the difficulty as the technological computational capacity increases, reductions in thresholds are implemented in the Bitcoin protocol. Although the main rule so far has been that the difficulty increases, it is also possible that the difficulty level reduces if the average time taken to find a new block increases. The difficulty is set such that a new block is found on average every 10 minutes.

[19] Nakamoto (2008) provides a mathematical model representing this competition and provides probabilities for the success of such an attack as a function of share the computing power controlled by the attacker. See Østbye (2018c) for a discussion of the validity and reliability of this model.

[20] This associated with a 51-percent attack, referring to the situation where an entity controlling more than 51 percent of computing-power can control the ledger. However, in a PoW scheme such an attack can be successful by an entity controlling less than 51-percent of the computing power – it is just that it is unlikely. A question explored is whether a group with less than 51 percent validation power still can perform successful attack by other means, such as the exploitation of a central network position relative to other block validators. See, for instance, Conti et al. (2017) for a survey of possible attacks on the Bitcoin blockchain. See also Narayanan et al. (2016), Chapter 5.

[21] For a detailed analysis of PoS schemes, see Bentov et al. (2016).

variables could also be taken into account.[22]

The PoW and PoS schemes are suitable for public blockchains as they provide resistance against so-called sybil-attacks, which involve the creation of several nodes to gain influence. Such a risk is less of a concern for permission-based DLTs, where the validators can be fixed by the protocols. Such systems may rely more exclusively on BFT voting procedures among validators to provide robustness against faulty nodes.

For many cryptocurrencies, interacting with the outside world is necessary; the cryptocurrency operations need information exogenous to the system such as some state of the world for the execution of smart contracts. Such information can be provided by oracles, and the information provided by oracles must be subject to incentive mechanisms to facilitate truthful information. In the end, it is the validators that execute the transactions based on the information provided by the oracles.

Finally, the ecosystem surrounding cryptocurrencies and their participants is briefly discussed, although the primary focus of this paper is the participants within a cryptocurrency ecosystem. There are several ways users can acquire cryptocurrencies from the owners. Such acquirement can follow from, inter alia, bilateral private exchange, brokers and professional exchanges, and as payment for goods, services, and labor. In addition to the direct trade with cryptocurrencies, there is an ecosystem of third-party service providers, such as wallet providers for users to administer their cryptocurrencies, payment service providers, consulting services, and investment services. Such services allow users to not participate as nodes in the system, since agents can appear as custodians for the users with their own nodes. Such custodians share similarities with banks and, in fact, some traditional banks are providing such services.

3 Harms from a cryptocurrency scheme

This section will discuss potential harms from the operation of a cryptocurrency scheme. There are other harms associated with cryptocurrencies such as harm from investment fraud associated with trading, and harm from cyber-attacks on exchanges holding cryptocurrencies for customers. Such harms will not be discussed further. We can distinguish between two types of harms associated with the operation of a cryptocurrency. Firstly, we have internal harm within a cryptocurrency ecosystem, where some participants' actions cause harm to other participants. Secondly, we have external harm beyond a cryptocurrency's ecosystem, which affects persons not directly involved in the cryptocurrency.

A typical example of internal harm is validators exploiting other users by charging excessive transaction fees, excluding the transactions of certain persons, and by performing various forms of double-spending attacks. Participants may also suffer from an operational failure that may result in data loss, loss of funds, compromised privacy, or lack of stability in a cryptocurrency designed for stability. Such harms may be subsequent to poor protocol design/development or smart-contract design.

There are several types of external harm from a cryptocurrency scheme. We can distinguish between harms to individuals and harms to the society more abstractly – that is – harms to public interests. The operation of a cryptocurrency scheme may result in direct harm to persons completely outside the cryptocurrency scheme. An example of such harm may be persons or entities subject predictions in prediction markets organized as cryptocurrencies. Prediction markets may create incentives to accelerate a predicted

[22] Analytically, we can describe a general validation mechanism as $p_i = p(W_i, S_i, A_i, C_i, ...)$, where p_i is the probability that agent i is chosen as a block-validator. W_i is the invested work of agent i, typically solving a cryptographic puzzle as in Bitcoin. S_i is the stake of agent i. A_i is the age of the stake of agent i. C_i is the contribution of agent i according to some measurable criteria. This could be the provision of services to the platform the currency is operating upon, such as code and protocol development. A way to implement protocols with a variety of inputs to the validation probability is to make the difficulty of the PoW dependent on these inputs.

event. In the worst case, a prediction market prediction may result in an assassination. External harm to individuals may also be present if personal information about persons not involved in a cryptocurrency is stored on the ledger, and an operational failure result in the data being compromised, for instance if a firm uses a distributed ledger to store its customer-data. Information subject to intellectual property rights held by persons outside a cryptocurrency ecosystem might be distributed via the distributed ledger.

Several possible harms to public interests have been identified. A major concern for financial regulators and prosecutors is the use of cryptocurrency schemes to facilitate crime, in particular money laundering.[23] Prediction markets, as just mentioned, is another case, as the predicted event might involve a crime. Financial regulators are also concerned with the potential of cryptocurrencies to threaten financial stability.[24] Although this concern is mainly related to the investment aspects of cryptocurrencies, it also has an operational side. If a cryptocurrency gets a prominent role in the payment system, either for retail payments or for settlement between financial institutions, the consequences of a failure may be on a systemic scale with adverse effects for the economy.

The use of energy in PoW schemes is also a concern for environmental reasons. There are various estimates in the literature on the energy use involved in the validation of bitcoin-transactions.[25] Estimates have shown that the energy-use are comparable to the energy consumption of Ireland[26] and Belgium.[27] However, it would be an invalid inference to conclude outright that this is a harm to society different than the harm produced by anyone that uses energy. The relevant question is if market failures are present such that the environmental harm is greater than the benefits. This may happen in (at least) two ways. One may be that validators do not pay the full cost of electricity, reflecting also the harm to the environment. This might happen if market prices are too low, or if the validators find ways to steal electricity, such as by hacking other people's computers for validation (cryptojacking). Another way may be that the technology used for validation is inefficient, but market failures prevent the transformation to more efficient technologies. For instance, the network effects associated with certain energy-demanding cryptocurrencies, such as Bitcoin, may create entry barriers for other cryptocurrencies.[28]

Many of the harms listed above are addressed by general and specialized laws and regulations, such as general criminal law and tort law, contract law, fiduciary law, anti-money laundering regulations, and antitrust law. Normally, if the harms were produced by a single legal entity, either physical persons or companies, these would have been subject to criminal and civil liabilities.[29] In some cases, they would also be subject to licensing regimes involving regulatory scrutiny. Hence, if a centralized entity performed the same task without a license, they could be sanctioned for that fact alone. At its face, such laws and regulations also apply to participants in a cryptocurrency scheme.[30] However, liability requires that there is someone to hold liable. Liability, in turn, requires some sort of causal link. When actions are performed by single institutions or legal persons, such causal links may be apparent. However, the operation of a cryptocurrencies is performed by a crowd on a distributed network, where each crowd member can claim their contribution as insignificant and unnecessary. The question as to whether this removes a causal link necessary for responsibility is the topic of the remainder of this paper.

[23] See, for example, Østbye (2018b).

[24] See, for instance, Østbye (2018b) and Ali et al. (2014).

[25] See, for instance, de Vries (2018). Rauchs et al. (2018b) provides a survey.

[26] See de Vries (2018).

[27] See Rauchs et al. (2018b).

[28] Such network-effects may be direct or indirect. They are direct in the sense that having more users increase the value of the cryptocurrency, and indirect in the sense that more validators brings more security to the cryptocurrency which increases the value for the users.

[29] Actually, this is not the full truth as persons are able to organize themselves out of liability in limited liability companies.

[30] See Zetzsche et al. (2017) for a discussion.

4 Participants' causal link to harms

Responsibility is about the duty or ability to take action – to respond – and about the consequences of failing to do so. Causation is a central element in the assessment of both legal and moral responsibility for an action. It is hard to be responsible for something that is out of any sort of causal control. In this section, we will seek to determine to what extent the actions by the participants in a cryptocurrency scheme can be considered a cause for a harm from the scheme. In other words, we will investigate under what causal concept participants' actions can be considered a cause for a harm. In the legal literature, causal analysis is often separated into "cause-in-fact" and "legal cause".[31] Cause-in-fact shares methodology with causality concepts in sciences and philosophy and seeks to determine what can be considered a cause for an event (not "the" cause, as an event may have many causes). Legal cause are refinements and specifications to make causes sufficient for legal liability. Legal cause will be discussed further in Section 5, while the topic in this section will be cause-in-fact.

The cause-in-fact question is whether an event conceptually is the cause of some other event.[32] The starting point for considering en event as a cause for another event is usually the "but-for" counterfactual analysis. The question is if some harm would have happened without the presence of the action. For some harms associated with cryptocurrency operations, such a causal link exists. If a particular validator causes environmental harm, this harm would not exist but for this validator. However, normally, such a requirement would rule out the actions of most participants at the network level of a public cryptocurrency, as causes of harm. The essence of decentralized operations is that no participants are necessary for correct validation, and that no single participant is necessary for a transaction to be broadcasted to the network. Consequently, neither validating nor transmitting nodes can be considered causally responsible for illegal transactions to be broadcasted to the network and entered into the ledger.

However, this might change in a permission-based system. If there is voting according to BFT consensus mechanism, one might end up in a X+1 versus X vote to validate a particular transaction. In this case, each validator on the winning side was necessary for the outcome. The necessity of other participants at the network level for the outcome may be more crucial. For instance, oracles feeding external information into a DLT may be more centralized than validators. If there is voting, there is a risk of having a X+1 versus X vote, similar to validators in a permission-based system, making each voter necessary for the outcome.

At the protocol level, there seem to be more obvious candidates for persons being necessary for harm produced by a cryptocurrency. If a protocol's governance mechanism design fails and causes harm, the protocol developers may seem necessary for this harm to occur. If there are errors in the software implementation of a protocol, the software developers appear as necessary for the harm. If an illegal transaction is completed by a privacy-centric cryptocurrency such as Zcash, the developers of Zcash may seem necessary for the existence of this transaction. However, there are also obstacles singling out individuals as necessary for the harm in these contexts. The development of cryptocurrencies is often performed in community projects based on open source code development. No single participants in this community may appear as necessary for a particular protocol and software design. This might be different for permission-based cryptocurrencies and for arrangements not based on open source. Also, in some cases, protocol developments are organized within organizations as described in Section 2, which may be considered responsible as an entity.

A general obstacle for considering any participants in a cryptocurrency necessary for harm is that a single cryptocurrency is just one among many that could have been used to produce the same harm, such

[31] See Moore (2009).

[32] For a thorough discussion of causality concepts, see Illari and Russo (2014).

as the production of an illegal transaction. There are, for instance, many privacy-oriented cryptocurrencies. If one was not available, another could easily substitute.

The organization of cryptocurrencies is an obstacle to causally link individual actions to harm according to a "but-for" counterfactual test. However, other causal concepts can overcome this obstacle. Weaker causal links than the "but-for" requirement are plentiful and are also not unfamiliar in legal contexts. A concept called INUS has been developed with legal applications in mind.[33] INUS is an abbreviation for Insufficient, but Non-redundant part of an Unnecessary but Sufficient condition. Put simply, if x is a non-redundant part of a set, X, sufficient for Y to occur, then x is a cause according to the INUS concept. In such a case, if a participant's action was part of a set of conditions sufficient for the harm, the participant cannot get away with claiming that the action was not necessary for the harm due to the assumption that the action would have been performed by some other participant, as long as the participant actually was a part of the production of the harm.

Causal concepts might also be derived from inquiring mechanisms. By using a mechanism concept, a participant's action, such as a validator's validation, the causal role might be apparent because it is included in the mechanism that led to the harm. The logic would be that the validator validated the illegal transaction and, hence, the validator is a part of the mechanism leading to the transaction being validated. Such a mechanism approach may expand the causal role of validators and other participants substantially compared to counterfactual analysis. If we consider a cryptocurrency as an ecosystem, many participants are contributing to its security. If one validator were not present, the system would be little less secure. Hence, each validator contributes to the security of the ledger and the (probabilistic) immutably of a transaction. In a privacy-centric cryptocurrency scheme, all users (potentially) contribute to increasing the privacy of a coin-mixing arrangement.

A cause might be probabilistic or deterministic. Probabilistic causes may be particularly applicable for establishing a causal link between validators' actions and the outcome in competitive validation. Although there is nothing intrinsically stochastic in who validates a block, the competition is so complex that we might consider it random. In a PoW validation scheme, a particular validator's probability of validating a block ex ante is a function of its computing power. Hence, any active validator could be causally linked to the validation by probability. This is closely related to statistical general causes versus individual causes. There might be a statistical tendency for rain to make the streets wet, but it may not be the cause of the wetness of a particular street (for instance, because it was washed just before the rain). By the same token, there might be a statistical relationship (not just a theoretical probability-model argued relationship) between the computational capacity of a validator and the likelihood of validation. An objection to such a general ex ante assessment of causality based on probability models or statistical analyses is that what matters is the actual cause. For instance, for the actual cause, it matters who actually validated, and not who could have validated according to a probabilistic formula or a general statistical relationship. We will return to this issue in the analysis of legal liability below. However, it is worth mentioning that more advanced probability calculus or statistical models may be invoked to get closer to the actual cause at the individual level.[34]

Capacity as causality is a concept based on very weak causal links. According to a capacity concept, an event can be considered a cause if it has the capacity to produce the harm. Hence, a participant can be seen as a cause of harm by dint of having the capacity to produce the harm according to other causal concepts. Hence, this is a weakening of other causal concepts (which may already be weak).

It should also be noted that in situations of multiple causes, causes can be weighted according to suitable criteria, where some causes are weighted heavier than others. In a cryptocurrency ecosystem, a

[33] Developed by Mackie (1974).

[34] See Halpern (2016) and Pearl et al. (2016), respectively.

large validator will typically have a larger causal weight in keeping the system secure than a smaller one. A participant controlling many wallets in a privacy-oriented cryptocurrency has a larger causal weight in the facilitation of anonymous transactions. Such weights and their construction may be useful for certain applications such as determining responsibility and liability. Typically, responsibility would require a weight above a certain threshold.

There is arguably an abundance of causal concepts to link the various participants of a cryptocurrency to cryptocurrency harm. However, this does not necessarily mean that that such causal links are sufficient for legal responsibility and moral responsibility. This will be discussed next.

5 Causal links and legal responsibility

The "legal cause" requirement specifies those causes that are legally liable. Such liability may follow from general criminal law or tort law or from specialized regulations. In this section, we will first look at those general requirements to causality required for legal liability established by practice in criminal law and tort law, which are also influential when assessing liability in specialized regulations. However, in specialized regulations, the legislator may be more creative when it comes to causal requirements for liability. This will be discussed after the general discussion.

Legal liability often depends on the faultiness of the person causing the harm. In some cases, intent or recklessness is required, while in negligence is sufficient in others. In some cases, there is strict liability, where only causality is needed for liability. Criminal liability often requires at least recklessness for liability. Hence, it may be crucial to elucidate the distinction between recklessness and negligence. This distinction might be particular crucial to draw for validators and users propagating transactions in coin-mixing schemes to facilitate privacy. If a privacy-centric cryptocurrency is proven to mostly facilitate criminal transactions, it might be hard for users and validators to argue that it is not reckless to participate in the scheme.

Strict liability is best known in product liability cases and cases involving the use of dangerous things, such as vehicles.[35] Strict liability might be applicable in the case of cryptocurrencies and DAOs capable of causing severe harm. However, protocol- and code-developers are more likely to be subject to the usual negligence standard. As argued by Choi (2018) for algorithms in general, this might be implemented by a requirement of crashworthy code, which means that the code must properly respond to a fault. The code behind the The DAO organization described in Section 2 may be considered negligent. The developers behind it were saved by a hard fork, but it is likely that victims would have taken legal actions for damages otherwise. Note that the standard may be higher when it comes to internal harm in a cryptocurrency ecosystem than harm outside. Within a cryptocurrency ecosystem, the participants can be considered to have implicit contractual obligations towards each other, implicating mutual loyalty duties.[36] In particular, protocol developers may be considered to have fiduciary duties towards other stakeholders.[37] In these cases, the threshold for holding other participants liable for harm may be lower.

Another legal cause requirement is proximity, which means a sufficient closeness in space and time between the cause and the consequence.[38] Such a requirement of proximity is likely to rule out liability for participants in a cryptocurrency distant from the harm in these dimensions. Proximity is closely re-

[35] This started out as a doctrine for liability for dangerous animals. Strict liability has a law and economics justification, as strict liability is a way to make people take into account the full cost of their activities and by this adjust their activity level. More on law and economics below.

[36] Zetzsche et al. (2017)

[37] See Walch (2018) and Seira (2018)

[38] For some claims there are statutory limitation rules.

lated to foreseeability as a requirement for legal liability for caused harm. For instance, for the developers of Bitcoin, its spectacular evolution, and many of the possible harms from the scheme, were probably neither foreseeable nor proximate, although they may be considered as a cause for these harms.

A principle originating from economic analysis of law is the least-cost avoider principle. The economic idea is that the liability for a risk should be carried by whoever can mitigate the risk at least cost. This provide incentives to effectively mitigate risk. Assume that a possible harm of $100 can be prevented by person A at the cost of $ 50 and by person B at the cost of $ 1. By placing the liability on B, B is incentivized to prevent the harm at the least cost. In more advanced forms with multiple participants, liability could ideally be distributed such that each participant is incentivized to take efficient care. In reality, it is difficult for courts and other allocators of liability to do the exact calculations assumed in economic models, but the least cost avoider principle may be approximated by various other principles. For instance, a harmed person who by negligence contributes to the presence or magnitude of harm may share the liability and in some cases remove the liability for the one who caused the harm.

The assignment of liability according to the least-cost avoider principle may be crucial for the operation of a cryptocurrency. This is particularly important if one also recognizes the benefits of a cryptocurrency. If liability is wrongly allocated, such benefits may be prevented.[39] This can be illustrated with the facilitation of illegal transactions by a privacy-centric cryptocurrency. An enforceable liability on every participating node for the harm from an illegal transaction would effectively shut down the cryptocurrency. This would not be productive if one recognizes that the harm could be prevented more efficiently by other participants. Such participants may be validators or oracles (for instance, the protocol could give oracles the authority to block transactions used for ransom). Only if normal users choose to participate in cryptocurrency schemes not entailing such governance mechanisms in their protocols may they be the least-cost avoiders of harm, and be liable accordingly. If some participants imprudently create a smart-contract, the least-cost avoiders would normally be the designers of the smart contracts and not the designers of the cryptocurrency, as long as the cryptocurrency is prudently designed. By the same reasoning, it would be wrong to place liability on validators for validating faulty smart contacts. This might, however, change if the smart-contract is already running and the prevention is out of the control of the originators of the smart contracts. We will return to this issue in Section 6.

So far, we have seen how general principles narrow and specify liable causal links. The law can also broaden liability where causal links are weak or narrow the liability when causal links are are strong to promote some public policy. If the law broadens liability, this can either be seen as allowing for liability in the presence of weak causal links or be considered as non-causal liability. Although the former is perhaps more descriptive, the latter term is often used to describe such liability. One way of expanding liability in the presence of weak causal links, is to consider contribution as liable. In criminal law, this is often known as “accomplice” and “aiding and abetting”. The typical example is the driver taking the criminal to the crime scene. This principle may seem to serve well for liability in cryptocurrencies. For instance, being a validator means to contribute to the network security even though they are not doing the actual validation. Being a node in a privacy-centric cryptocurrency is another example. Being such a node contributes to blurring transactions whether or not this particular node is actually involved in the mixing of a particular transaction.

Another related way of establishing liability for activities of complex organizations where the causal link for each participant to harm is weak (or when a causal link is hard to document) is to assign liability to a group, where the group itself is liable or each member is liable for the actions of the group based on group membership alone. A prime example of assigning liability to a group is corporate liability, where

[39] By this, I don’t take the position that persons should not be liable for harm as a means to spur innovation that characterizes the current legal system, most generally by the limited liability regime. I only take the position that liability should be effectively allocated.

the corporation is assigned legal personhood and can be liable on its own.[40] Some cryptocurrencies, in particular those that are permission-based, may fit into a corporate structure for corporate liability. Also, the initial development of many public cryptocurrencies is organized within corporate structures, which can be held liable in case of harm. By the same token the development of many cryptocurrencies are organized within foundations, such as the Ethereum Foundation, which can be hold liable as a unit. For many cryptocurrencies validation and the adminsitration of mining-pools is performed by companies that can be held liable.

However, group liability via an identifiable unit assigned legal personhood seems inadequate for addressing harm from many cryptocurrencies.[41] It seems more viable to make group members individually liable for the group's actions based on participation. Some general instruments exists in law, such as the principle of joint liability in tort law. Under joint liability, each tortfeasor is liable up to the full amount. This means that someone who has contributed "a little" to some harm still can be liable for the full amount. In some cases, joint liability is the default if a joint venture is set up. Some permission-based cryptocurrencies are likely to fit into this category, rendering certain participants jointly liable for each other's actions. Joint liability can also be established by specific regulations, such as joint liability for certain violations of the antitrust laws. Specific regulation will be discussed next.

Legislators may also choose to use special regulations to establish liability for conduct with a sufficient general causal link to harm without the need to establish a causal link at the individual level. Such a regulation may even be invoked on the basis of capacity. The idea is that actual harm from the individual conduct is so likely that it is inefficient to investigate the details of every individual case. In this case legislators only have to establish general causality between the conduct and the harm to justify the regulation.[42] Such regulations would typically be to prohibit certain conduct and group membership. For instance, if cryptocurrency activity in general is considered harmful, any activity involving cryptocurrencies could be prohibited.[43] Rules may be more specific such as prohibiting involvement with privacy-centric cryptocurrencies, public cryptocurrencies, PoW-based cryptocurrencies, or non-licensed cryptocurrencies. Several jurisdictions have implemented, or are considering, various prohibitions in the cryptocurrency sphere.

6 Causal links and moral responsibility

There should be some correspondence between the perception of moral blameworthiness and liability. The legitimacy of a legal system is likely to be undermined if someone perceived as morally blameworthy

[40] Assigning liability to a corporation makes the shareholders of the corporation as a group de facto liable for the actions of the administration of the corporation. There is a causal link here, as the shareholders are supposed to exert governance on the administration, but also minority shareholders with no influence are indirectly made liable. Corporate liability, however, also serves as an instrument for narrowing liability even when causal links are strong, because corporations often benefit from limited liability regimes effectively transferring the harm caused to the victims. Limited liability is intended as an instrument to spur innovation and progress by capping the liability of the participants. The discussion as to whether this is an instrument beneficial to society, or another clever way of increasing the wealth of the already wealthy, is beyond the scope of this paper. However, those concerned with the lack of liability for participants in a cryptocurrency are advised to compare it to the lack of liability in already-existing organizational constructs.

[41] It has been discussed in the literature as to whether artificially intelligent autonomous algorithms, including decentralized autonomous organizations could achieve legal personhood and be held liable for their actions. At the moment, this is science-fiction and will not be discussed further here. Although interesting from a philosophical point of view, this is not an available option for allocating responsibility for harm from cryptocurrency schemes for now.

[42] This is assuming a modern democratic legal system where legislation needs justification. In authoritarian regimes, the legislators may not need to justify laws or may not need a legislative basis for interference at all.

[43] This would probably be a very inconsiderate rule.

escapes liability due to constraints in what is considered legal cause, especially if this happens on a grand scale. For instance, someone escaping caused harm based on the least-cost avoider principle may be contrary to moral standards. If a smart contract is poorly programmed and creates harm if executed, it may appear contrary to moral standards if validators with the ability to prevent or mitigate severe harm from the smart contract escape liability based on an application of the least cost-avoider principle rendering the smart-contract designers as the least-cost avoiders.

Although some general principles can be applied, the principles for assessing the moral responsibility of cryptocurrency participants are not well developed. There is probably little consensus on such principles. While Satoshi Nakamoto may qualify for a saint status for the creation of Bitcoin to those who have been exploited and have a grudge against the existing financial system, extortion victims required to pay bitcoins or other cryptocurrencies in ransom may have another opinion. Moral responsibility for cryptocurrency participation is an interesting topic that moral philosophers hopefully will pick up and develop. Principles for moral responsibility for group and crowd actions will probably be useful in this context.[44] The moral debate in general will be beyond the scope of this paper. However, a brief discussion based on more-or-less objective scientific approaches to moral responsibility for caused harm by cryptocurrency participants is provided here.

Lagnado et al. (2013), with the help of a modeling approach elaborated in Halpern (2016), divides causal responsibility into criticality and pivotality in an empirical experiment.[45] Criticality is about how important the agent is perceived to be for the outcome, before any actions are taken.[46] Pivotality is about whether an agent made a difference to the outcome.[47] Hence, criticality is forward-looking, and pivotality is backward-looking. Both are important for how agents morally assign responsibility. It is easy to see that many participants in a cryptocurrency scheme, especially at the network level, score low on criticality. This is the essence of the decentralized nature of cryptocurrencies. However, hey are likely to score higher on pivotality. Especially in permission-based cryptocurrencies, single participants may be crucial, as discussed in Section 4. When we look at the protocol layer, protocol developers may score high on criticality, but maybe less for pivotality. They are responsible for the cryptocurrency coming into existence and as to whether the protocol provide governance mechanisms to prevent harmful transactions, but are not much in control to distinguish harmful from legitimate transactions in the actual operations. However, if the harm is caused by a flaw in the protocol design, the score might be high on both criticality and pivotality. Also, protocol developers that are also participating at the network layer, as in certain permission-based cryptocurrencies, may score high on both criticality and pivotality and, hence, be more blameworthy for harm.

Moral responsibility is likely to be much more complex than can be represented by models. However, since the cryptocurrencies themselves are heavily based on game-theoretical models in their governance structures, models of responsibility may prove particularly useful.

7 Concluding remarks

Although cryptocurrencies are novel innovations with many potential benefits, they can cause harm. Despite of the decentralized nature of cryptocurrencies, participants cannot conceptually, legally, or morally

[44] This is often referred to as collective responsibility in the philosophy of moral responsibility. A related issue is the problem of many hands, which characterize situations where a group hand be hold responsible despite no responsible individuals; see Van de Poel (2015). There is also an emerging literature on moral responsibility for online platforms; see Helberger et al. (2018).

[45] Lagnado et al. (2013) traces these two concepts also to the legal literature.

[46] See Lagnado et al. (2013).

[47] See Lagnado et al. (2013).

use lack of causal links as a justification to evade responsibility for such harm. There is an abundance of causal concepts available to link participant actions to harm to be applied in establishing legal and moral responsibility. With the benefits of avoiding trusted third parties come responsibilities for the participants otherwise absorbed by the trusted third parties. A single participant may easily become causally responsible for the harm produced by the whole scheme.

References

Ali, R., Barrdear, J., Clews, R., & Southgate, J. (2014). The economics of digital currencies. Bank of England Quarterly Bulletin, 54(3), 276-286.

Al-Naji, N., Chen, J. & Diao, L. (2018), Basis: A Price-Stable Cryptocurrency with an Algorithmic Central Bank.

Antonopoulos, A. M. (2017). Mastering Bitcoin: Programming the Open Blockchain. O'Reilly Media, Inc.

Bahn, D., & Dressel, D. (2006, October). Liability and control risks with open source software. In 2006 International Conference on Information Technology: Research and Education (pp. 242-245). IEEE.

Bentov, I., Gabizon, A., & Mizrahi, A. (2016, February). Cryptocurrencies without proof of work. In International Conference on Financial Cryptography and Data Security (pp. 142-157). Springer, Berlin, Heidelberg.

Blockchain Luxembourg, S. A. (2018). The State of Stablecoins.

Carbonara, E., Guerra, A., & Parisi, F. (2016). Sharing Residual Liability: The Cheapest Cost Avoider Revisited. The Journal of Legal Studies, 45(1), 173-201

Choi, B. H. (2018). Crashworthy Code. Washington Law Review, Forthcoming.

Cloakteam (2018). Enigma v2.1 A private, secure and untraceable transaction system for CloakCoin.

Conti, M., Lal, C. & Ruj, S. (2017). A Survey on Security and Privacy Issues of Bitcoin. arXiv preprint arXiv:1706.00916.

Dai (2017). The Dai Stablecoin System, December 2017, https://makerdao.com/whitepaper/DaiDec17WP.pdf

Duffield, E.(2017). Dash v13 Evolution Design Overview Rev 1.

Duffield, E., & Diaz, D. (2014). Dash: A privacy-centric crypto-currency.

Goertzel, B., Giacomelli, S., Hanson, D., Pennachin, C., & Argentieri, M. (2017). SingularityNET: A decentralized, open market and inter-network for AIs. https://public.singularitynet.io/whitepaper.pdf

Halpern, J. Y. (2016). Actual causality. MiT Press.

Helberger, N., Pierson, J., & Poell, T. (2018). Governing online platforms: From contested to cooperative responsibility. The information society, 34(1), 1-14.

Illari, P., & Russo, F. (2014). Causality: Philosophical theory meets scientific practice. OUP Oxford.

Lagnado, D. A., Gerstenberg, T., & Zultan, R. I. (2013). Causal responsibility and counterfactuals. Cognitive science, 37(6), 1036-1073.

Mackie, J. L. (1974). The cement of the universe: A study of causation. Oxford University Press.

Moore, M. S. (2009). Causation and responsibility: An essay in law, morals, and metaphysics. Oxford University Press on Demand.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.

Narayanan, A., Bonneau, J., Felten, E., Miller, A., & Goldfeder, S. (2016). Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction. Princeton University Press.

Popov, S. (2017). The Tangle, October 1, 2017 Version 1.3, https://iota.org/IOTA_Whitepaper.pdf

Pearl, J., Glymour, M., & Jewell, N. P. (2016). Causal inference in statistics: a primer. John Wiley & Sons.

Rauchs, M., Glidden, A., Gordon, B., Pieters, G. C., Recanatini, M., Rostand, F., ... & Zhang, B. Z. (2018). Distributed Ledger Technology Systems: A Conceptual Framework.

Rauch, M., Blandin, A., Klein, K., Pieters G., Recanatini, M. & Zhang, B. (2018b), 2nd Global Cryptoasset Benchmarking Study, Cambridge Centre for Alternative Finance.

Sasson, E. B., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., & Virza, M. (2014, May). Zerocash: Decentralized anonymous payments from bitcoin. In Security and Privacy (SP), 2014 IEEE Symposium on (pp. 459-474). IEEE

Schäfer, H. B., & Müller-Langer, F. (2009). Strict liability versus negligence. Tort law and economics, (1991), 3-45.

Seira, R. (2018), Blockchain Protocol Developers are not Fiduciaries: An Analysis of the Cryptoeconomics of Open Source Networks and the Role of Protocol Developers in Public Blockchain Network Governance

Tjong Tjin Tai, E. (2018). Liability for (Semi) Autonomous Systems: Robots and Algorithms.

de Vries, A. (2018). Bitcoin's Growing Energy Problem. Joule, 2(5), 801-805.

Van de Poel, I., Royakkers, L., & Zwart, S. D. (2015). Moral responsibility and the problem of many hands. Routledge.

Walch, A. (2018). In Code (rs) We Trust: Software Developers as Fiduciaries in Public Blockchains.

Zetzsche, D. A., Buckley, R. P. & Arner, D. W. (2017). The Distributed Liability of Distributed Ledgers: Legal Risks of Blockchain (August 13, 2017). University of Illinois Law Review, 2017-2018

Østbye, P. (2018). The Case for a 21 Million Bitcoin Conspiracy. Available at SSRN: https://ssrn.com/abstract=3136044 or http://dx.doi.org/10.2139/ssrn.3136044

Østbye, P. (2018b). Will Regulation Change Cryptocurrency Protocols? Available at SSRN: https://ssrn.com/abstract=3159479 or http://dx.doi.org/10.2139/ssrn.3159479

Østbye, P. (2018c). Model Risk in Cryptocurrency Governance Reliability Assessments. Available at SSRN: https://ssrn.com/abstract=3195686